SPATIAL MODE CONTROL AND ADVANCED METHODS FOR MULTI-PLATFORM QUANTUM COMMUNICATION

BY

ALEXANDER D. HILL

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2018

Urbana, Illinois

Doctoral Committee:

Assistant Professor Virginia O. Lorenz, Chair
Professor Paul G. Kwiat, Director of Research
Professor Jen-Chieh Peng
Professor Emeritus John Stack

Abstract

Though state-of-the-art quantum computers are currently limited to only a handful of physical qubits, a quantum computer large enough to perform prime factorization of modern cryptographic keys, quantum simulation, and quantum-enhanced searching algorithms will likely become viable within a few decades. Such computers demand communication networks that preserve the qualities of the quantum states used as inputs and outputs; they also herald the end of the flavors of classical cryptography reliant on the complexity of factoring large numbers. As a result, future networks must include channels which preserve the states of single photons over useful distances (e.g., using quantum repeaters), and must deploy quantum-safe cryptography to ensure the safety of classical information passing over the network.

Here we discuss strategies affecting several areas of a future quantum-enabled network: first, we demonstrate a technique for adaptively coupling single photons from point sources into single-mode optical fiber and apply the technique to coupling from quantum dots (a popular candidate for a future quantum repeater); secondly, we discuss various methods for simulating the effects of atmospheric turbulence on quantum cryptographic protocols in the laboratory, critical for understand the challenges facing free-space implementations of quantum communication. Thirdly, we demonstrate a technique that enables quantum cryptographic networks over free space channels to function in the presence of strong atmospheric turbulence using a multi-aperture receiver. Finally, we discuss our efforts to miniaturize a quantum key distribution system and operate a key distribution channel between flying multirotor drones.

To my parents

Acknowledgments

This thesis is the crystallization of the work of so many people, work that I have been given the honor of presenting as my own. Here I hope to thank a terribly incomplete subset of them. A great deal of thanks is owed to our collaborators at Duke University and The Ohio State University: Taimur Islam, Joseph Szabo, and their advisor Prof. Daniel Gauthier, with whom I have worked from afar for almost my entire tenure as a graduate student. A warm thank you as well to our collaborators at the University of Michigan, Alexander Burgers and Uttam Paudel, and their advisor, Prof. Duncan Steel, for their hospitality, advice, and support in applying our adaptive optical ideas to their quantum dots. I am grateful to all of my labmates, especially Brad Christensen and Trent Graham, for their guidance, and Dalton Chaffee, for his support. Thank you as well to Prof. Mike Goggin for his help and guidance. Bo Liu of the Intelligent Robotics Laboratory was absolutely invaluable in providing guidance and wholesale help when we struggled to make sense of our drones. Thank you to Prof. Lance Cooper for his unrelenting helpfulness, and to all of the University of Illinois Physics staff, especially Betsy Greifenkamp, Cheryl Sabas, and Wendy Wimmer, who have each rescued me multiple times using their own forms of administrative sorcery. Many of my projects would have been impossible without the help of the departmental machine shop, Ernie Northen in particular. I am deeply grateful to have worked with my advisor, Prof. Paul Kwiat, who has pushed me to succeed academically and opened so many doors for me. A final thank you to my parents, who have been fielding questions from extended family about what I have been doing with myself for so many years with unwavering support.

This work was supported by the National Defense Science and Engineering Graduate Fellowship (NDSEG), the U.S. Army Research Office MURI award #W911NF0910406, and the Office of Naval Research MURI on Fundamental Research on Wavelength-Agile High-Rate Quantum Key Distribution in a Marine Environment, award #N00014-13-0627.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

QKD	Quantum Key Distribution.
HWP	Half-wave Plate
QWP	Quarter-wave plate
TT	Time Tagging Unit
SPAD	Single-photon Avalanche Photodiode
SWAP	Size, Weight, and Power
QBER	Quantum Bit Error Rate
PAT	Pointing and Tracking
UAS	Unmanned Aerial System (drone)
BB84	1984 Bennett-Brassard Quantum Key Distribution Protocol

Chapter 1

Introduction

1.1 Future-proof Communication

Quantum computers are potentially only a short time from being able to break the most popular forms of encryption [1]. Recent years have seen an explosion in investment in quantum computing research (Google [2], IBM [3], and Microsoft [4] all have serious quantum computing investments), a trend driven by the promise of super-classical search [5], polynomial-time cracking of cryptographic keys [6], quantum simulation of complex many-body systems [7, 8], and quantum-enhanced optimization solvers [9]. This new model of computing demands a new model of communication, one capable of preserving and securing the information traveling to and from a quantum computer and resistant to the types of encryption attacks where quantum computers are known to excel [10].

Classical optical communication assumes a few basic features of a communication channel that do not apply in a potential quantum-enabled future: that states of light may be measured and amplified (low-noise repeaters), and that the information transmitted over the channel is secured at an algorithmic level, even if intercepted by an adversary. The former assumption is the question of physics, while the latter is the question of cryptography. The ability to amplify arbitrary bright classical states of light is an unassailable advantage of classical communication over quantum communication; however, future quantum computers should not be connected with classical light, as an adversary can ascertain information encoded in classical light sources using simple tomographic techniques. Quantum communication, in contrast, allows both for interaction with quantum systems, e.g., quantum computers, and enhanced message security, e.g., from quantum key distribution, by encoding information into the quantum state of single particles or the joint states of multiple particles. This differs in one important respect from the classical case in that the states of the particles are not clonable [11, 12]; Thus, although information-bearing particles cannot be amplified to improve transmission, the states encoded in each particle can be rendered safe from undetectable tampering or measurement.

A sketch of a future quantum-enabled network is shown in Figure. 1.1. A network of users wishes to

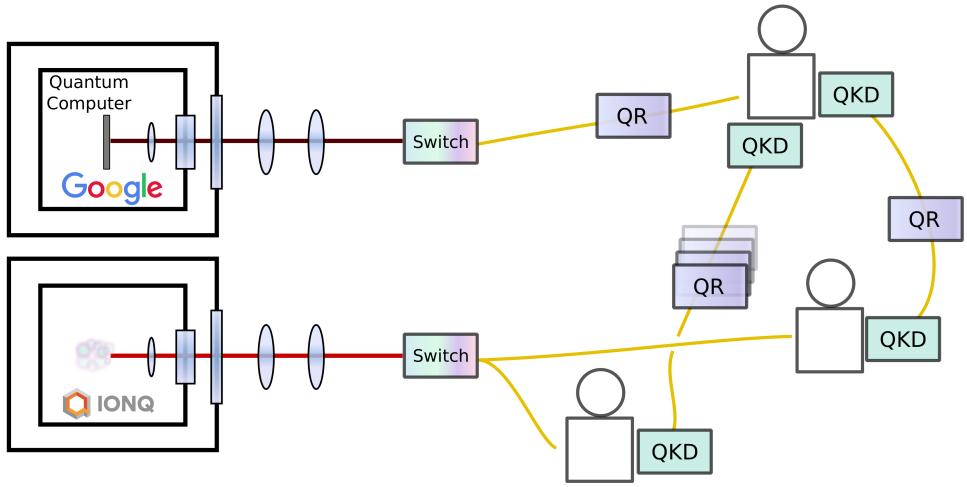

Figure 1.1: Future quantum-enabled networks will require a number of new technologies, such as quantum repeaters (QR), quantum key distribution (QKD) channels, and quantum computers.

exchange information with quantum computers of various flavors. They can encode their instructions for the computers as classical information, which may then be interpreted as states and gates for the purposes of computation. The users might also wish to send information encoded into the states of individual particles, which would necessarily be photons, the "flying qubit." Being able to shuttle users' states securely into and out of quantum computers is useful for blind quantum computing (BQC) protocols [13], where a possibly malicious quantum computer is able to perform a desired computation without being able to infer any information about the nature of the computation itself.

Users may also want to communicate with one another using a classical communication network featuring provable security. As we discuss later, the only truly secure form of encryption is the "one-time pad," and quantum key distribution (QKD) allows two parties to share a one-time pad in a "provably secure" manner. QKD exploits the no-cloning property of quantum states to share random cryptographic key bits between two parties without the possibility of undetected tampering by an adversary.

Establishing and maintaining a communication link capable of transmitting photons without corrupting the information they carry is a tricky task, especially in cases where high-end optical telecom infrastructure cannot be used (e.g., between airplanes). Real-world non-idealities frequently reduce the quality and security of quantum cryptographic links, thereby reducing their usefulness. Here we outline the difficulties with implementing enabling quantum-enhanced networking, specifically quantum cryptography, in a variety of useful situations. We begin with a brief introduction to quantum cryptography, then discuss how these protocols break down in practice.

1.1.1 Quantum Cryptography

Cryptography is the keystone of modern communications – virtually all sensitive communications, from bank statements to social media login data, are transmitted and secured secretly and electronically. Current public-key encryption schemes rely on the difficulty of prime-factoring large numbers; the most commonly used schemes, such as RSA, may be broken in finite time, but long keys may be impossible to factor in a reasonable amount of time using classical computers. Thus, the security of schemes like RSA relies solely on limitations in cracking algorithms and hardware speed.

However, in 1994, Peter Shor showed that computers could exploit quantum mechanics to implement quantum algorithms capable of prime-factoring integers in polynomial time [6]. Given that several platforms are approaching the quantum error-correction threshold for computation, and that quantum systems are being manipulated on larger and larger scales, it is not unreasonable to believe quantum computers in some form may be viable within our lifetimes. Government and civilian agencies alike have demonstrated the will and financial ability to produce novel large-scale cryptographic attacks [14]; we should expect the same would be true of any feasible quantum computer, regardless of cost. The advent of quantum computation could thus mark the end of standard public-key cryptography.

Several post-quantum schemes exist which are potentially robust against quantum-enabled attacks [15]. The only perfectly secure cryptographic technique is the one-time pad, where the cryptographic key, e.g., a random binary string, is as long as the secret message and is discarded after one use [16]. A typical algorithm relies on transmitting the XOR of the key (the one-time pad) and the target message [17]. The resulting encrypted message is indistinguishable from noise and renders all messages of the secret message's length equally probable. A significant drawback of the one-time pad is that both parties must know all encryption keys in advance; they must not reuse the key. Transmission of a new one-time pad over classical communication is not useful, as encrypting new keys would deplete all current keys.

Quantum Cryptography, more accurately called Quantum Key Distribution (QKD), solves this problem by allowing separated parties to generate a shared secret one-time pad (by assuming they start with a small amount of secret shared key to enable "authentication," QKD is sometimes described as a secret key *expansion* protocol). As discussed below, the no-cloning theorem of quantum mechanics forbids any eavesdropper from intercepting information shared between parties using single particles without being detected.

1.2 Quantum Cryptographic Protocols

There are a variety of protocols for quantum key distribution, but here we restrict the discussion to the most common and relevant scheme. The first and most straightforward quantum key distribution protocol – often referred to as BB84 – was proposed by Bennett and Brassard in 1984 [18]. Two parties, typically referred to as Alice and Bob, attempt to distribute a random cryptographic key using a 2-dimensional state space of single photons. For example, Alice prepares photons in a polarization randomly chosen from

$$
\begin{aligned}
|H\rangle &= |\rightarrow\rangle \\
|V\rangle &= |\uparrow\rangle \\
|D\rangle &= |\nearrow\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle + |V\rangle\right) \\
|A\rangle &= |\searrow\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle - |V\rangle\right),
\end{aligned}
$$

and sends them one at a time to Bob. H/V and D/A form two separate bases of orthogonal polarizations; because the projections of each basis element onto the elements of the conjugate basis are $1/\sqrt{2}$, these two bases are called *mutually unbiased*. Bob randomly chooses one of these bases for measurement; if he chooses the incorrect basis, he has a 50% probability of measuring either element of the basis. If he chooses the correct basis, his measurement should match what Alice prepared. Alice and Bob publicly compare the bases in which they prepared or measured each photon and discard measurements for which they used different bases. They then map the elements of each basis to a bitwise value to form a binary key (e.g., H→0, V→1, D→0, A→1), called the *sifted key* (i.e., the key after mismatched-basis measurements have discarded). For example,

Alice Sends	A	D	A	H	A	V	H
Bob's Basis	D/A	H/V	H/V	D/A	D/A	H/V	H/V
Bob Measures	A	H	V	A	A	V	H
Correct Basis?	✓	×	×	×	✓	✓	✓

Raw Key: AAVH → 1110

Because the only public information available to an eavesdropper (sometimes called 'Eve') is the basis choice, she cannot determine any of the key from the basis announcement, which comes only after the photons have been detected by Bob. Furthermore, if Eve attempts to intercept a photon, measure the polarization, and send a copy to Bob, she will incorrectly select the measurement basis 50% of the time. In this case, the

raw key shared between Alice and Bob will differ in 25% of the bits. Alice and Bob detect these errors by publicly comparing a subset of their raw key; a large bit error rate alerts Alice and Bob to the presence of an eavesdropper, and they can discard the entire key. For smaller error rates, Alice and Bob can perform classical error correction to obtain an identical key and "privacy amplification" to eliminate any information accessible to Eve. As a result of these procedures, the final, secure key will be shorter than the original sifted key (maybe much shorter – or even zero – if there is a high error rate in the raw key).

This form of BB84, where Alice prepares single photons in a state of her choosing and Bob later makes some sort of measurement on the state of that photon, is known as a *prepare-and-measure* scheme. Prepare-and-measure protocols typically rely on preparing coherent pulses with a poisson-distributed number of photons per pulse: $|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle$, where n is the number of photons in a pulse and α is the mean photon number. In this case, every photon in the pulse has the same polarization (or time mode, or frequency mode, etc.) as every other; therefore, an eavesdropper can perform a simple attack (called the photon-number splitting attack) wherein she attempts to remove at most one photon from the state, e.g., using a beamsplitter, and send the remaining photons along to Bob. The photons in the state Bob receives will not have been perturbed by the eavesdropper's measurements, so she can gain information about the state without Alice and Bob noticing (the eavesdropper is assumed to be able to teleport the untouched photons to Bob with arbitrarily low loss; in this way, she can remain undetected as long as the bypassed channel transmission from Alice to Bob is equivalent to her measurement success probability). As it is impossible to know which pulses have photon numbers of over 1 photon per pulse (it is undefined), Alice and Bob are forced to use extremely low mean photon numbers per pulse to reduce the probability that multiple photons are produced.

A solution to this problem is known as the decoy-state method [19]. In this protocol, multiple states of known mean-photon numbers are produced by Alice and sent to Bob. Though neither Alice nor Bob can know the number of photons in any given pulse, they do know the statistics of each. If an eavesdropper attempts a photon-number splitting attack, the relative success probabilities of each decoy state will be different than for the pure-loss channel, as the eavesdropper can only manipulate a single statistical moment of the photon number distribution – the mean – but not any higher-order moments. This allows Alice and Bob to bound the probability of an eavesdropper attempting to siphon photons from the channel, and enables them to operate using much higher mean photon numbers. In fact, using decoy state methods can restore the speeds of secure key generation to approach the pure single-photon version of BB84. As such, for modern prepare-and-measure schemes, decoy states are a requirement.

The BB84 protocol may be implemented using entangled pairs of photons, for example, pairs in the joint

state $|\Psi\rangle = (1/\sqrt{2})[|HH\rangle + |VV\rangle]$. This state is equivalent to the state $|\Psi\rangle = (1/\sqrt{2})[|DD\rangle + |AA\rangle]$. If Alice and Bob each share one photon of the entangled pair and make projective polarization measurements on the state of their photon, the results of their measurements will be perfectly correlated when they measure in the same basis. Entanglement-based BB84 is formally identical to the prepare-and-measure version [20]. Such protocols can also leverage Bell's theorem [21] for a further security check against tampering by an eavesdropper, as projective measurements destroy the two-photon entangled state and prevent a violation of Bell's inequality [22].

Entanglement-based protocols, while perhaps more complicated to implement, offer a distinct advantage over prepare-and-measure protocols: chiefly, side-channel leakage directly results in an error as it manifests as a state impurity, guaranteed genuine randomness in the outcomes of joint measurements on Alice and Bob's particles, and the fact that high source brightness results only in a lower rate (due to increased noise) but does not allow an eavesdropper to perform a photon-number splitting attack. Most high-speed sources of entangled photons rely on nonlinear interactions between light fields in a crystal that result in the probabilistic generation of photon pairs entangled in a number of degrees of freedom (discussed briefly in Chapter 4). As the process is probabilistic, multiple pairs may be generated simultaneously. However, as entanglement is monogamous, measurements on the photons in one pair do not reveal any information about photons in the other pairs; each state is purely a two-photon state. This means that one must operate in a regime where only one pair at a time is produced to prevent the situation where losses cause Alice to detect photons that are uncorrelated with those reaching Bob.

1.3 Considerations from Communications Theory

In 1948, Claude Shannon presented formal notions of the amount of information transmissible in classical channels [23]. For example, in classical communications, messages may be communicated using sequences of binary characters (e.g., $\{0, 1\}$). The amount of information contained in a message is inversely proportional to the likelihood of receiving it; for example, messages of 8 bits containing only 1 (1111111) communicate less information than messages where either 1 or 0 is equally likely – there are 2^8 times more of them. The amount of information communicable per symbol in an alphabet A (here, $\{0, 1\}$) is measured in bits and is quantified by the Shannon entropy $H(A) = -\sum_{a \in A} p_a \log_2 p_a$, where p_a is the probability of obtaining symbol $a \in A$ (e.g., 0, or 1). For the case of binary alphabets where p_0 is the probability of obtaining element 0, the Shannon entropy takes the form $h(p_0) = -p_0 \log_2 p_0 - (1 - p_0) \log_2 (1 - p_0)$.

As the goal of QKD is to distribute a one-time pad without leaking any information to an eavesdropper,

it is important to quantify how much information the eavesdropper can extract, and what detectable errors that extraction might induce. After generating a raw key, Alice and Bob perform a series of classical error correction and privacy amplification procedures to eliminate information leaked to an eavesdropper. This may be quantified in part by the raw key fraction, r, which is the ratio of the sifted key length to the number of channel uses (the number of times Alice and Bob both record a detection in the same basis), and may be calculated by

$$r = \max\{I(A:B) - I_E, 0\},$$

where I_E is the information leaked to an eavesdropper and $I(A:B)$ is the shared entropy between Alice and Bob, also known as the shared information [24]. The shared entropy depends on the individual entropies of Alice and Bob and their joint entropy $H(A,B)$ and is given by

$$I(A:B) = H(A) + H(B) - H(A,B)$$
$$H(A,B) = -\sum_{a,b\in A} p(a,b)\log_2 p(a,b),$$

where $p(a,b)$ is the joint probability distribution for obtaining elements a, b from the transmission alphabet. When $I(A:B) < I_E$, it is impossible to generate any secret key between Alice and Bob – too much information is leaked to the eavesdropper.

The security of QKD depends on accurately quantifying the amount of information the eavesdropper can obtain about the secure key. This is frequently discussed as the maximum tolerable bit error rate (the difference between Alice and Bob's sifted keys) given the QKD protocol and certain assumptions about the abilities of the eavesdropper. If the eavesdropper is allowed only to intercept single photons, attempt to determine their state, and resend an attempted copy of that state, the maximum allowable bit error rate she can introduce reaches 17% before preventing Alice and Bob from generating any secure key. The maximum allowable bit error rate for BB84 drops to only 11% if the eavesdropper is capable of the most general physically allowed attacks [25].[1]

1.4 Challenges

Real-world QKD systems are subject to loss and degraded signal purity. They rely on single photon counting detectors (SPCDs) which typically erroneously fire many times per second (due to detector dark counts), and

[1] This assumes Alice and Bob are using only 2 mutually-unbiased bases (e.g. H/V and D/A) and a 2-level encoding. If, for example, they use all of H/V, D/A, and L/R, an eavesdropper will introduce more errors per bit of information gained, leading to a slightly higher QBER threshold of 12.7% [26]

it is impossible to distinguish a noise event from events arising from collection of a signal photon (assuming the noise event occurs during the expected arrival time of the signal). One way to lessen the effects of noise is to use strict coincidence timing to determine when an entangled pair of photons is collected by Alice and Bob: the timing of noise will be uncorrelated with the signal photons coming from a source of entangled photon pairs, so noise will be suppressed while counting in coincidence.

1.4.1 Fiber Optics versus Free-Space

Photons can be transmitted either through a waveguide (e.g., optical fiber), or through free space. Each channel type has advantages and disadvantages for the purposes of quantum communication and quantum key distribution. Fiber optic networks are fixed, and are therefore not susceptible to dynamic alignment issues; however, fiber is lossy, with losses on the order of 0.5 dB/km at telecom wavelengths ($\sim$ 1550 nm). For visible light this can be much worse, on the order of 1-3 dB/km [27]. Transmitting over practical distances in fiber (typically 20 km or more) can easily result in losses on the order of 10-20 dB, depending on wavelength and fiber type. Long fibers also introduce rapid fluctuations in the polarization and, in some cases, the distance between subsequent signal pulses, resulting in bit errors from fiber expansion. QKD protocols over fiber therefore typically use time-mode analogs of BB84 to avoid polarization scrambling. Furthermore, although fiber networks are fixed and therefore do not need to be actively aligned, they cannot be reconfigured easily. Moving platforms, such as aircraft, ships, and ground vehicles cannot access fiber links.

Free-space operation, on the other hand, does not introduce unpredictable polarization rotations, and can be used to connect rapidly-moving platforms [28, 29]. The atmosphere is not birefringent, so polarization states can be used instead of time-mode states. The attenuation in the atmosphere can be significantly better than fiber as it falls off as only $1/r^2$ due to diffraction (in the best seeing case), however, visibility is degraded by weather-dependent patterns of aerosols and vapors, and time-dependent temperature gradients in the atmosphere result in turbulence that can further degrade visibility (e.g., mirages over a hot road). We discuss this effect in more detail in Chapter 3. Additionally, free-space systems must have a well-engineered subsystem for ensuring the signal states are projected toward the receiver, called "pointing and tracking." Loss in free-space systems is dominated by turbulence and the ability of the pointing and tracking system to maintain a robust connection between the transmitter and the receiver.

1.4.2 Sources

Experimental QKD systems frequently minimize or elide the possibility of side channel attacks, in which an eavesdropper could obtain information about the transmitted states by measuring a correlated variable [30]; for example, if, in a BB84-based system, the source of $|H\rangle$ photons operates at 655 nm while the source of $|D\rangle$ photons operates at 656 nm with a 0.5-nm bandwidth, the eavesdropper could gain complete information about the basis used by making a simple spectral measurement (e.g. with a narrowband filter), thereby rendering BB84 insecure. One pernicious example that this author finds particularly amusing is given in Ref. [31], in which the authors demonstrate that a promising technique for rendering the spectra of multiple diode laser sources indistinguishable using an external seed laser opens another avenue of attack, as an eavesdropper may be able to manipulate the behavior of Alice's source lasers by applying a polarized seed laser of her own. Any other distinguishing feature of the sources used to generate the single-photon states is exploitable (for example, pulse shape variability). It is difficult, however, to engineer sources that operate quickly and uniformly enough to obviate side-channel attacks while being robust enough to deploy on a variety of platforms.

1.4.3 Loss and Quantum Repeaters

Quantum key distribution is not inherently dependent on loss. If only a few photons are received by Bob, then the BB84 protocol simply results in only a few bits of secure key material. However, if there is any noise (or if the states received by Bob are not perfectly correlated with what Alice transmits or measures), the error rate observed by Alice and Bob will increase rapidly as the signal-to-noise ratio is reduced. Eavesdroppers introduce bit errors, and, to be conservative, Alice and Bob must interpret all noise as a loss of secure bits. Furthermore, our ability to estimate the bit error rate is dependent on the number of bits received due to finite sample statistics, often called the "finite-key effect". The number of truly secure key bits that can be extracted after performing multiple rounds of BB84 therefore drops exponentially when transmission is low, and, as shown in Ref. [32], is estimated by the relation

$$\ell \leq n(q - h(Q_{\text{tol}} + \mu)) - \text{leak}_{\text{EC}} - \log \frac{2}{\varepsilon_{\text{sec}}^2 \varepsilon_{\text{cor}}},$$

where

$$\mu = \sqrt{\frac{n+k}{nk} \frac{k+1}{k} \ln \frac{2}{\varepsilon_{\text{sec}}}}.$$

Here, l is the final secure key length, n is the number of successful channel uses (e.g., how many times Bob receives a photon), q is the source preparation quality (close to 1 for high-quality sources), h is the binary

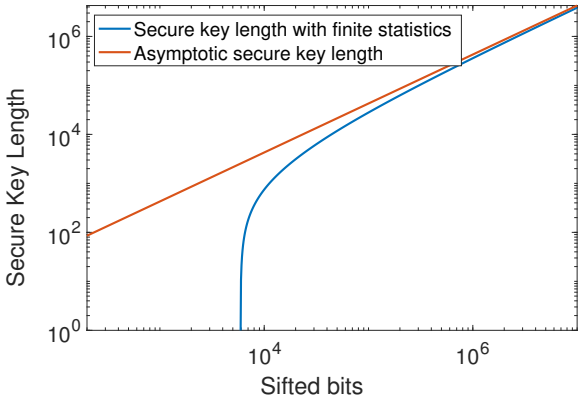

Figure 1.2: We are only able to approximate values of Q given a finite number of input sifted bits; underestimating Q (or, more importantly, the number of bits of information leaked to an eavedropper) results in insecure bits in the final key. To account for this effect we must scale down the number of secure bits on the output. Here we assume the true QBER is $Q = 5\%$, that half of the sifted bits are used for parameter estimation (in practice this may be a much smaller fraction), and that $\varepsilon_{\text{sec}} = 10^{-14}$ and $\varepsilon_{\text{cor}} = 10^{-10}$. For short input sifted keys ($<$ approximately 5000 bits) we are unable to generate any secure key at all, and for lengths slightly above this cutoff we are unable to use a significant fraction of the input bits. One solution to the finite-key problem, if possible, is to aggregate enough data for QKD postprocessing that finite statistical effects are minimized.

entropy function, k is the number of bits used to estimate bit errors, and the ε are parameters establishing the maximum tolerable probability of an eavesdropper guessing the correct final key given all other parameters and are set by hand (typically to $\varepsilon_{\text{sec}} = \varepsilon_{\text{tol}} \approx 10^{-10}$). Additionally, Q_{tol} is the channel error tolerance (e.g., bit error rate), and leak_{EC} is an estimate of the amount of information leaked to the eavesdropper during the classical error correction stages of the protocol, usually estimated to be $\text{leak}_{\text{EC}} \approx h(Q_{\text{tol}})$ as a compromise between various extant error-correction models. Figure 1.2 shows the effect of finite statistics on the secure key length. This relation between ℓ and Q may be understood as Q being artificially increased when the number of detection events used to estimate Q is small, i.e., Q is assumed to be the worst it could be given the poor statistics used to determine it, given some security thresholds ε. The final extractable key length is therefore significantly reduced over long, lossy QKD channels, where the number of bits received is too small to bound the information leaked to an eavesdropper.

QKD has a distinct distance/rate tradeoff. As such, research in quantum communications is concerned not only with the maximum achievable rate, but the maximum rate achievable at the longest possible distance. To give a sense of the state of the art, over laboratory distances, the highest rate achieved over 20 km is 26.2 Mb/s (1.07 Mb/s over 83 km)[33]. Fiber-based QKD has reached distances of 400 km at rates of 3.2×10^{-4} bps (unusably low) [34]. These current bests for QKD are nowhere close to commercial telecommunication link lengths and speeds, so the practicality of QKD hinges on our ability to built loss-tolerant protocols and

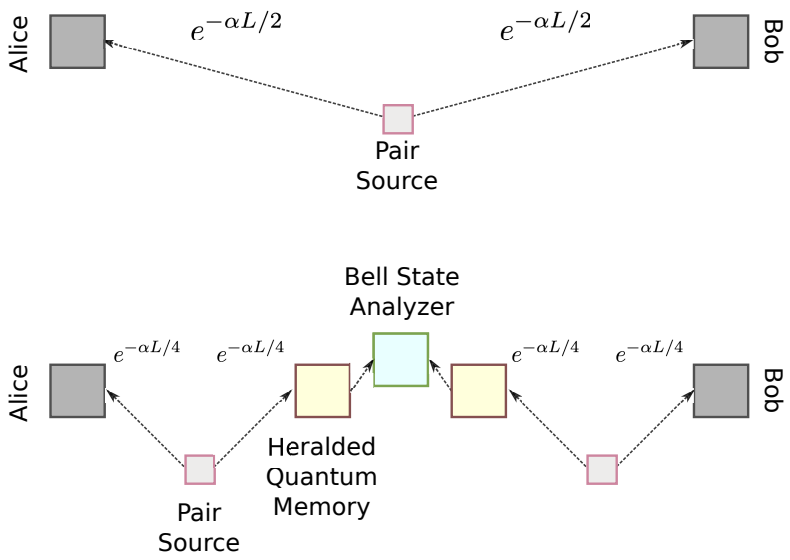

Figure 1.3: Sketch of a channel using quantum repeaters. (top) Model for a standard entanglement-based protocol, where paired photons must traverse half of the total channel length. (bottom) To enhance the probability that Alice and Bob can each receive a photon from an entangled pair, the channel is split using multiple entanglement sources and quantum memories, called quantum repeaters. The quantum memories are used to store photons until adjacent quantum repeaters or sources produce viable entangled pairs. Each repeater then swaps the state of the received particle onto the particle to be propagated down the channel to the next repeater using a Bell state measurement (a projection onto the four Bell states; with linear optics only three states of the four may be distinguished) on the received photons. If all entanglement swaps are successful, Alice and Bob share entangled photons. This procedure results in effective loss that scales much better than attempting to traverse the entire channel with a single pair source [37], at the expense of losses within each repeater due to input/output coupling and light-matter interactions.

networks.

One solution to this problem is a quantum repeater [35, 36], a device that is able to probabilistically extend the length of a quantum communication channel by storing the quantum state of one photon until another particle arrives (a quantum memory), allowing for more attempts at a single use of the channel. We will briefly discuss the function of repeaters to motivate future chapters, but for more details on implementations see Ref. [37].

In the abstract, quantum repeaters help extend QKD channels as follows: assume Alice and Bob are a distance L apart and would like to use an entanglement-based QKD protocol. For the protocol to be successful, each photon of the pair must propagate at least a distance $L/2$ from the source to Alice and Bob, for example, if the source is located at a central trusted third-party location. The protocol will only succeed when Alice and Bob both receive a photon at the same time. In the presence of simple absorptive loss, the success probability takes the form $e^{-\alpha L}$, where α is the absorption coefficient of the channel. If they have access to a series of $N-1$ quantum memories distributed evenly between them (Fig. 1.3), then the distance between each segment is $L/(2N)$ (Alice and Bob act as destination nodes). At each repeater there

is quantum memory capable of holding photons for longer than the the protocol run time with minimal state degradation and a source of entangled photons. The repeater generates entangled photons, sends one along the channel, and waits until a signal photon is received from an adjacent node, whereupon the repeater performs an *entanglement swapping* protocol (wherein the joint state of two unentangled photons is made to be entangled by performing a joint measurement on their individual partners, see Ref. [37] for more details). This generates entanglement between the photon the repeater already transmitted and the corresponding photon for the second repeater. When all repeaters have successfully and simultaneously swapped their entanglement with the next repeater, Alice and Bob share an entangled pair. The total success probability for a single channel use then scales as the success probability for the quantum memory and entanglement swapping operations, much better than attempting to traverse the entire channel in one shot [38]. However, the viability of this strategy for improving system transmission depends on the success probabilities of operations internal to the quantum repeater. As quantum memories require matter-light interactions and very high input-output system coupling, current state-of-the-art success probabilities are somewhat low [39, 40, 41].

1.5 Summary

In this thesis we address a few of the challenges with creating multi-platform quantum cryptographic networks:

- Quantum repeaters require extremely high input-output coupling efficiency to be useful (see, for example, Ref, [41]), as any losses in a single repeater can drastically lower the probability of successfully sending a photon from Alice to Bob. In Chapter 2 we demonstrate a technique for adaptively improving the coupling of a single photon-emitting point source (e.g., ions, quantum dots, etc.) into a single-mode fiber without using traditional optical wavefront sensing.

- Deploying quantum key distribution in free-space applications requires a deep understanding of how atmospheric turbulence impacts protocol performance and security; however, developing a physical free-space testbed is difficult. In Chapter 3, we discuss techniques for simulating the effects of atmospheric turbulence on propagating light fields and how they can be physically implemented in a laboratory setting.

- Future quantum cryptographic networks must be able to link mobile platforms, such as air- and watercraft. Quantum key distribution in free space over useful distances will be subject to unpredictable losses due to turbulence in the atmosphere. In Chapter 4 we demonstrate a new method that improves

the signal-to-noise ratio of QKD systems over very turbulent channels by allowing for post-selection of portions of the transmitted signal that are most likely to actually contain a signal photon.

- Free space quantum key distribution typically requires elaborate telescopes, well-conditioned single-photon sources, and motorized tracking systems that would be impossible to use on many platforms of interest, such as cube satellites and small aircraft. In Chapter 5, we discuss our progress in developing an airborne quantum key distribution system on twin multi-rotor unmanned aerial systems (UASs).

Chapter 2

Adaptive Optics

This chapter reproduced with permission from A. D. Hill, et al,"Optimizing single-mode collection from pointlike sources of single photons with adaptive optics," Opt. Express 25, 18629-18642 (2017) (Ref. [42])

2.1 Introduction

Efficient collection of photons emitted from point-like sources (e.g., trapped ions [43, 44], nitrogen-vacancy centers [45], quantum dots [45, 46], etc.) is often critical for the interrogation of single quantum systems and the efficient realization of quantum computing and communication protocols. Collection optics often introduce significant wavefront aberrations due to the clipping and focusing of collected light from an effective point source. Due to these aberrations and a small collection solid angle, coupling of the light emitted by an interrogated quantum system into a single-mode optical fiber (SMF) – the most natural way to connect separate quantum systems – may be extremely inefficient, despite the fact that the light is emitted into a single spatial mode (e.g., a dipole radiation mode). Recent work in producing radiation from pointlike objects in the desired mode has been quite successful using microcavities and antennae matched to the source [47, 48, 49]; however, they may be difficult to implement in existing experimental setups in general. Other purely optical approaches have proven successful [50], but are designed for isotropic sources. Note that as photon emitters are contemplated for use in multi-photon quantum information processing applications, e.g., multi-mode quantum repeater networks [51] or demonstrations of photonic integrated circuit based quantum algorithms [52], even modest improvements in coupling efficiency can lead to large net enhancements – especially in multiphoton experiments.

Adaptive optics (AO) allows experimental apparatus to correct for the aberrations created by large numerical aperture lenses and other collection system optics, and thus to optimize coupling to collection fibers (which can then transport the light to detectors, other quantum systems, photonic circuitry, etc.). AO is effective for optimizing spatial modes for coupling applications [53, 54] and nonlinear optics [55]; however, AO can be difficult in the low-intensity limit, where wavefront measurements are unreliable and

photon count fluctuations can thwart traditional gradient methods of system optimization. In this regime, AO must be performed entirely at the single-photon level [56, 57]. We have developed a "drop-in" AO-enhanced collection system that incorporates a novel noise-resistant genetic algorithm to optimize the shape of a deformable mirror for collection of single photons from a point-like emitter into a single-mode optical fiber (SMF) (although we use a SMF as the most common photonic transfer element, our techniques should also work in a large variety of other applications, e.g., enhanced coupling to/from waveguides, plasmonic devices, nano-antennas, etc.). We have tested our system extensively with a simulated point source (directing light through a sub-wavelength aperture) as well as with actual quantum dot photon emitters, and observed significant enhancements in both cases. After describing the genetic algorithm strategy we will present the results of optimizing collection from a sub-wavelength pinhole and from an InAs quantum dot.

2.2 Genetic Algorithm

We control the 69 independent electromechanical actuators of a deformable mirror (Alpao DM69) using an in-house designed genetic algorithm specially suited for AO applications with single photons [Fig. 2.1]. Our algorithm generates random mirror shapes based on previous configurations and weighted by the performance of the generating families; the algorithm first estimates the effect of deforming the mirror on the measured signal (e.g., photon counts coupled into single-mode fiber per second) by computing the count-rate variance observed while randomly permuting subsets of the basis elements. Basis elements may be the full 69 actuators or the 30 lowest-order Zernike polynomials (e.g., tilt, focus, coma, etc.) created over the 69-actuator space of the mirror surface, the choice of which leads to changes in optimization behavior (discussed below). "Child" shapes are constructed by randomly weighting all basis elements to form new generations of mirror configurations. After initialization, randomized deformations are then applied to the mirror surface, creating the "children" of the subsequent generation. The count rate for each of these children is compared; the best test deformations in each generation (the new parents) are weighted by performance and then probabilistically combined to form a new generation of mirror deformations. For our tests, generations were comprised of 20 children from 10 parents selected from the previous generation. Over time (of order 100 generations in our tests) the variance in the generated children is lowered, which allows convergence to an optimal mirror shape.

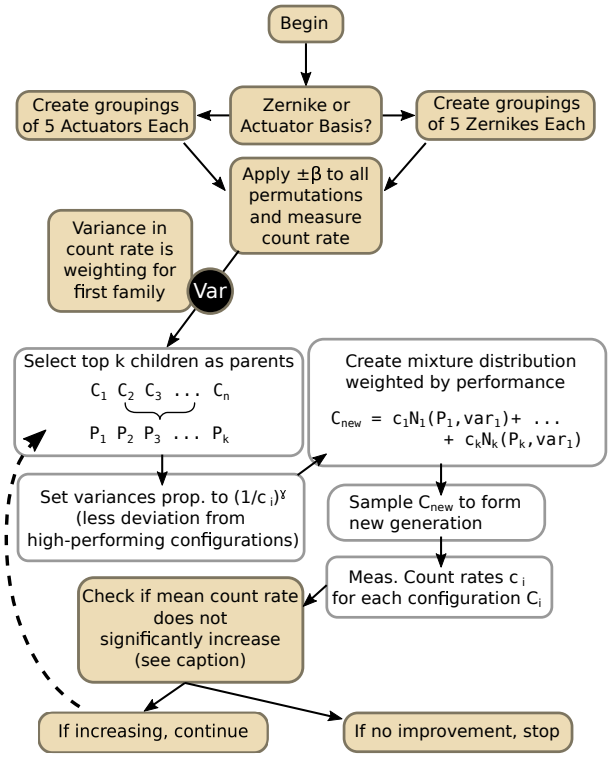

Figure 2.1: A schematic of the genetic algorithm. The first generation of mirror configurations is pseudo-randomly generated using a sum of gaussian distributions with variances established by the initialization procedure. Each new generation of mirror configurations is created by weighting gaussian random variables centered on each top-performing mirror configuration, according to count rate. An arbitrary number of children may be generated by sampling this mixture gaussian. The variance of the gaussian random variables is reduced in inverse proportion to configuration performance to allow the mirror generations to converge to an optimal solution. The algorithm completes when it is estimated that the mean has not statistically increased for at least 6 generations. *Reprinted with permission Ref. [42]*

2.2.1 Detailed Description of Algorithm

The algorithm begins by estimating the initial size of the space to search during optimization, which results in the number of families chosen in each generation. The d basis elements (in our case, 69 actuators or 30 Zernike polynomials) are grouped into sets of five. In the case of the Zernike basis, the elements are grouped by polynomials of nearest order (using the Noll ordering, groupings are $\{Z_1, \ldots, Z_5\}, \{Z_6, \ldots, Z_{10}\}, \ldots$). In the actuator basis the elements are ordered by device index and grouped similarly, with thirteen groups containing 5 basis elements and one containing 4, to reach the full 69 actuators. For each grouping, each basis element is assigned a magnitude $\pm\beta$, where β is a parameter set before optimization; basis elements outside the current grouping are set to a neutral (flat) position. Each set (6 for Zernikes, 14 for actuators) of 5 basis elements is permuted through 32 combinations of voltages ($\pm\beta$ for each of the 5 basis elements in each group) so that every combination of elements within each grouping is applied to the deformable mirror. The collected output power or count rate is recorded for each permutation, and the variance in the count rate between all permutations is used to initialize the intergenerational variances a_i, discussed below. Assuming β is large enough, this ensures that the initial set of n children (n set by the user; we typically used 20 children per generation) adequately spans the configuration space needed to optimize collection, since more effective (higher variance) basis elements are allowed to contribute heavily to the first generation of mirror configurations.

After initialization, in each generation we choose k future parents from the previous n children. The next n children are generated by breeding the chosen k parents using the following procedure. Let $\mathbf{Y}_i^t$ be a random vector of d elements (d being dimension of the space, as defined above) representing the test mirror configuration of the i-th parent at time step t. We select the k highest-performing configurations (those resulting in the most photon counts registered leaving the single-mode fiber in a fixed measurement interval) and form the new generation Y^{t+1} by sampling

$$\mathbf{Y}^{t+1} = \sum_{i=1}^{k} \omega_i N(\mathbf{Y}_i^t, a_i \mathbf{I}_d), \tag{2.1}$$

where $N(\mathbf{u}, \sigma)$ is a multivariate normal distribution with mean $\mathbf{u}$ and covariance σ, and $\omega_i = c_i / \sum_{i=1}^{k} c_i$ is a weighting given by the relative count rates c_i of each of the selected parents. The a_i are set to $a_i = (C/c_i)^\gamma \rho$, with γ and ρ set by the user; here C is an estimate of the maximum possible count rate in the absence of all aberrations, ρ sets the extent of the search space, and γ sets the sensitivity of the search on the count rate relative to the estimated maximum. In our tests, we typically used $n = 20$, $k = 10$, $\rho = 3.6 \times 10^{-6}$, and $\gamma = 0.5$. These values were influenced by the simulations shown below; however, ρ and γ are strongly

dependent on the experimental setup (very sensitive corrections require small ρ to avoid losing all coupling, for example).

The algorithm completes when the statistics of the count rate do not change significantly for several generations. First, we check if the mean count rate for the past 3 generations differs significantly from the mean for the prior 3 generations (comparing 6 generations total). If the mean has not significantly changed, we then take three new generations with twice the measurement time per mirror configuration to reduce fluctuations. If the mean does not change significantly between these three generations, we stop. In each case we consider the statistics of two generations to be significantly different if $|\bar{x}_n - \bar{y}_n| < 0.93\sqrt{2\bar{y}_n/n}$, where x_n is the set of the counts from generation x and y_n is the set of counts from the current generation y (the two-sample Z-test); 0.93 is a parameter that balances the probabilities of either accepting or rejecting convergence erroneously (for $p < 0.05$ rejection of both Type I and Type II errors, this would be 1.96; 0.93 is chosen to balance both errors).

The algorithm depends on a number of parameters that may be customized to the experimental implementation. For example, the size of the initial test mirror deformations should be close to the magnitude of aberrations encountered in the laboratory. Through numerical simulations of the algorithm we have indentified some general starting parameters (discussed later).

2.3 Laboratory Simulation of a Point Source

We physically simulate the collection of light from a point-like emitter using a 600-nm pinhole etched through a 200-nm thick aluminum film on the tip of a SMF [Fig. 2.2, inset]; the result is a highly-divergent source (NA > 0.8) of 675-nm light. A 4-mm pane of glass representing, e.g., the vacuum window of a simulated emitter's setup, separates the pinhole from either a 10X, 0.25-NA microscope objective or a 0.5-NA aspheric lens; these approximately collimate $\sim$ 500 μW (16%) of the collected light and also serve to define the effective numerical aperture of the collection system by clipping the mode emitted from the fiber. Using two-lens imaging we expand the resulting beam to cover the 10.5-mm-diameter surface of our deformable mirror (9.6° incidence). After reflecting from the mirror, the light is collected by a microscope objective and focused onto a single-mode fiber (Thorlabs 460HP). In order to obtain the optimal coupling without the AO mirror, the fiber position is first optimized using a precision piezo-controlled 6-axis stage (APT Nanotrak), though in practice most of the alignment is optimized using only the positional (X/Y/Z) degrees of freedom. The fiber output is then projected through a beamsplitter; the transmitted light is attenuated and sent to a free-space single-photon detector, while light from the reflected port is recorded by a photodiode to monitor

the coupling efficiency while the genetic algorithm optimizes single-photon counts.

Figures 2.3(a)-2.3(b) demonstrate the results of typical optimization runs to maximize the coupling efficiency into the final, single-mode fiber. It is important to note, however, that because the algorithm begins optimization with a randomized mirror configuration, the initial couplings recorded in these plots are typically one-third to one-half the coupling for the pre-algorithm, flat-mirror configuration. Thus, the actual overall improvement was somewhat more modest, as the first set of mirror configurations necessarily deviated from the initial optimized alignment with a "flat mirror" surface (a mirror configuration calibrated for near-flatness). Nevertheless, improvements of 100-200% in coupling efficiency over the flat-mirror configuration were typical, though the initial and final couplings are dependent on initial beam misalignment and beam-waist mismatch into the final fiber. Given the finite number of actuators in our deformable mirror, correction of wavefront aberrations is accurate to 14th-15th order Zernike polynomials (Noll index); higher-order aberrations from lens distortion and the resulting Airy pattern will not be completely corrected, which limits the maximum coupling achievable in our simulation.

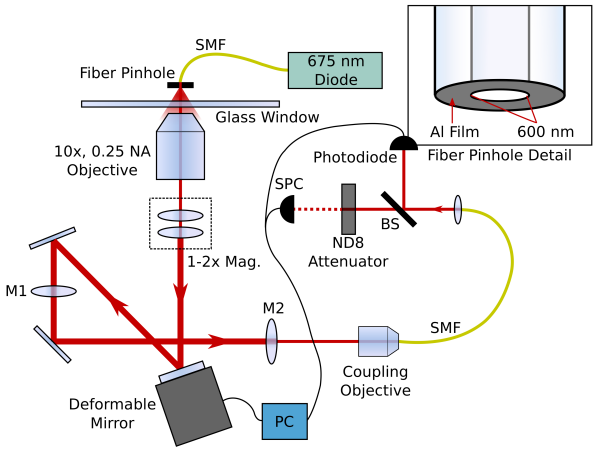

Figure 2.2: Schematic of the experimental setup. Laser light at 675 nm is collected from a pinhole, collimated, and manipulated using a deformable mirror before being contracted (minification lenses M1-2) and focused into a single-mode fiber. The coupling efficiency is monitored using a photodiode, but optimized using a single-photon counter (SPC) after attenuation. *Reprinted with permission Ref. [42]*

In Fig. 2.3(a) we show a comparison of the system's performance for various initial count rates. For lower count rates (10 kcps), the count rate is more significantly affected by shot noise and direct comparison of mirror configuration performance is difficult. Nevertheless, our system is able to optimize collection almost as well as for higher count rates (75 kcps or greater) with a small decrease in optimization speed and an increase in the variability of the final count rate. As shown in Fig. 2.3(b), use of the actuator basis is better able to correct aberrations in the long term given the higher number of basis elements (all 69 actuators); in

contrast, the Zernike basis allows for a faster initial optimization due to a shorter initialization time. The actuator basis was used in all tests (decribed below) for consistency.

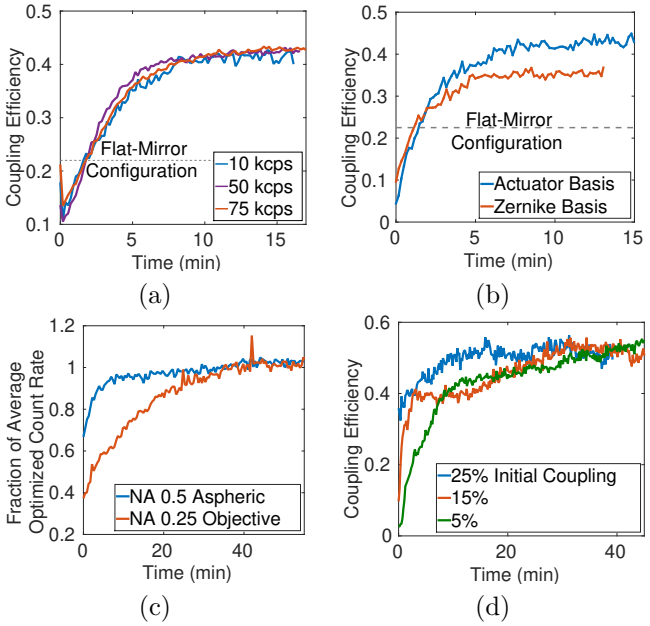

Figure 2.3: Performance of the algorithm (a) for various initial count rates (in the actuator basis after first optimizing the flat-mirror coupling), showing a robustness against signals affected heavily by shot noise and (b) for optimization in the basis of the first 30 Zernike polynomials or the basis of all 69 mirror actuators. (c) A comparison of AO correction for 1", 0.5-NA aspheric and 0.25-NA objective collection lenses using the actuator basis, normalized to the average count rate after optimization, in order to emphasize the convergence speed difference. The unnormalized final coupling efficiency with the aspheric lens ($< 5\%$) was actually much lower with the objective (50%). (d) Performance for system misalignment where the assumed 'best' alignment results in a coupling of 50%. The algorithm is able to correct even severely misaligned systems, with a small cost in optimization time. Each line plots the maximum counts per generation, with the classical coupling given by the monitoring photodiode. *Reprinted with permission Ref. [42]*

We also evaluated our technique's adaptibility to changes in both experimental parameters, such as the coupling enhancement for different collection optics, and various algorithmic parameters (e.g., number of children per generation, number of parents selected per generation, etc). In the former case, we tested both a 10x, 0.25-NA microscope objective and a 1", 0.5-NA aspheric lens to collect light from the fiber pinhole. We anticipated that the ability of the mirror and algorithm to correct for aberrations should depend heavily on the order of those aberrations, which in turn depend on the quality and numerical aperture of the collection lens. A comparison of the algorithm performance for both lenses is shown in Fig. 2.3(c). The microscope objective produces predominantly higher-order aberrations, while the aspheric lens also introduces significant lower-order aberrations; this is reflected in the coupling optimization time required for both cases: the system is able to correct for low-order aberrations (spherical aberration, coma, etc.) much faster than high-order aberrations due to errors introduced in higher-order corrections from the

finite number of actuators. The overall coupling efficiency for the aspheric case, however, was comparatively low ($< 5\%$, versus 50% with the microscope objective), as the aspheric lens produces larger aberrations of all orders. Our particular deformable mirror, the Alpao DM-69, accurately corrects most low orders other than severe spherical aberration and secondary astigmatism. The maximum coupling efficiency into a single mode goes as $\eta = 1 - (2\pi/\lambda)^2\sigma^2$, where σ is the RMS wavefront error [58]. For example, for our mirror, addressing a peak-to-valley spherical aberration of 7.26 μm results in an RMS residual error of only 107 nm, but even this corresponds to a maximum theoretical coupling efficiency of only 5% *after* optimization. For other orders this effect is much less pronounced. This is a limitation of the device itself, and should be reduced significantly for higher-resolution models.

We investigated the effect of the initial system alignment (initial coupling) on the algorithm's performance. For these tests, the algorithm achieved a 50% coupling efficiency after a manual alignment; the system was then intentionally misaligned varying amounts by tilting the final mirror before the single-mode fiber. Finally, the algorithm-controlled deformable mirror was used to try to compensate for this misalignment. The system was able to correct these issues (Fig. 3d), though the optimization time increased for poorer alignments.

The parameters controlling the genetic algorithm itself significantly alter its performance. We identified the total time to optimization (defined as the time for the algorithm to reach a steady state with respect to coupling efficiency), and the ratio of AO-corrected coupling efficiency to uncorrected coupling efficiency as metrics for evaluating overall system performance. We tested several strategies: holding the number of children constant in each generation, or linearly or exponentially increasing the number of generated children in proportion to the iteration number. The amount of time spent per generation increased in all cases, though the number of children per generation (and therefore the time per generation) was significantly larger at the end of optimization for the exponential case. Because our algorithm automatically reduces generational variance as the coupling approaches the user-specified theoretical maximum, this additional time required for these final runs does not necessarily improve the overall optimization efficiency. Monte Carlo simulations of the algorithm indicate that the fastest strategy on average (comparing using a fixed number of children for each iteration, or linearly or exponentially increasing the number of children per iteration) is to fix the number of children at 20 mirror configurations per generation. These simulations also suggest that, for our experimental setup, the algorithm should reach completion after approximately 100 generations (discussed later).

2.4 Application to Quantum Dots

2.4.1 Stationary Collection

The optimization of light collection is crucial for performing experiments using quantum optics techniques in QD systems [59]. Recent work on spin-photon entanglement using single charged QDs are largely limited by collection efficiency [60, 61, 62]. Adaptive optics could have a significant impact on experiments such as entanglement swapping via intermediate entangled photons, and were, in fact, used in recent work using NV centers to perform a loophole-free bell test [63].

We applied our technique to collecting photons from self-assembled InAs quantum dots (QD) grown using molecular beam epitaxy and embedded in a distributed Bragg-reflector (DBR) cavity for enhanced light collection. As shown in Fig. 2.4(a), the sample is cooled by a liquid-helium optical cryostat; excitation and collection are performed by a high numerical aperture lens (NA=0.68). A Ti:Sapph continuous-wave laser is tuned above the ∼890-nm band-gap, exciting carriers into the conduction band; these carriers then radiatively recombine at the exciton resonance. The excitation laser (<890 nm) is filtered out by a 925-nm long-pass filter while the QD luminescence (∼950 nm) is sent to a single-grating spectrometer where single dot signatures are seen on a liquid nitrogen-cooled CCD. In order to isolate a single QD for optimization, we use an etalon (10-nm free spectral range, finesse of 100) placed after the long-pass filter. The beam path is sent to a single-mode fiber via the deformable mirror; the fiber output is sent to the spectrometer to verify that only QD luminescence is coupled into the fiber. Finally, a fiber-coupled single-photon avalanche photodiode (SPAD) monitors the single photon counts from the dot and serves as the input for the genetic algorithm software.

The algorithm consistently improved coupling of single photons from the QD [Fig. 2.4(b)]. Each mirror configuration was tested for one second, with an initial count rate of approximately 15-20 kcps. Though the approximate 1-ns lifetime of the QD would imply a maximum photon production rate of 1 GHz, poor detector efficiency (2%), relatively low-NA collection optics (0.6 NA), and light lost in the sample itself resulted in a low detection rate (<50 kcps). Furthermore, due to the variability in the brightness of individual quantum dots and uncertainty in the alignment of driving lasers onto the QD, establishing an absolute coupling efficiency is impossible (though in this case the absolute efficiency is likely to be quite low). Better index matching from GaAs to air using a solid immersion lens could improve this absolute efficiency. Nevertheless, the algorithm was able to improve the count rate $35 - 50\%$ over the non-optimized (flat-mirror) case across all runs. The discrepancies between the runs were due to variability in the initial dot coupling and overall system alignment between each run, which the system was not able to completely correct.

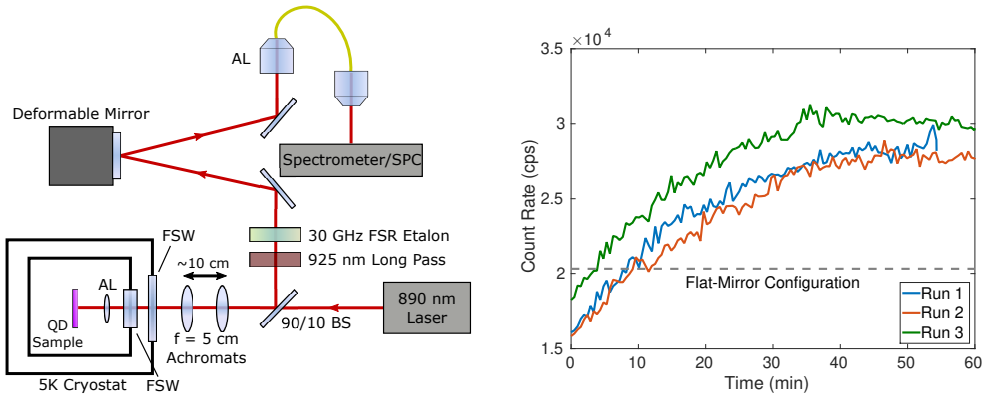

Figure 2.4: Optimizing collection from an InAs quantum dot (QD). (a) Photoluminescence light is collected from the QD and collimated with a focusing asphere (AL, Thorlabs 352330-B) before passing through two fused-silica cryostat windows (FSW, 0.2" and 0.125", respectively). Two 5-cm focal length achromatic lenses (Thorlabs AC254-050-B-ML) are used for fine adjustment of the beam. The dot light is then filtered with a 925-nm long-pass and a 30-GHz free-spectral-range etalon to remove residual pump light. The filtered QD light is manipulated by the deformable mirror and coupled into a single-mode fiber (SM980) with an aspheric lens (Thorlabs C260TME-B). The light is then sent to a spectrometer to confirm the collection of QD photoluminescence. After tuning the system for a particular QD, the QD emission is sent to a single-photon counter for optimization using our algorithm. (b) Several optimizations of the collection from a single quantum dot, showing up to 50% improvements in coupling. The optimization rate and final count rate were dependent on the overall system alignment and temperature of the dot. *Reprinted with permission Ref. [42]*

2.4.2 Simulation of Source Drift

In practice, almost all point emitters that emit into free space are likely to display some amount of drift, e.g., situations where effects such as thermal expansion or ion drift translate the photon source relative to the collection optics. Ideally, the adaptive collection system should compensate for such drifts. Here we evaluate how well our system achieved this. Because our experimental setup did not enable us to realize such a displacement of the actual source in a controllable way, we instead used a piezoelectric translation stage to apply a 3.95-μm peak-to-peak displacement in the horizontal position of the collection fiber with respect to the final coupling lens. The movement of the fiber results in an oscillation in the count rate as seen by the final single-photon detector, just as a lateral shift in the position of the quantum dot emitter would. In this way, we are able to emulate some of the aberrations caused by a slowly moving source.

We tested the behavior of our QD collection system in the presence of such simulated drifts of two frequencies. Figure 2.5 shows data with imposed sinusoidal "drifts" with periods of 1000 s (slow drift) and 200 s (moderate drift). For these drift speeds (3.95-μm full displacement, giving peak speeds of 0.12 μm/s peak speed for the 200s period case), we observed an improvement in the mean count rate and reduction in the oscillation magnitude when the AO stabilization was employed ($\sim$ 100% improvement in mean coupling

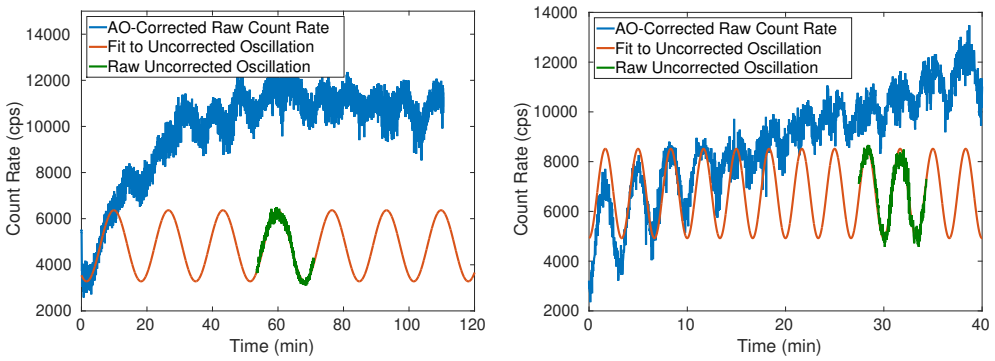

Figure 2.5: Behavior of the adaptive optic system in the presence of source drift (oscillating count rates of (a) 1000 s period (1 mHz) and (b) 200-s period (5 mHz). Sinusoidal fits to measured oscillations (green data) before correction are shown in red as a guide to the eye. For slow oscillations ($>$ 200-s period) the algorithm improved the mean photon count and was able to largely suppress the oscillation amplitude. For faster oscillations, the ability of the algorithm to reduce the oscillation amplitude and improve the mean count rate decreased approximately inversely proportionally to the "drift" rate. *Reprinted with permission Ref. [42]*

in both cases, with oscillations reduced from $\sim 50\%$ peak-to-peak variation to $\sim 25\%$ for the 1000-s period case); this indicates that our AO system should be able to provide significant automatic stabilization against the types of drift found in many experiments with point sources.

2.5 Simulations

The other parameters influencing our algorithm's behavior (the number of parents k, the step rate between generations γ, the scale of the step rate ρ, and the initialization voltage β) are also tunable by the user. Some of these parameters are situation-dependent; for example, in cases where coupling light from a point emitter is extremely sensitive to very small beam aberrations, using large initial deformations (large β) or allowing the children to differ significantly (large ρ) may cause algorithm performance to degrade. However, in some cases we are able to set general guidelines for choosing parameters as a result of numerical simulations of the algorithm. A summary of these results is presented in Fig. 2.6.

In general, simulations show that using $n = 30 - 50$ children per generation and selecting $\leq 50\%$ of the children as parents for the new generation is optimal. For our experiments, performed before these simulations were completed, we chose $n = 20$ and $k = 10$. The parameters ρ and γ must be chosen depending on the application. Simulations [Fig. 2.6(c)] show that ρ should not be too low (too little intergenerational variance) or too high (too much intergenerational variance); for our pinhole coupling application we chose 3.8×10^{-6}, which is related to the maximum actuator voltage allowed on the deformable mirror and the

magnitude of the aberrations being corrected. Furthermore, simulations show that γ should be large for short-term optimization, though long-term optimization is unaffected by γ [Fig. 2.6(d)]; in practice, however, large γ values may result in unstable behavior. For this reason we chose $\gamma = 0.5$ for all experimental tests.

2.6 Tests of Varying Family Sizes for Pinhole Coupling Optimization

Using our pinhole-simulated point source, we experimentally evaluated the algorithm's performance as the number of children per generation was varied [Fig. 2.7]. We compared holding the number constant at 20 children per generation, adding one child every generation (linearly increasing from 20 children), and exponentially increasing the number of children as $e^{0.04m} + 20$, where m is the generation number. The final optimized value was similar for all cases; however, the optimization speed of the linear case was significantly lower than for the constant or exponential cases. We chose to hold the number of children constant at 20 children per generation for all other tests.

2.7 Repeated Runs for a Single Quantum Dot

Coupling into a SMF efficiently requires flattening the phase profile of the aberrated beam. The mirror configuration that optimally couples into a SMF will therefore be conjugate to the wavefront of the coupled light. Figure 2.8 shows the final optimized AO mirror phase profiles for three repeated optimizations of a single quantum dot as discussed in the main text, as well as the corresponding Zernike decompositions of the final mirror configurations. From the average of the three runs we see a tendency to correct significant defocus (index 4), spherical aberration (index 12), and some higher-order aberrations at indices 16 and 17 that may indicate correction of the typical "Maltese-cross" pattern emitted from the quantum dot. Qualitatively, there is some agreement in the general sign and magnitude of each Zernike order applied to the mirror in the optimal configuration; however, the variation between the runs may imply that either a global maximum was not achieved in each case, or the aberrations were not addressable by the mirror (e.g., aberrations were of too high an order). We anticipate that a higher-resolution mirror or longer run times may reduce the variation in final optimization conditions.

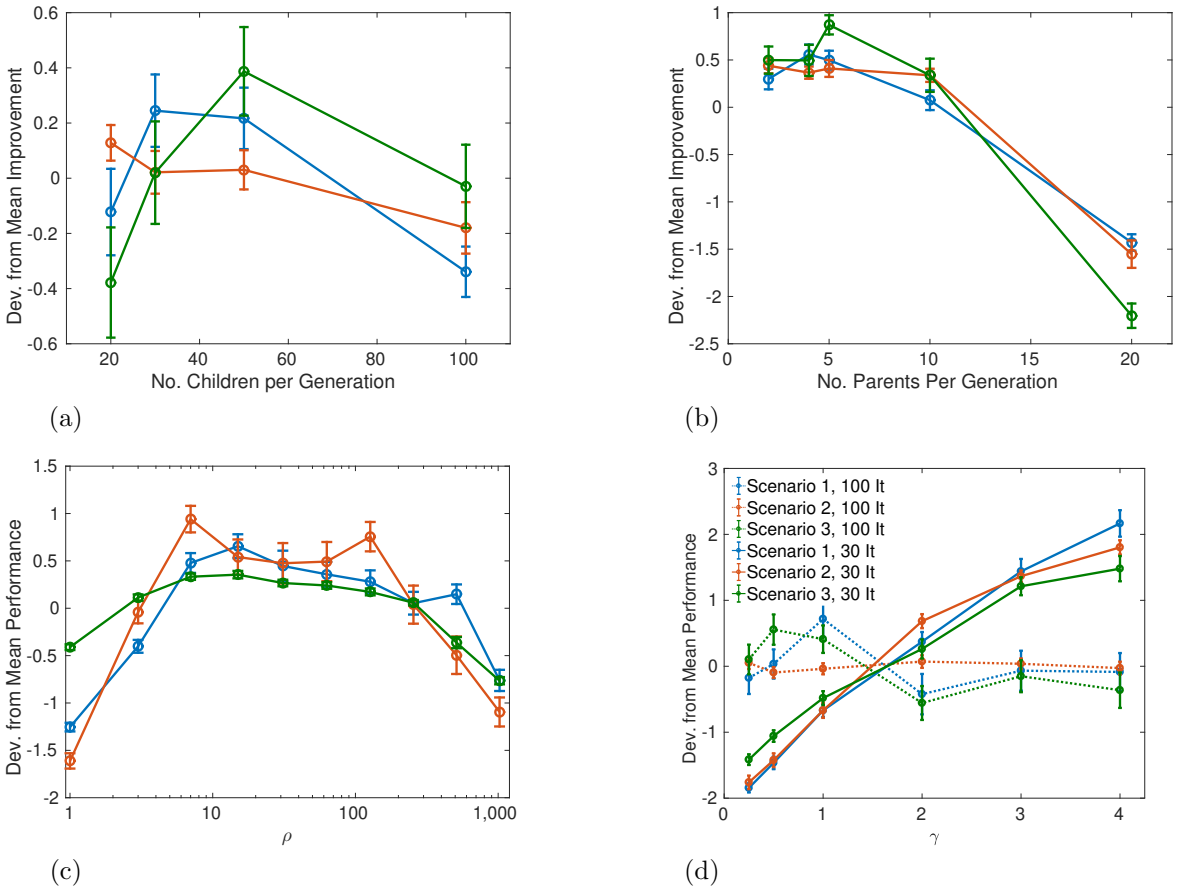

Figure 2.6: The results of numerically simulating the genetic algorithm while adjusting various parameters for three aberration scenarios (red, blue, green), adjusted for the final mean performance across the parameter space (i.e., ± 1 represents a $\pm 100\%$ difference in algorithm performance from the average case). Each plot represents the coupling improvement relative to the starting efficiency using 2000 test configurations (e.g., 100 generations of 20 children). In (a)-(c), the three colored curves represent three different sets of wavefront aberrations to correct. (a) Adjusting the number of child mirror configurations per generation n while keeping the total time constant (total number of test configurations across all generations, where the number of generations is set to $2000/n$) suggests that, given a constant 10 parents selected per generation, the optimal family size is 20-50 children. For all experimental tests we used 20 children per generation. (b) Selecting all 20 of the children to form the next generation (no culling) degrades performance. In order to balance speed and the final optimized coupling, in all other simulations we used 10 parents and 20 children per generation. (c) Simulations suggest ρ may be optimized for the experimental scenario (ρ as plotted has arbitrary units); for these simulations, an optimal value would be $\rho = 5 - 100$. (d) Finally, in the short term (solid lines, 30 test generations), large γ values (the power scaling used to determined the variance between each new generation) are better, though we have observed that the convergence may be unreliable for larger γ (not shown here). In the long term (dotted lines, 100 generations) there was no advantage for large γ. *Reprinted with permission Ref. [42]*

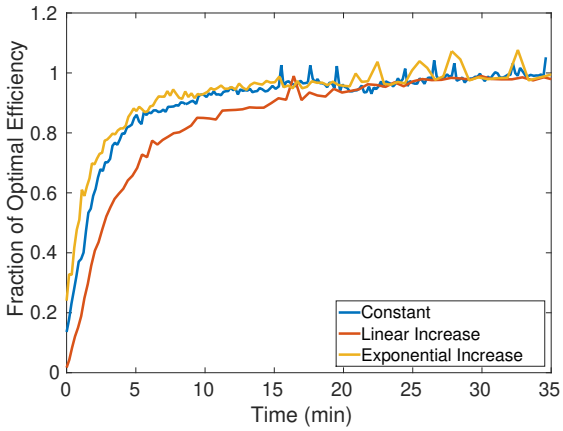

Figure 2.7: Experimental algorithm performance as the number of children per generation is varied (holding the generation size at 20 children per generation, increasing linearly from 20 children per generation, and exponentially increasing from 20 children per generation). Each run is normalized to the final optimized value, which was comparable across all cases. The speed of optimization was worst for the linearly-increasing case. *Reprinted with permission Ref. [42]*

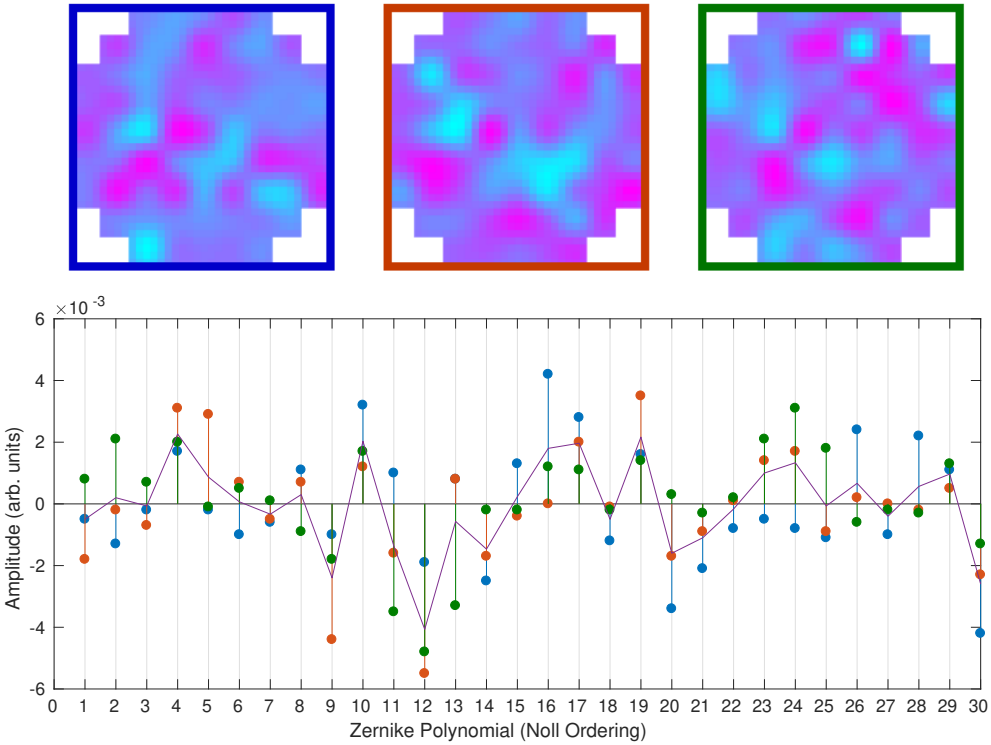

Figure 2.8: Zernike amplitudes after optimizing collection from a single quantum dot several times in succession (optimization data presented in main text). The three runs are presented in blue, green, and red, and the solid line displays the average of the three runs. The corresponding final mirror configurations for each run are reproduced above the chart. Actuator voltages of ±0.02 V are magenta and cyan, respectively. The absolute throw per volt is dependent on the overall membrane tension and therefore cannot be estimated directly, but the peak-to-valley displacement is on the order of a few microns or less. There is significant variation in the optimal configuration for each run, which may suggest the presence of higher-order aberrations which are not fully addressable using our finite-resolution mirror. *Reprinted with permission Ref. [42]*

2.8 Final Notes

In conjunction with a deformable mirror, our algorithm is capable of improving the single-mode fiber coupling of aberrant beams from a variety of sources, including from real-world stationary and drifting light sources. At least 1000 counts per test configuration is desirable for reliable performance; this number may be increased at the cost of overall optimization time (assuming the increase is simply due to longer accumulation times), leading to a reliably higher final coupling efficiency as long as the longer collection time does not approach the timescale of any system drifts. For some single-photon emitter applications, low count rates may increase the overall optimization time to unacceptable levels (e.g., molecules that emit only a finite number of photons before bleaching). One possible solution for some cases is to run the collection system backward, stimulating the emitter with high-intensity light reflected from the deformable mirror and optimizing the counts collected from the subsequent emitter fluorescence over a much larger solid angle, e.g., detected by a camera or large area photomultiplier tube. We have verified that our system performs identically in both the forward (e.g., collection from a pinhole or emitter) and backward (e.g., coupling into a pinhole, or driving an emitter) directions. This is clearly a strategy to be tested *in situ*, but one that could greatly improve the time for optimizing the coupling between the single-photon emitter and a fiber or other single-mode optical element.

Our technique is general and can be extended to other applications, such as side-on or oblique coupling of dipoles. The absolute theoretical coupling into *free space* under certain conditions may be near unity [64], though coupling performance into a single Gaussian mode (SMF) will be optics-dependent and therefore difficult to estimate generally. Indeed, though experiments in near-4π collection from dipole emitters with parabolic mirrors have shown possible experimental collection efficiencies of over 50% into free space [65], such experiments have additional optical limitations that can significantly reduce the measured efficiencies of the complete optical system ($< 25\%$), even without the additional constraint of fiber coupling. At least in principle, the single (but non-Gaussian) free-space mode can be mapped with high efficiency onto a SMF mode using adaptive optics.

Chapter 3

Simulating Turbulence in the Laboratory

3.1 Introduction

Distributing a cryptographic key using single photons over fiber is often impossible (e.g., between ships on the ocean, from Earth to/from a satellite, etc.). Such situations require *free-space* operation, where light is transmitted directly through the atmosphere. This approach introduces a number of complicating atmospheric effects (turbulence, absorption, background, etc.) that may deteriorate the optical signal to the point where quantum key distribution is difficult or impossible.

We would like to form a basic understanding of how the atmosphere affects single-photon transmission and quantum communication protocols; however, it often not experimentally practical to construct an actual free-space link that can generalize to every protocol of interest, and, despite the accuracy obtained by using a real link, the atmospheric conditions in that particular location may not generalize. In this chapter, we discuss a variety of techniques that can be used to simulate the effects of turbulence on the spatial mode of propagating laser light.

3.1.1 The Influence of Turbulence

Neglecting time dependence, depolarizing effects, and backscattering, we can model the propagation of an electromagnetic field $U(\mathbf{r})$ with wavenumber k_0 using the Helmholtz equation:

$$\nabla^2 U + k_0^2 n^2(\mathbf{r}) U = 0. \tag{3.1}$$

Density fluctuations in the atmosphere result in variations in its refractive index $n(\mathbf{r}, t)$ as a function of position and time. In general, $n(\mathbf{r})$ is a random function of position, the moments of which are determined by the strength of turbulence and the assumed model.

To determine the properties of the refractive index, we should first determine the properties of the air flow in the atmosphere. Here we assume the Kolmogorov model of turbulence [66], which is the earliest accurate

model of the statistics of airflow in the atmosphere. The Kolmogorov model is valuable for predicting standard atmospheric behavior within the so-called "inertial subrange" of atmospheric scales – above some small granular parameter l_0 and below the scale of large atmospheric movements characterized by the outer scale L_0. In this regime, optical effects are dominated by the behavior of turbulent eddies that form as cells of sizes near the outer scale subsequently break up into many smaller cells down to the inner scale.

Random fields such as the air velocity (and, in turn, the refractive index) may be characterized by the structure function $D(\mathbf{r}_1, \mathbf{r}_2)$ associated with the relevant physical quantity, which is more convenient (but less powerful) than correlation functions; in general, the structure function can be derived from spatial correlation functions, but not vice-versa. The air velocity field $v(\mathbf{r})$ can be written as the sum of its mean $m(\mathbf{r})$ and small fluctuations $v_1(\mathbf{r})$, and therefore the structure function can be approximated as [67]:

$$D_v(\mathbf{r}_1, \mathbf{r}_2) \approx \langle [v_1(\mathbf{r}_1) - v_1(\mathbf{r}_1 + \mathbf{r})]^2 \rangle. \tag{3.2}$$

For locally homogenous and isotropic fields, the structure function becomes a function of $r = |\mathbf{r}_1 - \mathbf{r}_2|$ alone. In the Kolmogorov model the air velocity is assumed to be locally homogenous and isotropic, and its structure function has the form $D(r) = C_v^2 r^{2/3}$, where C_v is a turbulence strength parameter determined by the magnitude of velocity fluctuations [66].

The power spectrum $\Phi(\mathbf{k})$ (equivalent to a Fourier transform of the structure function) tells us the spatial frequencies of fluctuations to expect in a propagating wave. The power spectrum for the air velocity field assuming Kolmogorov turbulence can be calculated from its structure function by:

$$\Phi_v(k) = \frac{1}{r\pi^2 k^2} \int_0^\infty \frac{\sin kr}{kr} \frac{d}{dr} \left[r^2 \frac{d}{dr} D_v(r) \right] dr = 0.033 C_v^2 k^{-11/3}, \tag{3.3}$$

where C_v^2 is a constant strength parameter for the velocity field [67]. As velocity fluctuations correspond to temperature (and therefore density) fluctuations, it is possible to find the power spectrum for the index of refraction,

$$\Phi_n(k) = 0.033 C_n^2 k^{-11/3}, \tag{3.4}$$

where C_n^2 is a parameter describing the strength of the refractive index fluctuations [67]. It is possible to mimic atmospheric turbulence by reproducing the Kolmogorov power spectrum for optical index fluctuations, as discussed in the next section.

Propagation through turbulence causes the transmitted field to have irradiance fluctuations, also known

$C_n^2(m^{-2/3})$	σ_R^2	$\sigma_R^2\left(\frac{2L}{kW^2}\right)^{5/6}$	r_0 (m)	Turb. Strength
5×10^{-18}	0.13	0.01	0.37	Weak
3.9×10^{-17}	1.00	0.114	0.11	Moderate
1×10^{-16}	2.57	0.29	0.062	Moderate/Strong
5×10^{-16}	12.8	1.47	0.024	Strong

Table 3.1: Channel characteristics for various local turbulence strengths C_n^2, assuming $\lambda = 700$ nm, $L = 30$ km, and $W = 30$ cm.

as scintillations. The Rytov variance for a plane wave is given by

$$\sigma_R^2 \equiv 1.23 C_n^2 k^{7/6} L^{11/6}, \tag{3.5}$$

where L is the transmission channel path length, k is the wavenumber, and C_n is the strength parameter describing the random index of refraction – characterizes the strength of scintillations induced in the turbulent volume; as shown in Table. 3.1, a low value of σ_R^2 corresponds to weaker turbulence and scintillation. These scintillations are also conventionally described by the Fried parameter r_0, defined by

$$r_0 \equiv \left[0.423 k^2 \int_C C_n^2(z')dz'\right]^{-3/5}, \tag{3.6}$$

where k is the wavenumber of the propagating beam [68]. The Fried parameter is a convenient metric for scintillation strength that describes the spatial coherence radius of an optical beam propagating through atmospheric turbulence: when turbulence is strong, r_0 is small (the beam changes rapidly in all transverse directions). Roughly speaking, r_0 is the size below which the air cells do not appreciably distort the beam, i.e., beams smaller than r_0 will not be significantly broken up, though they may still be deflected by the index variations.

By convention, turbulence can be split into weak and strong regimes: if W is the receiver-side beam radius, for the weak regime, $\sigma_R^2 \ll 1$ and $\sigma_R^2(2L/kW^2)^{5/6} \ll 1$, and for strong turbulence, $\sigma_R^2 \gg 1$ and $\sigma_R^2(2L/kW^2)^{5/6} \gg 1$ [69]. Table 3.1 summarizes some of these quantities for a 30-km turbulent channel transmitting a 700-nm wavelength Gaussian beam with a 30-cm diameter beam waist; the transition to medium/medium-weak turbulence for this channel occurs around $C_n^2 = 3.9\times10^{-17}m^{-2/3}$.

One parameter of interest is the probability of *fades*. Given a "bucket detector" inside a receiver with some effective diameter (either from imaging optics or actual size), a fade is an event where the power incident on the detector falls below some threshold value; in other words, a fade is when the signal is lost. Given the statistics of turbulence in the channel, it is possible (though cumbersome) to derive approximate distributions for the probability of fades for a propagating plane wave, beam wave, etc. Here we consider

the simplest case, the plane wave [67].

For moderate turbulence, fades are well-described via the $\Gamma - \Gamma$ distribution, which is the probability distribution governing the intensity of light incident on the detector (the fade probability itself, then, would be the integral over the $\Gamma - \Gamma$ distribution from a chosen threshold value to infinity). Let X and Y be independent random variables describing the so-called "small-scale" and "large-scale" scintillations due to the atmosphere, such that the intensity received is $I = XY$ (see Ref. [67] for additional detail). Define parameters α, β as

$$\alpha \equiv \frac{1}{\exp\left(\sigma^2_{\ln X}\right) - 1}, \beta \equiv \frac{1}{\exp\left(\sigma^2_{\ln Y}\right) - 1}. \tag{3.7}$$

where the $\sigma^2_{\ln X,Y}$ are the variances of the logarithm of the random variables X and Y. The probability of receiving intensity I at the receiver is then given by

$$p_I(I) = \frac{\alpha\beta^{(\alpha+\beta)/2}}{\Gamma(\alpha)\Gamma(\beta)I}\left(\frac{I}{\langle I(0, L + L_f)\rangle}\right)^{(\alpha+\beta)/2} K_{\alpha-\beta}\left(2\sqrt{\frac{\alpha\beta I}{\langle I(0, L + L_f)\rangle}}\right), \tag{3.8}$$

where Γ is the gamma function, L is the channel length, L_f is the distance from the receiver lens to the bucket photodetector, I is the received intensity at the output plane, and K_i is the i-th modified Bessel function of the second kind. σ_I^2 is the scintillation variance – at its simplest just a function of the Rytov variance σ_R^2. For large apertures, the scintillations must be integrated/averaged out over the whole aperture, leading to a dependence of σ_I^2 on the aperture size. $\langle I(0, L + L_f)\rangle$ is the average intensity received at the receiver.

When the turbulence is weak, this distribution takes on a somewhat more tractable form as the log-normal distribution:

$$p_I(I) = \frac{1}{I\sigma_I(0, L + L_f)\sqrt{2\pi}} \exp\left\{-\frac{\left[\ln\left(\frac{I}{\langle I(0,L+L_f)\rangle}\right) + \left(\frac{1}{2}\right)\sigma_I^2(0, L + L_f)\right]^2}{2\sigma_I^2(0, L + L_f)}\right\} \tag{3.9}$$

Both distributions give us a metric by which to evaluate the efficacy of our laboratory turbulence models: we accept a simulation as accurate if the intensity/fade statistics are well-fit by the relevant probability distributions.

3.2 Simulation of Turbulence in the Laboratory

3.2.1 Thin-screen Models

Developing solutions to issues in free-space communication requires some form of laboratory modeling. Modeling light propagation through turbulence perfectly would require a model for the complete turbulent volume. This is computationally and experimentally impractical, so a simpler approach is required. One common simulation technique is splitting the propagation into multiple discrete steps consisting of applying a two-dimensional phase mask to a beam and then propagating the beam a certain distance through empty space [70]. The refractive index $n(\mathbf{r})$ results in a random phase $\phi(\mathbf{r})$ that is distilled into a two-dimensional approximation $\phi(x, y)$ and applied to the propagating wave at each computation step. The Kolmogorov power spectrum specifies how these phase screens should be generated to produce the correct statistics. First, one selects a particular scintillation strength σ_R^2 and determines the power spectrum $\Phi(k)$ that produces it. One generates Gaussian noise for each individual screen in the frequency domain and filters it, so that all of the frequency-domain masks together match the target power spectrum. One then performs an inverse discrete Fourier transform (IFFT) on each frequency-domain mask to generate the corresponding phase mask $\phi(x, y)$. The FFT technique does not accurately reproduce low-order aberrations such as tip/tilt, coma, etc., due to sampling limitations near the origin in the frequency domain, but other data processing techniques can improve this issue without significantly affecting computational speed [71]. A sample phase screen generated by this procedure and the scintillations produced are shown in Table 3.2.

When implemented in a laboratory setting, the physical probe beam will not be as large as would be deployed in practice (typical free-space communication beams are on the order of 10 cm in diameter or more), as this would overfill the devices used to perform the turbulence simulation. Special care must be taken to consider how the true strength of the simulated turbulence scales as the beam is reduced to diameters on the order of centimeters, so that the scintillations in the smaller beam accurately represent the scintillations in the larger beam. We achieve this by matching the scintillation variance σ_I^2 across the two size scales: we first select a target σ_I^2 that matches a real-world turbulence channel of length L_{real}; then, we find the value of C_n^2 that replicates that σ_I^2 over the tabletop path length L_{lab}. For example, if we assume that $\sigma_I^2 = \sigma_R^2$,

$$\frac{C_{n,\text{real}}^2}{C_{n,\text{lab}}^2} = \left(\frac{L_{\text{lab}}}{L_{\text{real}}}\right)^{11/6}. \tag{3.10}$$

$C_n^2(m^{-2/3})$	Two phase mask simulation	CCD Output
0	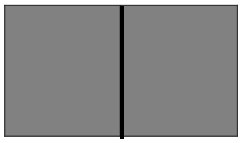	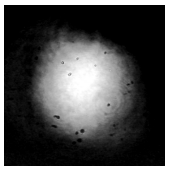
7.7×10^{-18}	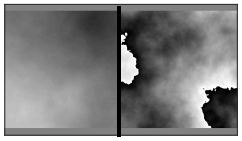	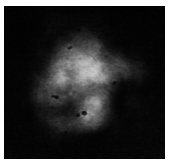
7.7×10^{-17}	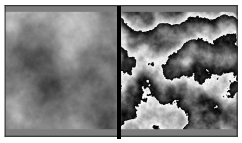	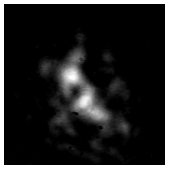

Table 3.2: Sample turbulent phase mask pairs and 700-nm wavelength beam output for a 30-km turbulence channel. The beam acquires the phase from the lefthand mask, propagates some distance (here, $\sim$ 1 m) through free space, acquires the phase from the righthand mask, and then propagates to the collection aperture (in this case, a CCD). Increasing C_n^2 increases the degree of scintillation in the beam.

3.2.2 Phase-only Spatial Light Modulators

Previous work, for example, Ref. [72], has investigated the use of phase-only spatial light modulators (SLMs) to imprint beams with a thin-screen approximation to atmospheric turbulence. SLMs operate through the individual modulation of liquid-crystal pixels aligned in a 2-dimensional grid: light enters the SLM face, propagates through a layer of liquid crystal material (thereby accumulating a controllable polarization-dependent phase, usually in the range $0 - 2\pi$), strikes a backplane mirror and propagates back out through the same layer of liquid crystal material. Each pixel is individually addressed and controllable, so applying a spatially-varying phase to a laser is as simple as modulating the voltage applied to each pixel (provided by an external driver following a PC video interface). Because spatial light modulators retard the phase of the propagating beam directly, they are able to closely mimic thin-screen simulations of turbulence in laboratory settings. They are also not particularly lossy, as they do not rely on diffraction or holography to perform the imprinting of the desired phase distribution. We discuss an application of SLMs for turbulence simulation in the next chapter, but here we briefly describe their merits and drawbacks to provide context to the other methods of simulating turbulence we discuss.

A distinct advantage of SLMs is their high resolution (up to 1920x1080 for the off-the-shelf Holoeye PLUTO spatial light modulator, for example), and large sensor area ($\sim$ 2 cm^2). Large SLMs can be divided

into two halves, with sufficient resolution and space in each half to support multiple passes of a single beam on one SLM face. As discussed previously, the accuracy of thin-screen turbulence models depends on the number of phase screens used. A typical guideline is to select the number of phase screens such that the mean acquires only 10% of its total, end-of-channel scintillation variance σ_I^2 between each phase screen. Though a single SLM can only provide two passes, this is an improvement over only one (provided, for example, by loss-limited DMD-based technology, described in the next section).

A major disadvantage of phase-only SLMs is that they typically only work for one incident polarization, e.g., along the long axis of the SLM face. Phase-only SLMs therefore cannot be used to simulate the impact of turbulence on polarization-based quantum communication protocols, as the turbulence phase profile would only be imprinted onto a single polarization. The atmosphere, however, is not significantly birefringent [73], so simulation of a polarization-based protocol in the presence of turbulence, e.g., polarization-based BB84, requires a simulation device that is polarization-independent.

3.2.3 Digital Micromirror Devices

Digital Micro-mirror Devices (DMDs) are intensity-only spatial light modulators commonly used in digital light projection (DLP) applications. They consist of an array of very small ($\sim$ 10 μm wide) rectangular mirrors which can rapidly switch from one discrete angular position to another (e.g., $\pm 12^\circ$). Incident light may be thereby steered in one of two discrete directions. In DLP applications, the DMD allows for the interleaving of lamps of various colors within the projector to create a per-pixel color variation. In our application, however, the DMD may be used as an intensity hologram which can be programmed on the fly. DLP applications frequently demand polarization insensitivity (e.g., for 3-D films), so DMDs have a high reflectivity for all polarizations (in our tests of a TI DMD, discussed below, reflectivities were $> 99\%$ for H, V, D, and A polarizations).

As described previously, micro-currents in the atmosphere result in small variations in the index of refraction. These variations result in local phase variations in the wavefront of an electromagnetic wave. Phase-only SLMs can simulate these variations by approximating chunks of the channel as thin phase screens; however, intensity-only holograms cannot modulate the phase directly. To convert from an intensity hologram to a phase hologram, we use a technique known as Lee holography [74]. Lee holography uses a simple setup (Fig. 3.1) to transform from an intensity profile to a phase profile. A collimated beam is reflected from the DMD surface; after reflection the beam propagates through a 4-f lens system, with the input plane at the plane of the DMD. At the focal point, each diffractive order from the DMD is easily separable using a beam block (e.g., iris). The first diffractive order is selected by blocking all other orders and allowed to propagate

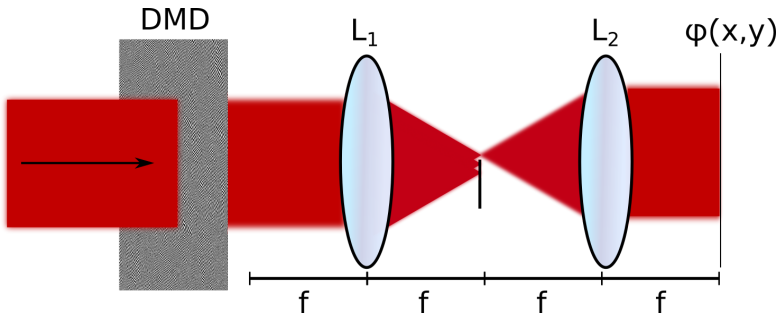

Figure 3.1: Setup for Lee holography using a DMD. Light is reflected (depicted as transmission) from a DMD, which performs an intensity modulation according to Eq. 3.11. The resulting diffraction pattern is sent through a $4-f$ imaging setup. All diffractive orders are filtered except the $+1$ order, which produces the desired phase profile at the output of the optical system.

to the final collimating lens. At the output plane of the 4-f system the DMD has effectively performed the same operation as a phase-only SLM.

The desired phase profile must be modulated before being applied to the DMD according to the relation

$$f(x,y) = (1/2)\left[1 + \cos(2\pi(x-y)\nu_0 - \phi(x,y))\right], \tag{3.11}$$

where $\phi(x,y)$ is the target phase spatial distribution, ν_0 is the blazing frequency used to separate each diffraction order at the focus of the 4-f system, and $f(x,y)$ is the output intensity distribution. Because $f(x,y)$ is continuous and the DMD is binary only, $f(x,y)$ must be converted to a truth value I_{ij} for each DMD pixel via $I_{ij} = \{f(x_i, y_j) > 1/2\}$[75]. This produces an intensity hologram that not only separates the diffractive orders from one another, but also applies a spatial modulation to the ± 1 diffractive orders that results in the desired phase modulation at the output plane.

Results

An example DMD turbulence-simulation mask is shown in Fig. 3.2(a), corresponding to a turbulence strength of $\sigma_I^2 = 0.16$, or moderate turbulence. To verify the ability of the DMD to reproduce realistic turbulence we measured the power collected into a bucket detector (optical power meter) behind a fixed aperture (iris) for a sequence of randomly-generated phase screens. For this and for subsequent tests we used the TI DLP4500 WXGA DMD, with a resolution of 912×1140 and micro-mirror pitch of 7.6 μm. The DMD is driven using a modified TI DLP LightCrafter 4500 evaluation module with all optics and illumination sources removed. The effectiveness of the technique is displayed in Fig. 3.2(b), where PDFs of the statistics of fades into a 1 mm aperture for simulated turbulence at a beam diameter of 0.75 cm agree well with the expected $\Gamma - \Gamma$ distribution. As this work was completed, we became aware of another similar work, Ref. [76], in which

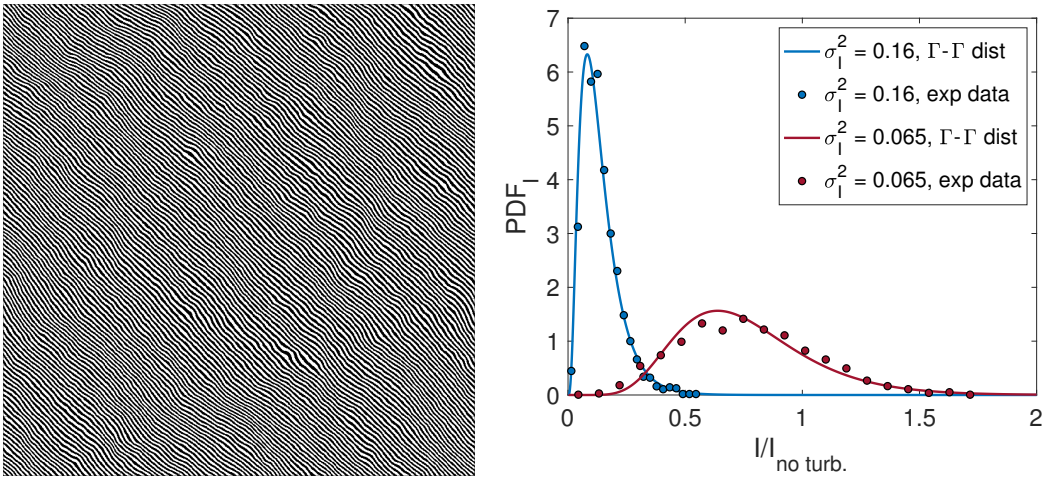

(a) (b)

Figure 3.2: (a) A sample randomly-generated DMD pattern used to generate turbulence. (b) 100 samples of the probability density function of intensity measured into 1-mm aperture (∼0.75-cm diameter beam) for two turbulence strength values, along with theoretical $\Gamma - \Gamma$ distributions. There is a good qualitative agreement between the DMD-simulated tabletop distribution and the expected intensity fluctuation distribution.

DMDs are used for turbulence simulation.

3.2.4 Acrylic Phase Wheels

The two technologies previously discussed are easy to use, adaptable, and accurate. However, they are expensive to implement, as high-resolution spatial light modulators and DMDs for light steering applications can reach into many thousands of dollars for a single unit. Spatial light modulators are much less lossy than DMDs for phase-application purposes as they do not rely on intensity holography/diffraction, but are polarization-sensitive.

Glass phase wheels are a standard technique for simulating thin-screen turbulence. Instead of a reflective hologram, each phase wheel is a transmissive disc with variations in thickness etched into the surface or modulations of the optical index imprinted into the bulk of the disc material. This technique is more similar to the actual process of light propagation through the atmosphere, as the fluctuations in optical index can be continuous over the beam (although, frequently, commercially-available glass phase wheels are digitally etched). Furthermore, each wheel can be rotated to produce time-dependent (though periodic) atmospheric beam distortion. We have investigated an inexpensive, low-loss, polarization-insensitive strategy for generating turbulence in the laboratory using acrylic films. As acrylic is not intrinsically birefringent, the discs would allow us to simulate polarization-based QKD protocols.

We repeated a simplified version of the procedure outlined in Ref. [77]. 12-cm diameter acrylic discs were

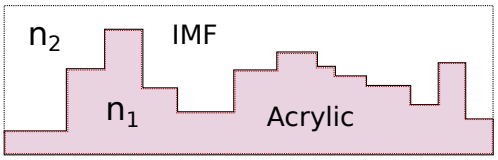

Figure 3.3: Modulating the strength of an acrylic phase screen using index-matching fluid (IMF). The IMF reduces the optical path length variations through the plate, resulting in smaller phase variations and therefore weaker simulated turbulence.

prepared with a mounting hole in the center. Each test disc was coated with a clear acrylic fixing spray (Krylon 1303A), varying the number of total coats over the entire surface from 1-10. The time between coatings was also varied from 10-60 seconds to determine if the setting time of the acrylic fixing spray played a role.

Results

We found that the discs, while promising in some cases (Figure 3.4(a)), were difficult to duplicate, and many of the discs were unusable due to inconsistent coating thickness. Unlike Ref. [77], we did not have access to an automated spray hood; all coatings were performed manually in a painting fume hood. However, the discs can be generated cheaply and rapidly, so the statistics of the turbulence simulated by each disc can be compared until one that approximates the desired distribution (e.g., fade or intensity fluctuations onto a bucket detector) is found. Considering that producing these discs is several orders of magnitude cheaper than purchasing commercially-available glass turbulence simulation wheels, they are a reasonable way of simulating turbulence in the laboratory, despite the uncertainty in fabrication.

A decided disadvantage of these plates is that they cannot change the strength of turbulence they simulate after being fabricated; however, because these plates apply different phases to different parts of beams by varying the amount of physical material in the beam path, we can alter the strength of the plate by applying a phase-matching fluid (e.g., an oil with an optical index close to that of acrylic) to modulate the total phase accumulation through the plate (Figure. 3.3). If the acrylic has index n_a and the index-matching fluid has index n_i, then the total accumulated differential phase across the beam (not including constant offsets over the whole beam) for a plate with thickness variation profile $\delta_a(x, y)$ will be given by $\phi_{tot} = 2\pi(n_i\delta_i + n_a\delta_a)/\lambda$. However, the fluid mates perfectly with the acrylic surface, giving $\delta_i = -\delta_a$. Then, the total phase is given by $\phi_{tot} = 2\pi(n_a - n_i)\delta_a/\lambda$. The index-matching fluid allows for a reduction in the phase applied by the wheel, extending the usefulness of the fixed-strength plates. Also, this method works to reduce the effect of imprecise fabrication tolerances: a $\lambda/2$ optical path length can be produced from a $10\lambda/2$ path length through the acrylic ($n = 1.49$) using a coating layer with index $n = 1.34$. Figure 3.4(b) demonstrates the

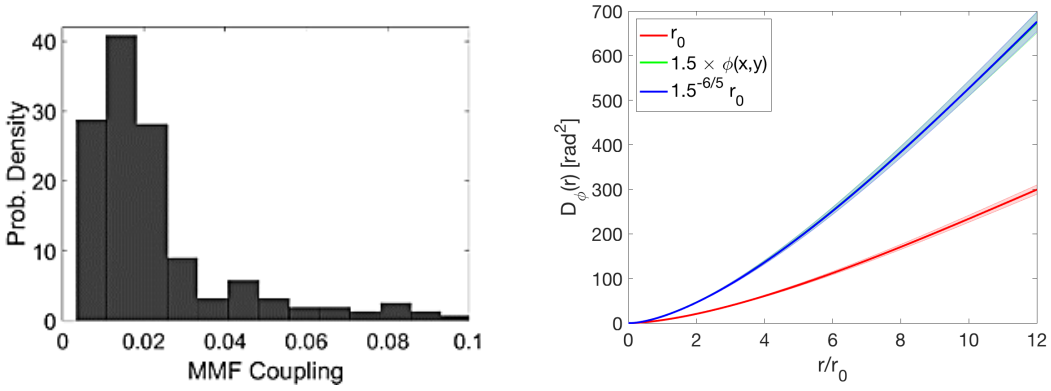

(a) (b)

Figure 3.4: (a) Intensity distribution for light propagating through a single disc-based phase screen, propagating 1 m, and coupling into a multi-mode fiber (MMF). The MMF was used to approximate light clipping from a hard aperture, and the intensity of the light coupled into this fiber was measured using a photodiode. The distribution mimics that of the log-normal distribution, discussed earlier, but further tests are required to verify that such plates can consistently produce the necessary statistics. (b) A study of the phase structure function D_ϕ for a 1.5-m aperture with different assumptions about the underlying turbulence model. The baseline, given in red, is for a Fried parameter of $r_0 = 15$ cm, while the blue line is the result of scaling the original turbulence strength so that the new Fried parameter is $r_0' = 1.5^{-6/5} r_0$. The green curve (covered by the blue curve) shows the effect of simply scaling the phase used in thin-screen models by a constant factor of 1.5. They agree perfectly, implying that we can re-use fixed-strength phase wheels by reducing the differences in optical path length accumulated across each beam (by, for example, using index-matching fluid).

validity of simply modifying the phase by a constant for changing the strength of simulated turbulence.

3.3 Conclusions

We have a presented a variety of effective techniques for simulating turbulence in the laboratory: phase-only spatial light modulators, digital micro-mirror devices, and fixed-strength acrylic plates. Each has their own advantages and disadvantages: SLMs are high-resolution and low-loss, but only function at a single polarization; DMDs work for all polarizations, but require significant loss of light due to the selection of a single diffractive order for operation; acrylic plates are inexpensive, high resolution, and polarization-insensitive, but are difficult to produce consistently and cannot be easily altered without the use of index-matching fluid. Despite their rather different methods of operation, each is capable of accurately modeling weak to moderate turbulence. For the purposes of simulation of turbulence for QKD applications we have applied both SLMs and DMDs in a laboratory setting, discussed in the next chapter.

Chapter 4

Selective Deactivation

4.1 Introduction

As discussed in the previous chapter, light propagating through the atmosphere will develop variations in its intensity distribution as a result of small optical index fluctuations. Over long channels, or in the presence of strong turbulence, a signal beam may not couple efficiently into single- or even multi-mode optical fiber due to significant fluctuations in the spatial mode. For quantum key distribution purposes, turbulence-induced scintillations in the signal beam may drop the collected signal brightness below the acceptable signal-to-noise ratio, resulting in a dropout in the channel or, on average, an increased quantum bit error rate (QBER). Using bucket detectors or wide-core multi-mode fibers with large, high numerical-aperture coupling lenses may improve overall performance, but will introduce additional noise from background light (e.g., solar photons scattered from the atmosphere). Previous work has demonstrated that monitoring channel performance and omitting data where it is suspected that turbulence has degraded the signal can improve key generation rates in poor seeing conditions [78]. In this chapter, we discuss a new technique for operating a multi-mode QKD channel without introducing excessive noise. The technique can enable key generation in turbulence conditions that would otherwise preclude any secure key material from being distributed.

4.2 Selective Deactivation

Selective deactivation is a scheme we have developed for filtering out QKD data which is unlikely to contain 'good' signal photons given a signal beam with a large scintillation variance (strong turbulence). The technique may be applied to free-space QKD channels where the signal beam is significantly aberrated by the atmosphere, or where detector noise or background light leakage is significant enough to cause a degradation in the QKD signal-to-noise ratio (SNR).

An illustration of the problem and our selective deactivation scheme is provided in Figure 4.1(a). Light

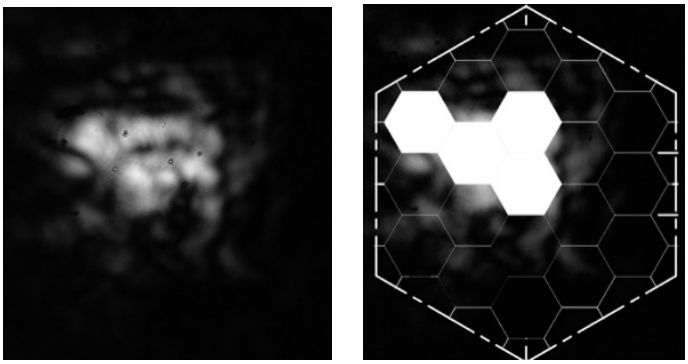

(a) (b)

Figure 4.1: Outline of Selective Deactivation. (a) A signal laser acquires scintillations – or fluctuations in the intensity spatial distribution – at the output of a turbulent channel. (b) By dividing the output plane into smaller subapertures (each associated with its own detector) we selectively examine only channels which are likely to receive a signal photon, thereby improving the signal-to-noise ratio of the QKD channel. Because the specific spatial channels themselves do not result in any information transfer between Alice and Bob (they are aggregated in the QKD postprocessing), this setup does not leak any exploitable information to an eavesdropper.

is transmitted over a long, turbulent atmospheric channel. At the receiver, the signal beam is imprinted with a number of optical scintillations. Efficient collection either requires significant adaptive optical treatment (difficult for long-distance channels where latency can be an issue), multimode operation, or both. If the receiver is truly multimode, it is susceptible to collecting background light from the scattering of sun/moonlight off aerosols or turbulence.

Our solution is to break the large multimode receive aperture into very small subapertures, each with a separate single-photon detector. This allows for collection of all modes into the receiver, at the expense of detector noise. To improve the SNR in the QKD post-processing, a bright classical beacon co-propagates with the signal single-photon beam and is monitored with an imaging sensor to determine which collection detectors are most likely to receive a true signal photon (Figure 4.1(b). Previous work has demonstrated similar techniques, with an emphasis on selecting the threshold that generates the most secure key for a given single-detector channel [79, 80]. Our scheme is a multi-spatial mode extension of this strategy that can allow for the generation of a secure key when single-aperture versions cannot.

4.3 Experimental Design

Initially, we examined the viability of the selective deactivation technique without the polarization analysis required to perform full BB84 QKD. The experimental setup for this initial test is shown in Figure 4.2. We use a heralded source of single photons to estimate the noise characteristics of a standard BB84-style

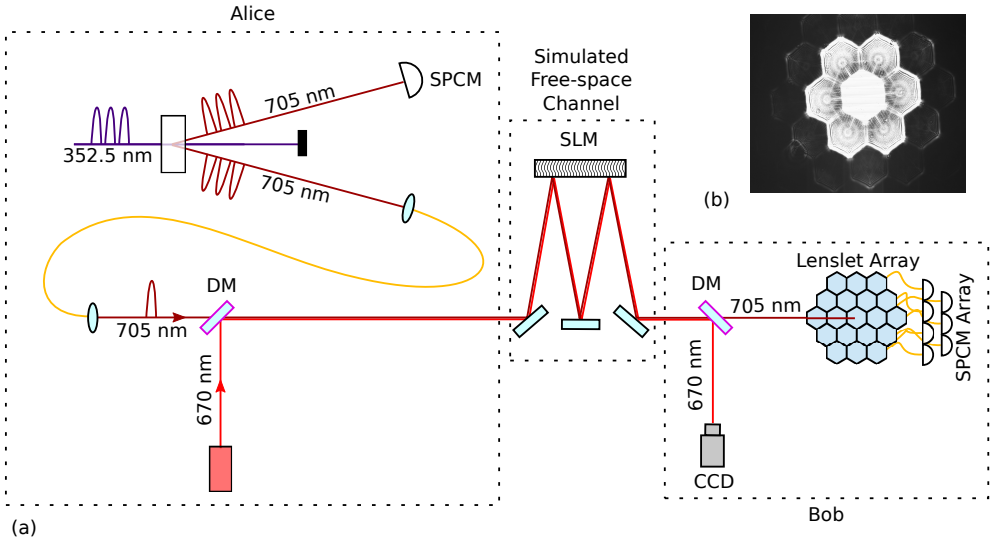

Figure 4.2: Preliminary experimental setup for selective deactivation using a spatial light modulator (SLM). (a) Alice generates 710-nm photon pairs through spontaneous parametric downconversion of a 4-W pulsed 355-nm pump and transmits one photon of each pair to Bob, along with a bright 670-nm beacon. Both the signal and beacon travel through a scaled-down 30-km simulated turbulence path. The beacon is then separated from the signal via a dichroic mirror and imaged with a CCD. Subapertures associated with low intensity regions in the CCD image are consequently omitted from the QKD analysis to improve the signal-to-noise ratio. (b) Image of laser output from the fiber-coupling lenslet array.

protocol in the presence of atmospheric turbulence when using selective deactivation. Though a detailed description of how these entangled photons are generated is somewhat outside the scope of this discussion, we provide a brief explanation of each of the experimental elements.

Entangled photon pairs are generated via spontaneous parametric downconversion (SPDC), a nonlinear optical process in which a high-energy 355-nm wavelength photon (produced by a 120-MHz pulse rate, 5-ps pulse width, mode-locked, frequency-tripled Nd:YAG laser [81]) is converted into a pair of 710-nm wavelength photons (called the signal and idler photon). Downconversion is achieved using a 200-μm-thick bismuth borate (BiBO) crystal (cut at $\theta = 141.8°, \phi = 90°, \gamma = 0/90°$, from Newlight Photonics), a biaxial crystal with a large $\chi^{(2)}$ nonlinearity. For details of the selection of cut and angle, see Ref. [81].

To downconvert 355-nm pump photons, both energy and momentum must be conserved (so-called *phase matching*):

$$E_{\text{pump}} = E_{\text{signal}} + E_{\text{idler}},$$

$$n_{\text{pump}} f_{\text{pump}} \hat{k}_{\text{pump}} = n_{\text{signal}} f_{\text{signal}} \hat{k}_{\text{signal}} + n_{\text{idler}} f_{\text{idler}} \hat{k}_{\text{idler}}.$$

Dispersion within the BiBO crystal makes these impossible to satisfy simultaneously without exploiting additional degrees of freedom in the crystal. BiBO is intrinsically birefringent, however, so phase matching

(specifically, tuning n_{signal} and n_{signal}) can still be achieved by orienting the optical axes of the crystal relative to the desired output photon polarization to counteract the effects of dispersion (BiBO is, in fact, biaxial, meaning it has two optic axes which must be oriented properly to achieve phase matching). For Type-I downconversion, this enables us to generate pairs via the process $e \rightarrow o + o$ (e for extraordinary ray, o for ordinary); for example, a horizontally-polarized pump photon may downconvert to vertically-polarized daughter photons $H \rightarrow V + V$ (or $V \rightarrow H + H$).

By sandwiching together two BiBO crystals rotated 90° from one another, we can simultaneously generate two orthogonal polarizations, with the possibility that there is entanglement between the two photons [82]. For example, in the ideal case, if the pump polarization is set to $|D\rangle = (1/\sqrt{2})\,[|H\rangle + |V\rangle]$, the downconverted two-photon state could be the maximally-entangled polarization state, in addition to wavefunction terms describing correlations in frequency, time mode, etc., as demanded by conservation of energy and momentum:

$$|\Psi\rangle = \frac{1}{\sqrt{2}}\,[|HH\rangle + |VV\rangle] \otimes \int \Psi_\nu(\nu_1, \nu_2)|\nu_1\rangle|\nu_2\rangle d\nu_1 d\nu_2 \otimes ....$$

Because each BiBO crystal has a finite thickness, it may be possible to distinguish between the photons generated in the first crystal in the sandwich versus the second, meaning $|\Psi\rangle$ will not be a coherent sum of $|HH\rangle$ and $|VV\rangle$ terms, and will not be perfectly entangled. This is because there is a slight relative time delay (and therefore a phase) between photons generated in the first and second crystals; alternatively, we can interpret the loss of interference as due to a frequency-dependent birefringent phase shift: averaging over the frequency bandwidth of the photons leads to effective decoherence between the $|HH\rangle$ and $|VV\rangle$ terms. To counteract this effect, we can pre-delay the $|H\rangle$ component relative to the $|V\rangle$ component in the pump polarization (temporal compensation) using a 550-μm piece of beta-barium borate (BBO), a birefringent crystal, to pre-shift the the horizontal and vertical components of the pump so that the $|HH\rangle$ and $|VV\rangle$ terms at the collection optics form the state $|\Psi\rangle = \frac{1}{\sqrt{2}}\left[|HH\rangle + e^{i\phi}|VV\rangle\right]$. This residual phase ϕ can be tuned using another thin birefringent optic, for example a zero-order quarter-wave plate (QWP).

This chain of optics can be seen in the upper left portion of Figure 4.6. The pump polarization is set to horizontal, $|H\rangle$, and rotated using a HWP to form $|D\rangle$. A BBO temporal compensation crystal and QWP correct for temporal walkoff and set the phase ϕ, respectively. The signal and idler photons are filtered (10-nm-wide bandpass filters with center wavelengths 766 nm and 660 nm, respectively) and collected into single-mode fiber (SMF) for use in our QKD protocol.

The single-photon detectors we used in this experiment are silicon avalanche photodiodes (APDs) of various makes and models (SLiK-based Perkin-Elmer / Excelitas SPCM-AQRH-FC SLiK, SAP500-based

MPD PDM APDs with both custom and stock electronics). These devices are able to detect single photons by biasing the detection photodiode beyond its breakdown voltage. When a single photon of sufficient energy (400-1000 nm) impinges on the surface of the diode, it is likely that an electron will be elevated into the conduction band. The intense bias applied to the diode causes this single, free electron to accelerate and excite other electrons into the conduction band, producing an avalanche (hence "avalanche photodiodes"). Fast electronics in the detector module quench the avalanche before the diode overheats, producing a pulse on the output. All silicon APDs are afflicted by dark noise / dark counts, timing jitter, and dead time. The dark noise is the rate of false detection events registered by the detector, which is a result of the finite probability of impurities in the silicon spontaneously producing free electrons that trigger an avalanche (i.e., electrons in impurities may tunnel into the conduction band) [83, 84]. The timing jitter, on the other hand, results from a deviation in the pulse centroid as a result of the readout electronics. The detector dead time is the amount of time that the module electronics must keep the detector off before re-applying the bias, which limits the maximum rate the detectors may fire; it is applied to avoid having a detector re-trigger itself repeatedly due to the abundance of free electrons (so-called "afterpulsing"). The SLiK-based[1] APDs have good all-around dark count characteristics ($<$ 200 cps) and timing jitter (200-300 ps), and dead times on the order of 80-100ns. The SAP500-based detectors, however, have slightly worse dark noise characteristics ($\sim$ 2000 cps), but superior timing jitter ($\sim$ 100 ps). Of at least equal importance, however, is the detector efficiency, which is the probability of recording a pulse event for each photon entering the detector. This is wavelength-dependent, but for the SLiK and SAP500 detectors the efficiency is about 70% and 29% at 700 nm, respectively. The quantum bit error rate in BB84 depends on the ratio of signal photons to noise events detected. As the detectors become noisier due to dark counts or less likely to record a signal photon due to poor efficiency, the bit error rate will increase; however, improved timing jitter better localizes the arrival time of each photon, allowing for superior noise rejection. When the signal count rate is large relative to the noise rate, for example, it may be beneficial to accept increase background noise/jitter and use detectors with a higher efficiency (e.g., an SNR of $C_{signal}/C_{noise} = 0.01 - C$ here representing the count rate – corresponds to a negligble QBER of 0.5%, so a higher-efficiency detector may be used to improve the key generation rate at the expense of additional noise).

Bob's photon is mixed with a bright beacon at 670 nm using a dichroic beamsplitter and transmitted through a double-pass setup using a phase-only spatial light modulator (Holoeye PLUTO), which directly applies the desired $\phi(x, y)$ phase distribution for turbulence simulation. The beacon is then stripped from the signal beam using an identical dichroic beamsplitter, and the beacon is routed to a CCD sensor. The

[1]Super-Low K; k is the so-called ionization ratio, a parameter describing the level of excess shot noise during the detection avalanche. Low k indicates low excessive noise.

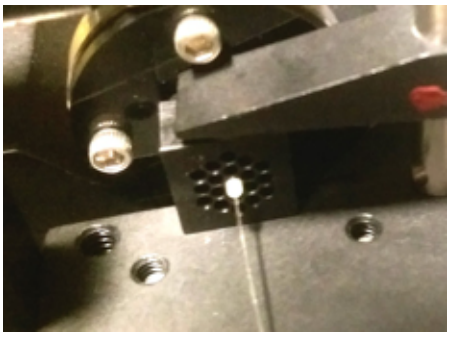

Figure 4.3: Closeup of unloaded fiber array mount, with a single, ceramic-ferruled fiber loaded in the center for scale. The mounted lenslet array is located behind the fiber array mount in this image.

signal beam is routed to a fiber/lenslet array, where each fiber is used as a separate channel for quantum key distribution. The lenslet array consists of a hexagonal grid of micro-lenses with a pitch of 1.5 mm and a focal length of 7 mm (Advanced Microoptic Systems, Figure 4.2(b)). The lenses focus onto a custom-fabricated fiber mount (Figure 4.3), with 19 fibers mounted at a 1.5-mm pitch in a hexagonal pattern. Bright portions of the beacon CCD image correspond to areas of the fiber array where the signal photon is most likely to be measured, allowing us to post-select on channels with a high likelihood of containing true QKD signal photons. Alice's photon functions as a herald for the signal transmitted to Bob; by analyzing the signal-to-noise ratio as a function of the number of subapertures used we can determine if selective deactivation is able to improve the single-photon SNR (and therefore ultimately the QBER and BB84 secure key rate).

4.3.1 Alignment Procedure

Small misalignments in the optical train and the slight mismatch in wavelength between the beacon and the signal photons (670 nm versus 710 nm, respectively) prevent perfect correlation between what is recorded on the CCD and what is coupled into each multi-mode fiber in the fiber array. Therefore, the CCD image of the beacon represents a distorted view of what is truly being coupled into each element of the fiber array. Here we describe a simple adaptive method for matching regions of the CCD image to elements of the fiber array (see Figure. 4.4).

We assume that every fiber corresponds to some distinct region of the CCD image, but we do not know which region that is, and do not even assume the spatial ordering is the same (for example, if the optical system introduced a radially-dependent rotation). First, we note that the phase transfer function for a thin lens is given by

$$\phi(x, y) \sim \exp\left(\frac{2\pi i}{\lambda}\frac{r^2}{2f}\right),$$

so a laser transmitted through a thin lens acquires a radially-dependent quadratic phase [85]. Because we

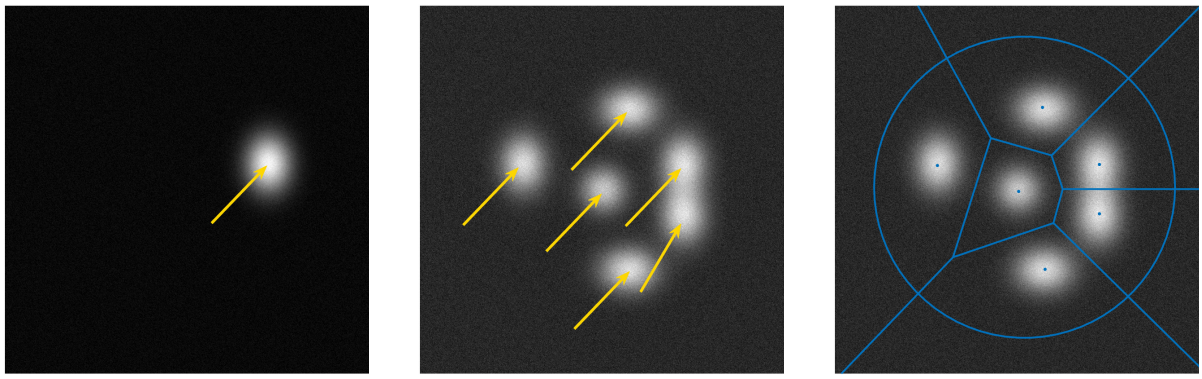

(a) (b) (c)

Figure 4.4: SD system calibration protocol (simulated example) (a) First, a CCD image is captured for the beacon beam for which the signal states are maximally coupled into a chosen element of the fiber array. This roughly corresponds to the region of the CCD that matches that fiber in the fiber array, the center point of which is indicated by the yellow arrow. (b) Images are taken for a range of fiber array elements and overlaid to determine the complete set of centroids corresponding to the complete set of fiber array elements. (c) The centroids are used to calculate a Voronoi diagram, which associates a cell with each point that equally subdivides the distance to each other point. After applying a bounding circle surrounding all centroids to reduce the effect of peripheral regions tending to occupy a larger portion of the CCD, we generate masks from the Voronoi cells which we then apply to turbulent images.

have control over the spatial distribution of the beam phase (using a phase-only SLM), we can control the output angle and focus of our signal and beacon beams by applying the appropriate intensity hologram onto our SLM and shifting f and the center point for r. We then scan f and the center point for r and monitor the signal count rate coupled into each fiber array element. When we have maximized the coupling of our signal beam into a single element of the fiber array, the CCD camera records a snapshot of the beacon spatial distribution. The centroid of this distribution defines the center of the region on the CCD corresponding to the target fiber in the fiber array. Identifying only the centroids for each fiber fails to account for situations where the optical system produces significant higher-order aberrations; however, for selective deactivation to function, the CCD map need only be correlated with the output of the fiber array, and the fiber array need not be exactly mapped to zones of influence on the CCD.

After this procedure is complete for all 19 fiber array elements, we possess a constellation of CCD intensity centroid points, one for each fiber array element. We then compute the Voronoi diagram for the point constellation, which produces a honeycomb-like network of 2-dimensional cells dividing all points evenly [86]. These cells form a set of masks that can be applied to subsequent CCD images (i.e., as an integration region) to determine the likelihood of measuring a signal photon in the fiber array element corresponding to each CCD mask.

4.4 Preliminary Results

We first tested the viability of selective deactivation using a simplified spatial light modulator (SLM) double-pass setup. In Figure. 4.5(a) we present a theoretical study of the impact of the selective deactivation strategy on the secure key rate (SKR). When there is no background noise from detector dark counts or background light, the selective deactivation strategy does not improve the secure key rate (one is always better off using every noiseless detector all the time, as there is no contribution due to noise). However, when the noise level is significant, in this case over 15 kcps per detector, no key may be generated when all detectors are used. The precise noise value where this begins to happen depends on assumptions about the source, such as repetition rate, and the channel, such as overall loss. Generating a secure key is possible when only some subensemble of detectors is used (the intermediate region); there is an optimal threshold (see caption for Figure 4.5), however, as the final secure key rate is slightly reduced when the threshold is too high and only a single detector is used at a time (as many legitimate photons are discarded in this case). The vertical axis – the secure key rate – is calculated based on the quantum bit error rate that one would observe in standard BB84 if we were confronted with the same level of measured background noise; that is, a noise count resulting in a coincidence detection between Alice and Bob will result in an error 50% of the time. The secure key rate itself is calculated using the standard BB84 asymptotic key rate, $r = 1 - 2h(Q)$, where Q is the bit error rate, and h is the binary Shannon entropy.

4.5 Full QKD demonstration

We have demonstrated that selective deactivation improves the SNR of a heralded single-photon source, with one photon transmitted through simulated turbulence. Our current effort is in reproducing these results with an actual polarization-based BB84 system. Here, we report our progress on a complete QKD demonstration of selective deactivation: instead of simply heralding Bob's photon using a paired photon sent to a single-photon detector on Alice's side of the channel, we use polarization-entangled photons to implement a standard BB84-style protocol.

The full setup is presented in Figure 4.6. Non-degenerate spontaneous parametric downconversion of a 355-nm pump produces daughter photons at 660 nm (Bob) and 760 nm (Alice). Here, we changed to non-degenerate downconversion (unequal daughter photon energies) so that Bob's photon has a wavelength compatible with the 400-700-nm operating range of our digital micromirror device (DMD). Because we are implementing the full polarization protocol, our single-polarization turbulence simulation setup using two passes from a phase-only SLM must be replaced with a single pass through a lossy – but polarization-

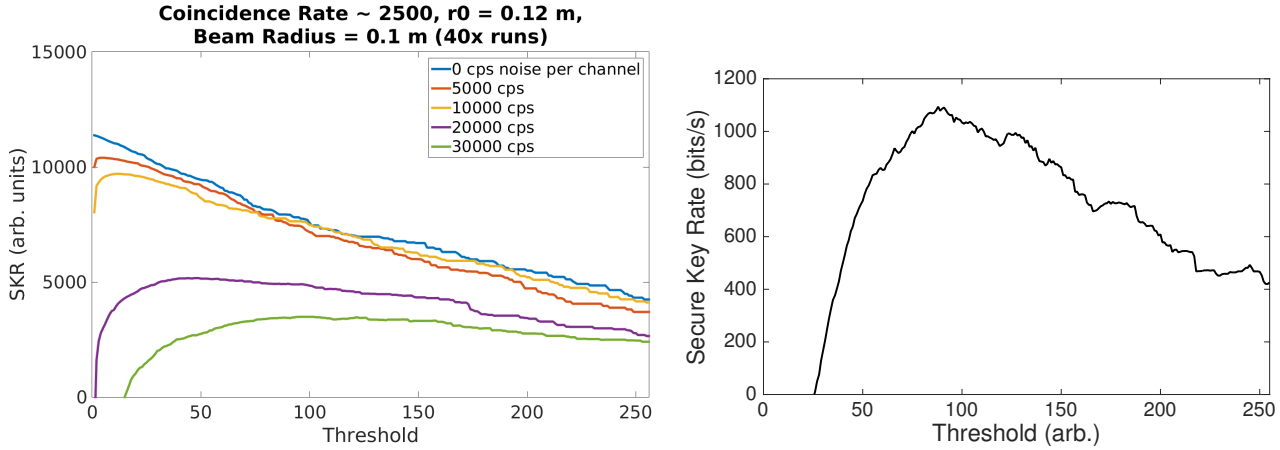

(a) (b)

Figure 4.5: (a) Estimated key rate from a simulation of selective deactivation using 19 fibers in the presence of moderate turbulence. The simulation was performed by generating a random thin-screen turbulent channel, applying the selective deactivation CCD masks to the scintillations produced after propagation of a beam using multi-step numerical Fresnel integration (see, for example, Appendix B), and calculating the signal-to-noise ratio assuming a total coincidence rate of 3500 cps between Alice and Bob. Repeating this procedure for 40 samples of the turbulence channel and scanning the CCD threshold below which detectors will be omitted from QKD analysis gives the results shown here for various noise values. The vertical axis units are presented as arbitrary, as they depend on assumptions about the source repetition rate (in this case, 120 MHz), but will scale linearly with this arbitrarily-chosen parameter for simulation in the asymptotic (non-finite) key limit. Intermediate threshold values improve the secure key rate over both keeping all detectors active (low threshold) and removing all but the brightest detector (high threshold). The threshold values are presented as a grayscale value, with 0 being completely dark (all detectors selected) and 255 being completely bright (only a single detector selected). Though there are only 19 discrete values of the threshold at which a detector is activated or deactivated, these thresholds are not known *a priori* for a given turbulence mask, so we must scan through all greyscale values to find the optimal threshold. The optimal threshold may be selected by forming an average of all samples. (b) Experimental results for $C_n^2 = 7.7 \times 10^{-17} m^{-2/3}$, with a noise rate of approximately 15 kcps per detector. We observe qualitatively similar features to the simulations in (a), with no secure key generated when all detectors are used (too much noise), some key generated when a single aperture is selected at a time, and an optimal rate when some subensemble of detectors is selected. As in (a), the "secure key rate" here is calculated assuming the standard BB84 secure key rate, though no actual key can be generated with this preliminary experimental setup, as it operates on a single polarization state (i.e., the SLM-based turbulence simulator only applies turbulence to horizontally polarized light).

independent – DMD setup (discussed in Chapter 3).

Alice's detectors may be seen in the upper right of Figure. 4.6. Optical fibers apply a unitary rotation to the polarization state of light propagating through the fiber due to stress-induced birefringence in the fiber and the accumulation of a geometric phase. As a result, after traveling through a SMF, the polarization state of Alice's photons will not be in the same state as when they entered (i.e., the global Alice-Bob polarization entangled state will be different), so we correct the polarization launched from the fiber using a HWP/QWP pair.[2]We then use a 50:50 beamsplitter (BS) to randomly choose between two measurement bases, H/V (reflected) and D/A (transmitted). After transmission, a HWP rotates the polarization by 45°, so that D/A may be analyzed using a polarizing beamsplitter (PBS), which transmits only horizontal polarization and reflects only vertical polarization with contrasts on the order of $500:1$. In this way Alice can measure both polarizations simultaneously in each basis, while selecting the basis randomly. She can verify which basis she "chose" in post-processing by looking for detector clicks in each respective basis's detectors.

On the lower arm of Figure. 4.6, Bob's photon is mixed with a bright beacon at 635 nm (with power $\sim 500\mu$W) using a dichroic beamsplitter (Semrock Brightline 649-nm edge FF649-Di01-25x36) and transmitted through our Lee holography setup, as described above. 10-cm focal length lenses (L_1 and L_2) separate the diffractive orders from the DMD and allow us to select only the first diffractive order using an iris. The remaining beam acquires the desired $\phi(x, y)$ phase distribution for turbulence simulation at the output plane of the Lee holography setup. The 635-nm beacon is then stripped from the signal beam using an identical dichroic beamsplitter, and the beacon is routed to a CCD sensor. The signal beam is routed to a fiber/lenslet array, where each fiber is used as a separate channel for quantum key distribution. Bright portions of the beacon CCD image correspond to areas of the fiber array where the signal photon is most likely to be measured, allowing us to post-select on channels with a high likelihood of containing true QKD signal photons. In a true implementation we would need a separate QKD polarization analysis system for each of the 19 channels; here, we simply use a *single* polarization analysis system, which we then switch from each of the 19 fibers coming from the array to the next and aggregate the data for each fiber after the experiment has concluded. One extreme advantage of our "programmable" turbulence emulator is that we can precisely repeat a given turbulence pattern as we cycle through each of the 19 fibers. In practice, it was more efficient to use the same set of turbulence screens while collecting QKD data for each of the 19 collection channels individually.

[2]The collection fiber is fastened to the optical table to fix the polarization rotation applied by the fiber. The downconversion light is set to a known state, e.g., $|HH\rangle$, by adjusting the pump polarization (e.g., to $|V\rangle$), and both Alice and Bob's calibration waveplates are rotated and tilted until H is detected on both with high probability. The waveplates do not need to be adjusted from measurement to measurement as long as the fiber positions remain fixed

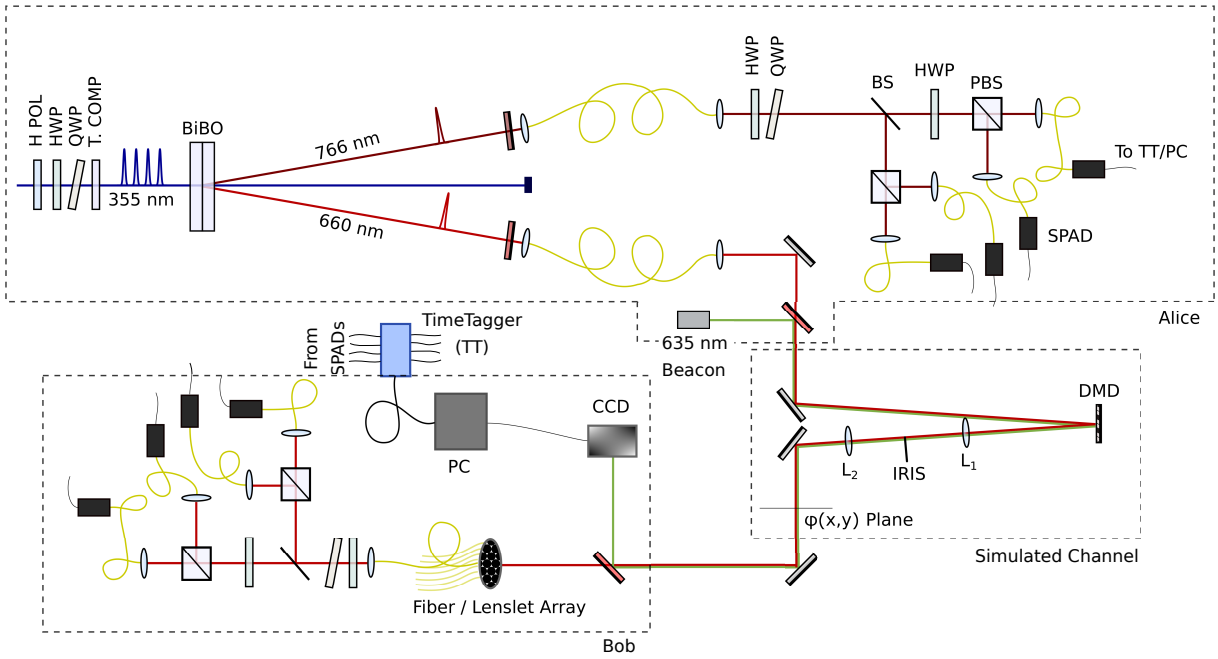

Figure 4.6: Experimental setup for the full selective deactivation QKD protocol. Alice generates 660-nm and 766-nm photon pairs through spontaneous parametric downconversion of a 4-W pulsed 355-nm pump (120-MHz repetition rate) and transmits the 660-nm photon of each pair to Bob, along with a bright 635-nm beacon. Both the signal and beacon travel through a scaled-down 30-km simulated turbulence path implemented using a DMD, as described previously. The beacon is imaged with a CCD, and subapertures associated with low intensity regions in the CCD image are omitted from the analysis to improve the BB84 error rate. Both Alice and Bob perform polarization measurements of their respective photons to complete the remainder of the standard BB84 QKD protocol.

4.5.1 Modified Alignment Procedure

The polarization-insensitive, DMD-based turbulence simulation setup is quite lossy ($\sim 3-6\%$ transmission). It is therefore not practical to calibrate each region of the CCD by brute-force scanning of the signal beam and beacon and observing when the collected photon counts are highest due to the increased integration times and background noise. Here we consider a slightly different approach that is able to recover the CCD centroid corresponding to each collection fiber array element without scanning.

We generate a sequence of strong turbulence screens ($r_0 \ll r_{\text{beam}}$) so that, at the CCD and fiber array, the beam is quite diffuse with significant random variations in intensity. For each turbulence screen j we record the CCD output $I_j^{\text{CCD}}(x, y)$ and the count rate from the fiber i under test $c_{i,j}$. The $c_{i,j}$ are measured in coincidence with the paired daughter photon on Alice's side of the experiment to reduce noise. Random scintillation patterns with regions coupling well into fiber i will result in large $c_{i,j}$, while regions with poor coupling will result in a low $c_{i,j}$. We then examine the composite image

$$I_i(x, y) = \sum_j c_{i,j} I_j^{\text{CCD}}(x, y)$$

and create a mask by thresholding: $M_i(x, y) = \{I_i(x, y) > 0.95 \max [I_i(x, y)]\}$. An example of such a thresholded image is shown in Figure 4.7; the output is used as a mask for selective deactivation instead of the Voronoi technique used previously. An advantage of this technique over the raster scanning method discussed previously is that it queries all areas of the CCD simultaneously and is therefore faster, at the expense of a introducing significant noise to the CCD masks used for the selective deactivation protocol. It is possible with this scheme for the masks to overlap, while in the Voronoi scheme it is not (the Voronoi diagram by definition evenly divides all centroids into cells, with edges equidistant from each centroid). This allows for ambiguity about which detector is likely to fire and allows for a smoother transition between detector activation/deactivation.

4.5.2 Results

We have demonstrated the ability to generate DMD turbulence screens that closely mimic realistic atmospheric turbulence (Chapter 3), and are improving our source to provide a suitable two-photon state for realistic quantum key distribution. Current progress toward the full QKD demonstration of selective deactivation has focused on optimizing the source of entangled photons and the efficiency of the DMD-based turbulence simulator. We can partially characterize our source and BB84 polarization analysis system via the visibility V, a rough measure of entanglement quality, with $V = |(C_{DD} - C_{DA})/(C_{DD} + C_{DA})|$; the

Figure 4.7: Example weighted CCD output for fiber #9 using 50 screens. The bright region corresponds to the region of the CCD most correlated with coincidence counts into fiber #9.

	H	V	D	A
H	1851	58	1097	1055
V	78	2081	1284	1381
D	606	506	190	1129
A	681	582	1301	256

Table 4.1: Coincidence counts over 5 seconds between each of Alice and Bob's detectors. For the state $(1/\sqrt{2})[|HH\rangle - |VV\rangle]$, the HH/HV measurements must be diagonal, while the DD/DA measurements are antidiagonal. The off-diagonal 2×2 subtable values should be equal, as they correspond to the probability of measuring a photon in either state of the conjugate basis.

C_{XY} are the coincidence counts for polarization analyzers X and Y. For the maximally-entangled state $(1/\sqrt{2})[|HH\rangle \pm |VV\rangle]$, the DD/DA visibility $V = 1$. We are able to generate photon pairs in the state $(1/\sqrt{2})[|HH\rangle - |VV\rangle]$ with a $|DD\rangle$ to $|DA\rangle$ visibility of 90%, which corresponds to a quantum bit error rate of 5%. This reduced visibility is due to a mismatch between the temporal compensation crystal (designed for the degenerate downconversion case), and the temporal walkoff observed in the non-degenerate case. When we incorporate the entire polarization analysis system, including fibers, polarization correction waveplates, and detectors (as shown in Figure 4.6), this visibility drops to 70% in the D/A basis due to the D/A basis-choice waveplates on both Alice and Bob's side not precisely matching the wavelengths of the signal/idler photons. This is summarized in the confusion matrix reproduced in Table 4.1, which shows the coincidence counts accumulated between each of Alice and Bob's detectors over a 5-s interval. In fact, the HH/HV visibility is also degraded (93% when analyzed by the full detector system, versus $> 97\%$ for a direct tomography of the source), implying that there are additional errors introduced by the polarization analysis system that must be eliminated before the full QKD demonstration of selective deactivation can proceed. Better matching of TC optics / waveplates will result in an analysis system suitable for a realistic QKD demonstration.

4.6 Conclusions

We have demonstrated a novel technique, selective deactivation, that is able to extract secure key material from a turbulent channel during post-processing. By monitoring signal disturbances using a classical beacon, we are able to select regions of the receiver aperture that are likely to detect single-photon signal states transmitted from Alice to Bob for QKD. This enables multi-mode operation without introducing noise from each of the modes due to background light and detector dark counts. Recent work has discussed the possibility of selecting the CCD masking threshold for a single pixel *a priori* [80]; it would be of interest to extend this work to our multi-aperture case to avoid significant post-processing, as the classical post-processing for QKD is frequently a bottleneck for high link performance. Nevertheless, in our demonstration we are able to extract the optimal key during post-processing by scanning the CCD threshold for each turbulence sample, resulting in an overall improvement in the key rate and, in some turbulence situations, distribution of a secure key when it would otherwise be impossible. Note that, when deployed in a realistic setting, the data used for thresholding would be set aside from the main QKD data, as Alice and Bob must perform multiple rounds of error rate estimation, possibly leaking information to an eavesdropper. The threshold-selection procedure would then repeat when the turbulence strength in the channel changes.

Chapter 5

Drone to Drone Quantum Key Distribution

5.1 Preliminaries and Motivation

Many systems have demonstrated quantum key distribution between mobile platforms, for example between a ground station and ground vehicle [28], ground station and aircraft [87], and satellite and the ground [29]. Additionally, a great deal of effort has been focused on miniaturizing QKD payloads for use in handheld applications [88, 89, 90]. However, in the former case, each setup is able to leverage elaborate ground stations with heavy, sophisticated telescopes and pointing and tracking (PAT) systems, while, in the latter case, the propagation distances are assumed to be quite small (less than a meter), so ultra high-performance PAT systems are not required.

Flying platforms are of particular interest for free-space quantum key distribution as they can avoid many of the atmospheric effects that plague low-altitude free-space channels. Turbulence strength, for example, drops rapidly with altitude (Figure 5.1) [91], and transient weather events such as fog or aerosols are typically located close (< 100 m) to the ground. Flying platforms can avoid the worst of these issues by simply transmitting above the layers where they are typically strongest. To that end, in this chapter we discuss our current efforts toward moving quantum key distribution systems aboard flying multi-rotor unmanned aerial systems (UAS), or drones. This work is a joint effort between our group at the University of Illinois (focusing on the PAT optics and the drone system itself) and Prof. Daniel Gauthier's group at The Ohio State University (focusing on the quantum key distribution signal sources).

The choice of drones, while catchy, has several appealing experimental upsides. For one, they are commercially available and relatively inexpensive. Secondly, their size, weight, and power (SWAP) requirements are not unlike what would be found on another QKD platform of interest: micro-cube satellites, or CubeSats. Such satellites are deployed cheaply and rapidly, but require months to years of development and are nevertheless quite costly, despite their size. Drones offer a way to address many of the same issues: compact, indistinguishable sources, robust PAT, and miniaturized optics, all within a suitable SWAP budget.

Furthermore, drones themselves are a useful platform for quantum key distribution. Such a system

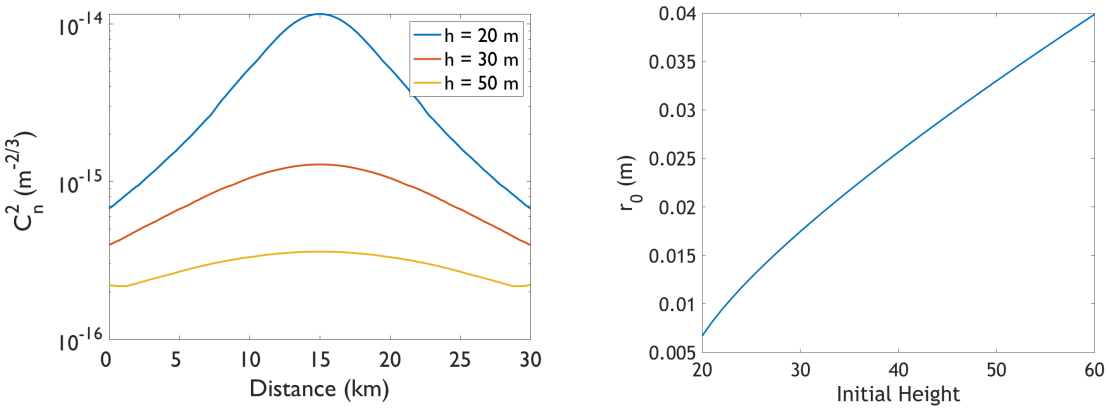

(a) (b)

Figure 5.1: Turbulence scaling as a function of height for a 30-km channel. (a) The turbulence strength parameter C_n^2 increases by an order of magnitude near the surface (data based on measurements by the Space and Naval Warfare Systems Command, or SPAWAR). When propagating from one terminal to another over long channels, the curvature of the Earth will bring the signal beam into this high-turbulence region. This effect is significantly reduced by launching from a height of 50 m versus 20 m. (b) The Fried parameter r_0 for a 670-nm beam propagating 30 km over the ocean as a function of starting height. Low starting heights correspond to small r_0, or a very small radius of spatial autocorrelation (significant scintillation).

must be small, as even the largest consumer-level drones have a maximum payload capacity of 5-10 kg. A QKD-enabled drone may then be a portable node useable by anyone in urban or rural environments, thereby democratizing a technology that is currently focused on industries with high-end cryptographic needs (e.g., banks, elections [92]). Quantum key distribution, despite its many flavors, only operates point-to-point. But drones equipped with both a receiver and transmitter could form a secure node in a larger network of drone-based QKD systems that can reconfigure as desired, as the transmitters can be relocated by individuals rather than remain aboard heavy aircraft or affixed to fiber networks. Perhaps even more important, developing a QKD-capable drone network is a necessary step toward a full quantum network based on distributed entanglement – the ability to faithfully transmit the various qubit states used in QKD protocols is sufficient to faithfully transmit entangled states.

5.2 Three-state protocol

Silicon single-photon detectors are heavy and power-hungry, with the smallest off-the-shelf commercial models weighing in at a minimum of 300-500 g and consuming 0.5-2 A of power continuously. Furthermore, bulk optics are heavy and require significant space on both the transmitter and receiver optics benches for the creation and detection of 4 polarization states (H, V, D, A) with high fidelity. Previous work, described in Ref. [93], has shown that Alice does not actually need to transmit all four states to Bob to achieve the full

secure bit rate of standard BB84; in fact she need only transmit 3; recent work has shown that this result extends to higher-dimensional states [33]. For our purposes, we would like to know if it is plausible, at least from an information theoretic perspective, to generate a secure key by transmitting *and* receiving only three of the four input states; for example, Alice measures/transmits R, L, and D while Bob measures R, L, and A, where R, L are the right- and left-hand circular polarizations $|R/L\rangle = (1/\sqrt{2})|H\rangle \mp i|V\rangle$. Here we use the circular polarization as the data basis because the circular states are invariant with relative rotations of the two drones; therefore, the relative drone motion does not contribute any error to the circular states. The single diagonal state used by each party acts as a check for the eavesdropper. We may anticipate, however, that the bound on the information leaked to an eavesdropper becomes more difficult to estimate when we are not able to make complete measurements in the diagonal basis.

To determine the effect of using only three states for both the transmitter and receiver, we mimic the approach of Ref. [33]. This work casts the task of estimating the information leaked to an eavesdropper as a convex optimization problem: let $\rho_{\mathcal{AB}}$ be the density matrix of the joint state shared by Alice and Bob (for this purpose we can assume that we are using an entanglement-based protocol, which is formally identical to the prepare-and-measure protocol actually implemented). Let $\Pi_n^A = |\alpha_n\rangle\langle\alpha_n|$ and $\Pi_n^B = |\beta_n\rangle\langle\beta_n|$ be operators describing the possible measurements, where A, B represent a basis choice (e.g, R/L and D/A), and α_n, β_n represent the individual operators within this basis, e.g., $\alpha_0 = R$. Let E_A and E_B be the error operators in the A, B bases, respectively, defined by $E_* = |*_0\rangle\langle *_1| + |*_1\rangle\langle *_0|$ (here $*_0$ and $*_1$ are placeholders for the first and second elements of basis $*$). Given an error rate in the A basis of e_A, we would like to maximize the error e_B achievable in the conjugate basis B using all possible $\rho_{\mathcal{AB}}$. This can be achieved using the following convex programming problem, as described in Ref. [33]:

$$\texttt{max} : \mathrm{Tr}(E_B \rho_{\mathcal{AB}}) = e_B$$

subject to the constraints:

$$\mathrm{Tr}(\rho_{\mathcal{AB}}) = 1,$$
$$\rho_{\mathcal{AB}} \geq 0,$$
$$\mathrm{Tr}(E_A \rho_{\mathcal{AB}}) = e_A,$$
$$\mathrm{Tr}(\Pi_n^a \otimes \Pi_m^b \rho_{\mathcal{AB}}) = p_{n.m}^{a,b},$$
$$\forall\{a, b\} \in \{A, B\}; n, m = 0, 1.$$

Here, the error rate e_A would be analogous to the bit error rate in the data basis (L/R for us), while e_B would be the (unknown) error rate in the check basis (D/A). Intuitively, this procedure estimates the most

error an eavesdropper can introduce in one basis (corresponding to an eavesdropper's maximum information gain) without being detected in the conjugate basis, optimized over all operators she could use according to quantum mechanics. All parameters except e_B are determined experimentally. From this optimized value of e_B the secure key fraction may be calculated via $K = 1 - h(e_B) - \text{leak}_{EC}$, where h is the binary entropy and leak_{EC} is the information leaked during classical postprocessing. This maximization problem may be solved numerically using standard convex programming packages (e.g., CVX for MATLAB). It is worth noting that this procedure provides a quantum information theoretic bound; whatever experimental system actually implements this protocol may introduce additional complexities that further influence the amount of secure key material generatable by the system.

The result of this calculation for three input and output states is shown in Figure 5.2(a). The secure key fraction for a given quantum bit error rate (QBER) is significantly smaller for the 3-state case than the standard 4-state case. For QBER values over about 5.7%, no secure key may be produced; for the case of 3 states prepared and 4 measured the full 11% QBER bound is regained [93]. Figures 5.2(b,c) show estimated link budgets for a drone-to-drone system assuming various aperture sizes, perfect PAT, and the limitations of our electronics.

5.3 System Architecture

An outline of the general structure of each node is shown in Figure 5.3. A large commercial film drone airframe (DJI S1000+) carries a flight controller (Pixhawk2.1), payload electronics, and onboard computer (Raspberry PI 3B+). The airframe is capable of a total takeoff weight of 25 kg; at this weight the drone consumes approximately 4 kW of power during takeoff and 1 kW during basic hovering operations. This necessitates the use of extreme lithium-polymer (LiPo) batteries, which operate at 25.2 V (GensAce Tattu 16000 mAh) and are capable of providing up to 240 A continuously, though during typical operation this does not exceed 40-80 A. These batteries make up a significant portion of the final payload weight at a mass of about 2.5 kg. There is a tradeoff: despite the extra weight, larger batteries result in a longer flight time; however, if the total takeoff weight exceeds the maximum that can be provided by the airframe then the drone will not be able to leave the ground.

Figure 5.4 depicts a general outline of the interaction of the two drones with one another and their environment. Each drone has a flight controller, which maintains orientation of the entire body dynamically (the "brain"). The onboard computer (Raspberry PI), communicates with the flight controller to maintain drone position indoors using a commercial motion-tracking system. The onboard computer also controls the

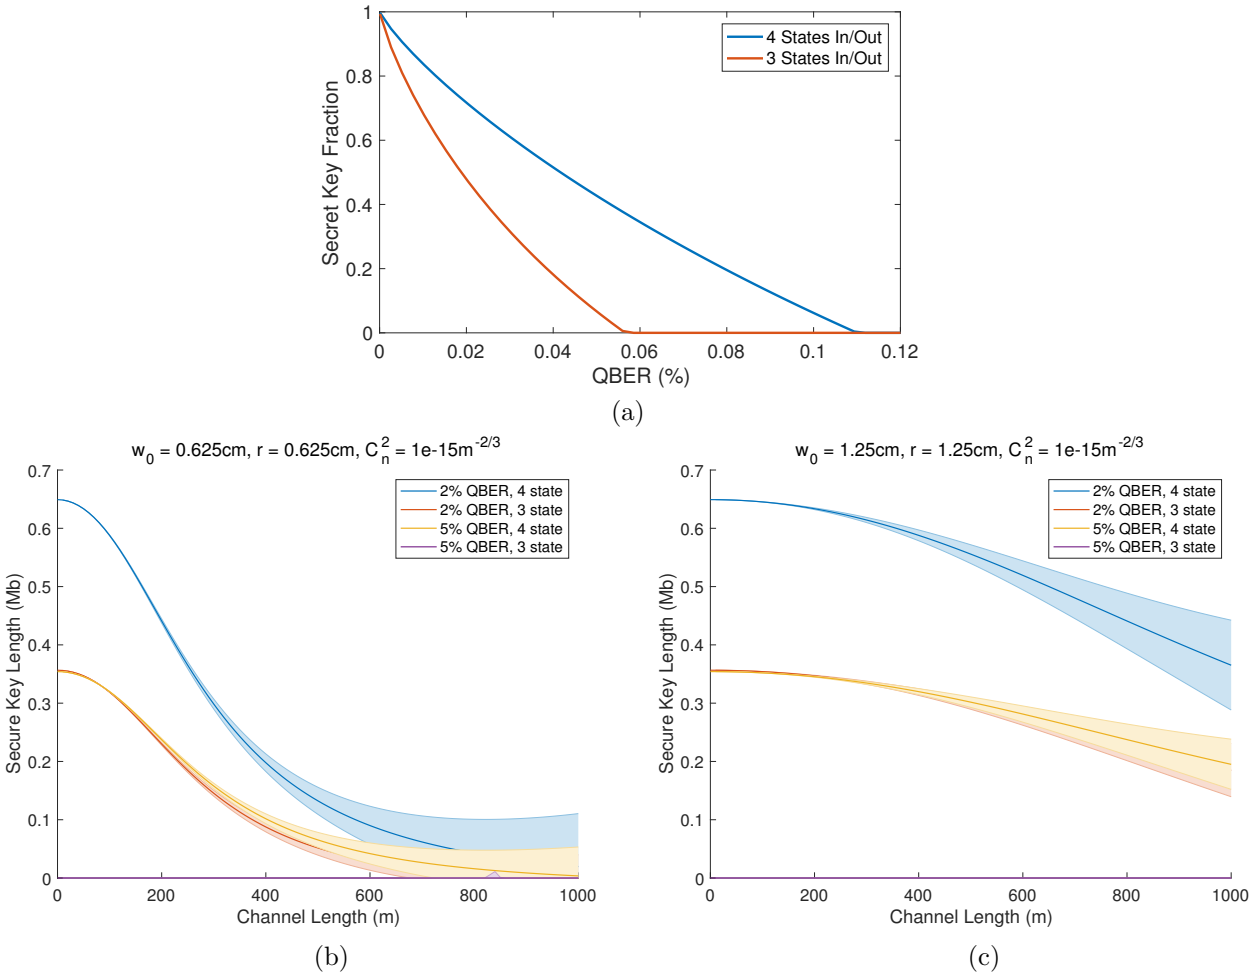

Figure 5.2: Link budgets for 4- and 3-state protocols. (a) A comparison of the best-case 3-state secure key fraction versus the standard 4-state protocol. For error rates above 5.7%, no secret key may be produced as too much information may have been gained by an adversary. (b) Estimated secure key length generated by a 100-s run with $\lambda = 656$ nm, $C_n^2 = 1 \times 10^{-15} m^{-2/3}$, detector efficiency of 37%, repetition rate of 50 MHz, mean photon number (at the receiver) of $0.025 \times 10^{-9.87/10}$, or 0.025 photons per pulse experiencing -9.87 dB loss, assuming finite statistics. The photon number distribution per pulse is Poisson distributed about the mean. The error bands represent the best turbulence case or worst case occurring for the entire run, not a variance based on the turbulence statistics. We also assume 50 kcps noise counts per detector and a 50:50 basis choice split. The input and output apertures are both 12.7 mm in diameter here. (c) The same, but with 25.4 mm-diameter apertures on the transmitter and receiver. The transmission at longer distances is much better due to decreased diffractive loss. For both aperture sizes, the 3-state protocol generates almost no secure key at $Q = 5\%$, limiting its viability for low-visibility channels. For channels where the bit error rate is low, however, it is a possible solution for reducing SWAP consumption.

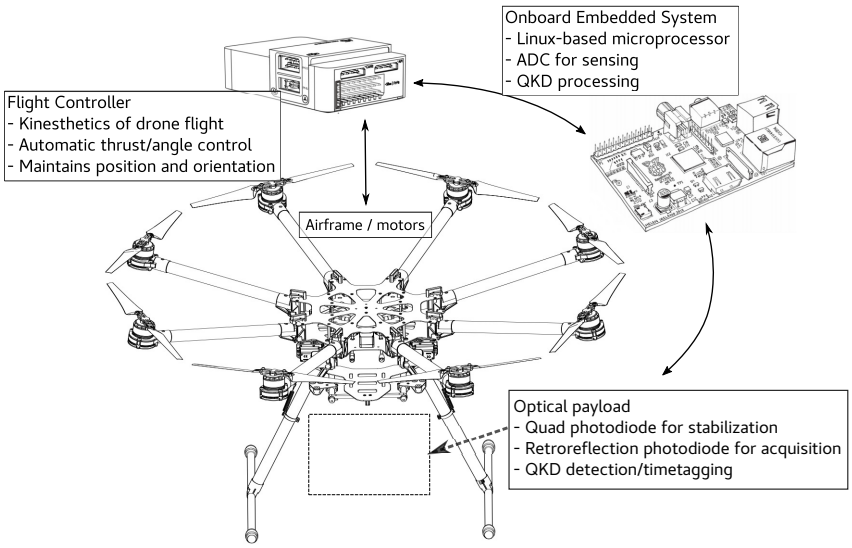

Figure 5.3: Node system design. Several subsystems conspire to align the multi-rotor UAS properly in our node model. The flight controller is responsible for maintaining proper position of the UAS on a kinesthetic level. An onboard microprocessor (in this case, an inexpensive Raspberry PI) instructs the flight controller on how to correct its position at a high level (e.g., *rotate 45 degrees*, or *translate 3 m horizontally*). The microprocessor also interprets sensor data from the optical payloads (such as quadrant-cell photodiode signals) and instructs the payload steering mirrors how to respond. (composite sources: DJI and Helen Ireland (CC))

function of the PAT payload, discussed later. To distribute a key, both drones will takeoff simultaneously from their launch positions and begin signal acquisition. Once a link has been established, they will being to transmit single photons for BB84 QKD. The QKD-centered electronics (FPGAs, microcomputers) are not connected to the stabilization and networking hardware to avoid the possibility of tampering; the drone must be physically accessed by the two parties after landing to begin the classical post-processing (sifting, error correction, and privacy amplification). Each drone remains in visual contact with the ground to preclude the possibility of physical access by an unauthorized party; in particular, there is no way to access the detector measurements on the receiver without landing and accessing the data directly. Because the drones themselves are secured by line of sight, they as an approximation to a "trusted node," a neutral third party that can be used to extend a quantum network.

The core of our effort is in developing a PAT stabilization payload that consumes minimal SWAP while being robust enough to close an optical communications link over useful distances (500 - 1000 m) between two moving platforms.

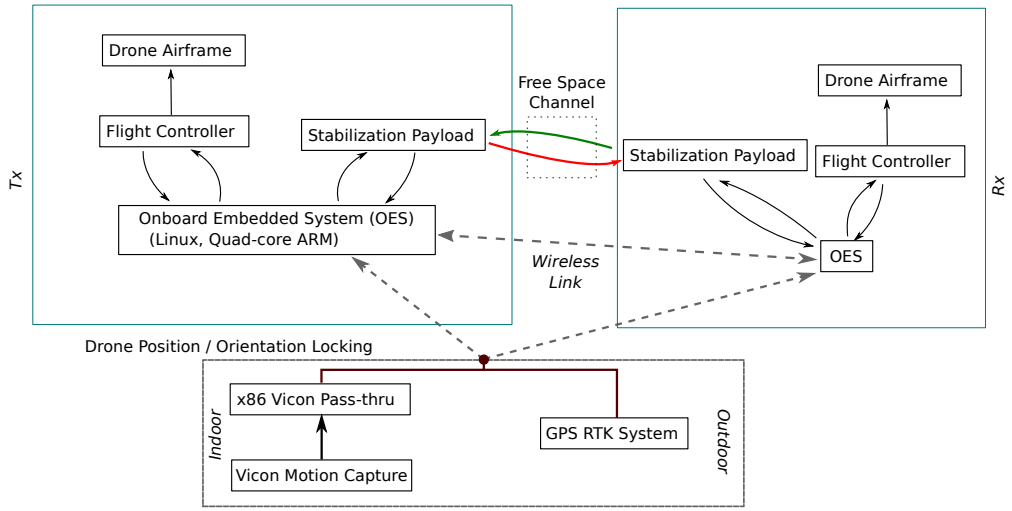

Figure 5.4: Sketch of the two-drone system. Each drone transmits and receives an alignment beacon (discussed in Figure 5.5) over a free-space channel while radioing position-correction data to its partner. Both drones receive location data from either a synthetic GPS source (Vicon motion-capture system), or a high-resolution GPS-RTK system. The former is used exclusively indoors to provide extremely high-resolution position tracking for the drone body and to allow the drone to locate itself precisely without access to a true GPS signal. The latter is used outdoors; the realtime kinematic (RTK) GPS system counteracts fluctuations in the received GPS signal using a stationary base station to improve location accuracy.

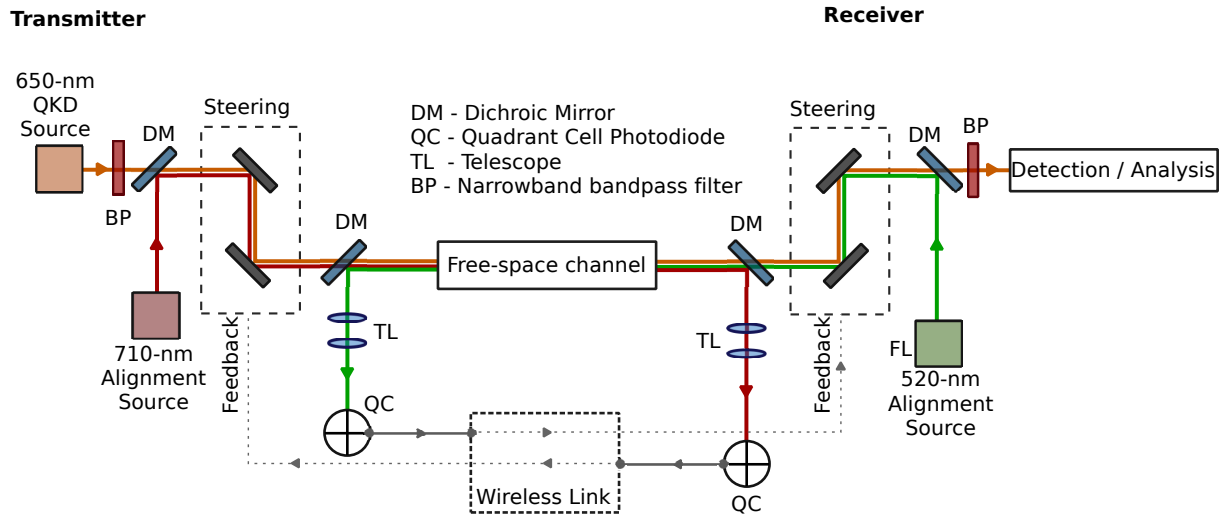

Figure 5.5: Schematic of pointing and tracking system. A 710-nm beacon is mixed with the 650-nm signal states, oriented using two fast piezoelectric steering mirrors, and transmitted through a free-space channel to a distant quadrant-cell detector on the partner drone. Position information from the quadrant-cell detector is reported by the target drone to the origin drone as a feedback error signal. The wireless link is ether a WiFi (indoor testing) or 915 MHz radio (outdoor testing) connection. The tracking system is symmetric: the target drone sends a back-propagating beacon at 520 nm and the origin drone reports beacon position information over the wireless link for feedback.

5.4 Technical Details

As established previously, quantum key distribution is loss-sensitive due to the deterioration of the signal-to-noise ratio when the QKD channel loses a significant fraction of the photons traveling from Alice to Bob. In free-space applications of QKD, this loss may be the direct result of poor performance of the *pointing and tracking* (PAT) system, which is the method used to ensure the signal photons are correctly aligned with the quantum receiver. QKD PAT systems have historically used quite large telescopes to expand the signal beam and prevent diffractive losses while propagating long distances, and typically require at least one ground station with sophisticated beam-tracking cameras and large motors used to slew the transmitter and receiver optics [87]. These systems usually share a common design: a bright classical beacon co- or counter-propagates with the signal beam (or both), while the ground station(s) monitor the beacon's position and/or entrance angle. When the monitored beam drifts from its center position, a corrective motion is applied on the either the transmitter or the receiver side that compensates for the drift. When the beacon(s) are centered, the signal states have a direct path from the transmitter to the receiver that couples into the receiver QKD analysis optics.

Drones do not have the luxury of carrying large telescopes, motors, and drivers due to extremely restrictive SWAP requirements; furthermore, the moving platforms studied in previous work typically move in a smooth, predictable manner, while drones do not, therefore changing the requirements of the PAT system. We have developed a small, lightweight payload, sketched in Figure 5.5, that is able to maintain an optical link between drones. Signal photons at 656-658 nm (produced by resonant-cavity LEDs via an optical bench developed in collaboration with OSU) are mixed with a forward-propagating 710-nm beacon using a dichroic mirror, which, in this case, transmits wavelengths shorter than 697 nm (Semrock FF697-SDi01) and reflects wavelengths above 697 nm. Both the signal photons and the beacon reflect from two mirrors (Thorlabs Protected Silver PF10-03-P01) oriented at 45° angle of incidence. These mirrors are mounted on fast, high-resolution piezoelectric stepping motors (Piezo LEGS 80mnm LR80) capable of positioning down to 0.05 μrad, and are oriented 90° askew to one another so that one motor produces horizontal-plane deflections while the other produces vertical deflections. The signal and forward-propagating beacon pass through a second dichroic filter before exiting the payload.

The receiver payload is identical in design, except that the backward-propagating beacon operates at 520 nm. On both the transmitter and receiver sides, the incoming alignment beacon is separated from the signal beam using a dichroic filter and routed through a Keplerian telescope before impinging on a quadrant-cell position-sensing detector. The Keplerian telescope serves two functions: first, the beacon is too large (approximately 25 mm in diameter) at the measuring quadrant-cell detector, as the detector face

is approximately 8 mm in diameter. The telescope reduces the beam to the appropriate diameter to not overfill the detector. Second, the telescope adds a degree of insensitivity to entrance angle. Imagine a two-lens telescope, shown in Figure 5.6, with an ideal lens of focal length f_1 followed by an ideal lens of focal length f_2. If these lenses are placed a distance $f_1 + f_2$ and light is input at an angle of incidence θ at a transverse displacement d, at the output plane at $f_2(f_1 + f_2)/f_1$ the offsets (d_f, θ_f) are given by:

$$\begin{pmatrix} 1 & f_2(f_1+f_2)/f_1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f_2 & 1 \end{pmatrix} \begin{pmatrix} 1 & f_1+f_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f_1 & 1 \end{pmatrix} \begin{pmatrix} d \\ \theta \end{pmatrix} = \frac{f_2(f_1+f_2)}{f_1} \begin{pmatrix} -df_2/f_1 \\ -f_1\theta/f_2 \end{pmatrix}.$$

The telescope therefore magnifies the beam by f_2/f_1 and the final output displacement is given by $-df_2/f_1$, which does not depend on the input angle θ. This contrasts with a Galilean telescope, which consists of a negative focal length lens followed by a positive focal length lens, produces a virtual image, and does not have a true focus anywhere between the constituent lenses. Such telescopes do not produce angle-invariant images, and in fact will introduce output transverse offsets from input beam angle offsets. This reduces the probability that the beacon will miss the quadrant cell detector if the drones are not pointing directly at one another. For real lenses, however, there is a slight coupling between the incident angle and the output displacement, but it is significantly reduced from what would be expected from the simple geometry of entering the receiver aperture at some angle, reflecting from a dichroic mirror, and propagating to the detector. Further, the above ABCD matrix calculation also implies that if we take the plane at which we desire the entrance angle to be invariant to be shifted back from the front face of the first lens, then the angle-insensitive output plane shifts slightly toward the final face of the second lens. The overall effect of this lens geometry is discussed in Figure 5.6.

After the beacon is reduced and shifted by the telescope, it propagates a small distance (~ 1 cm) and is recorded by a quadrant-cell detector (First Sensor QP50-6SD2). The detector consists of four photodiodes in a Celtic wheel cross pattern (bifurcated in the horizontal and vertical directions to form four quadrants). The detector outputs three voltages, corresponding to the total incident optical power on all four photodiodes, the difference between the power in the vertical half of the detector and the lower half of the detector, and the difference in power between the left half and right half. The photodiodes have bandwidths in the kHz, so the quadrant cell photodiode (quad cell) produces a high-speed, high-precision estimate of the centroid of an incident gaussian-profiled beam. If the beam is not gaussian, e.g., if turbulence disturbs the beacon wavefront significantly, the quad cell estimate of the beam centroid location may degrade significantly. The quad-cell voltages are polled at 3 kHz by a fast 24-bit analog-to-digital converter (ADC, Texas Instruments

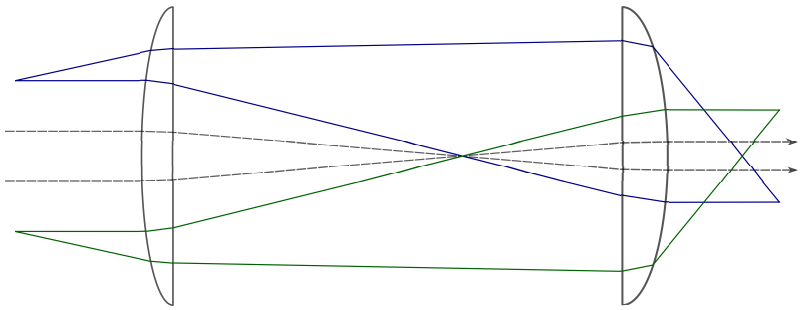

Figure 5.6: Function of Keplerian telescope for angle-insensitive beam offset measurement. A Keplerian telescope produces a real image at the output plane. This implies that, for an ideal telescope, the angle of incidence at a given input plane does not matter, at the output plane any input transverse offset will result in an output offset, scaled by the minification factor of the telescope. Though real lenses reduce the robustness of this setup, our optical train exploits this invariance to input angle to reduce the effect of angular mismatches of the beacons traveling between the transmitter and receiver.

ADS1256) attached to the onboard Raspberry PI microcomputer. The microcomputer normalizes each offset voltage by the total power voltage to establish an X/Y position estimate.

The X/Y positions are then radioed to the partner drone using a wireless AC1750 radio every 7 ms. The high performance of the AC1750 protocol produces transmission latencies between 3-10 ms, or approximately one cycle. Upon receiving an updated position estimate for its outgoing beacon, the partner drone calculates the offset error and attempts to reduce the error using a nonlinear proportional control to shift its local piezo mirror rotation. Given a feedback gain K, the command sent to each of the X and Y steering mirrors is scaled using the slightly nonlinear function Error $= K\text{sgn}(x)x^{6/5}$, which produces a small inflection near the zero-error point that can prevent ringing at high gains (Fig.5.7). The exponent $6/5$ was chosen to create an inflection point near $x = 0$ while approximating linear control for error values close to 1. Great care was taken to produce multi-threaded error signal transmission and receiving code that operates efficiently and at a low latency (See Appendix C).

The payload is 3-D printed from polylactic acid (PLA), a lightweight plastic-like material using an Ultimaker 3 Extended 3-D printer. Independent 3-D printed components are combined using hollow aluminum linkages to replace the usual stainless steel cage-mount rods. Though the general design of the payload is adapted from standard optical cage-style optic mounts, we save approximately 1 kg out of 1.8-2.0 kg in payload mass by moving to plastics (see Table 5.1). The entire assembly rests on a commercial cinema-grade camera gimbal (Gremsy T3), which automatically rotates to counter any motion of the drone during flight and dramatically reduces the angular fluctuations introduced by the drone attempting to maintain its heading.

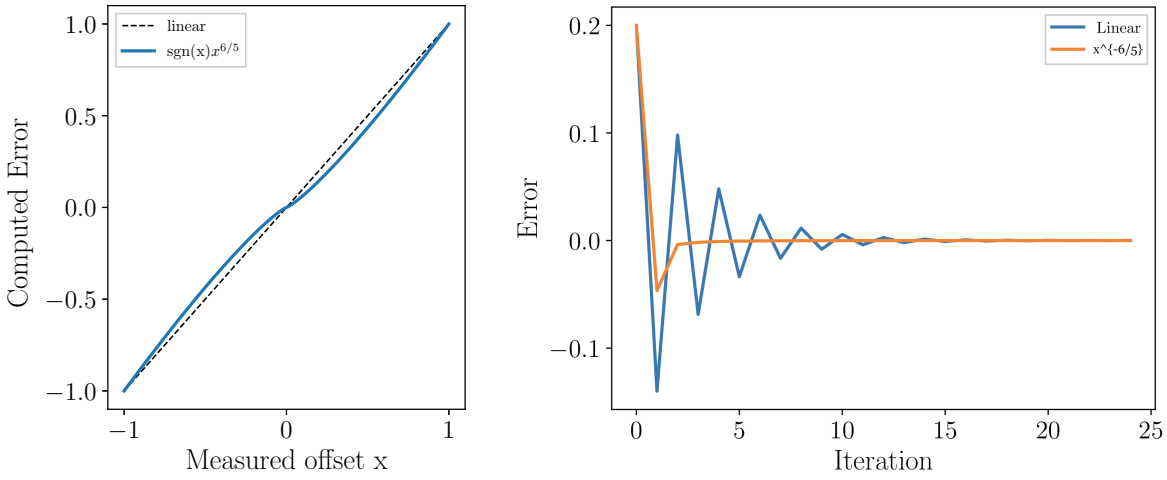

Figure 5.7: Comparison of standard linear proportional control with inflected proportional control (left). If input errors are normalized, the errors may converge to zero more quickly than for a standard linear proportional control (right).

5.5 Preliminary Results

Our initial testing has been conducted indoors in an arena equipped with motion capture cameras (Fig. 5.8(a)). The motion capture system (Vicon) tracks small, spherical, retro-reflecting markers attached to the upper side of each drone using 810-nm floodlight LEDs. This system is similar to those used for motion capture of actors for commercial films, and is capable of measuring the orientation and location of each drone at rates of up to 200 Hz and sub-centimeter resolution. A host computer records the motion capture data, re-routes it to a UDP broadcasting program that converts the position data into spoofed GPS coordinates, and sends them to the computer onboard each drone via WiFi. The onboard computers listen for the GPS data and convert them to messages the flight controller can understand before transmitting the faked GPS data to the flight controller over a wired serial interface. This allows the flight controller to use its built-in GPS position locking algorithms using high-precision position data from the motion capture system. While convoluted, this setup allows us to maintain a position hold of these large drones to within about half a meter, and almost precludes the possibility of the drones accidentally slewing into walls or the ceiling due to GPS drift inside the building.

To date, we have successfully achieved bidirectional optical PAT lock between one flying drone and one stationary drone in the laboratory. A sample of the transmission achieved during one such run is shown in Figure 5.9. The mean transmission of a bright 635-nm "dummy" beam was −9.87 dB, or about 10%, while the maximum transmission approached −2 dB, or 63%, while the stabilization system was operating over

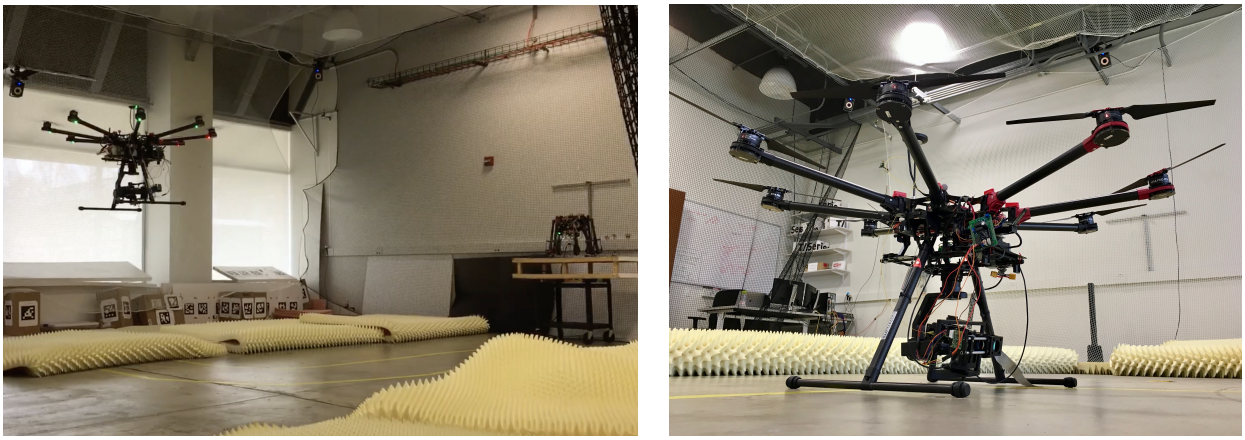

(a) (b)

Figure 5.8: (a) A photograph of the drones operating establishing an optical link, with one drone in flight and the other stationary (on table to the right). Data from a sample of these laboratory runs is shown in Figure 5.9. (b) Photograph of the transmitter drone and payload before flight.

about 40 s. When the system was not operating, no light was received by the stationary drone. With this setup the mean locking time was approximately 60 s, averaged over 3 tests.

As the objective of this experiment is to establish a QKD link exchanging single photons between drones at some distance (> 500 m), we are actively working to repeat the performance of our indoor locking system in an outdoor setting, with both drones flying simultaneously. To date we have achieved locking in a single direction at a time for time scales on the order of 5-10 s.

There are a number of circumstances which make outdoor flights more difficult than our indoor testing. First, real GPS systems cannot provide the precision or reliability of the indoor motion capture system. Even our higher-resolution realtime kinematic GPS systems (RTK), which use a stationary base station to counteract GPS noise, cannot locate the drones to better than 1-2 m. Furthermore, the variance of the relative velocities of the two drones is doubled when both are flying ($\mathrm{Var}[v_1 + v_2] = \mathrm{Var}[v_1] + \mathrm{Var}[v_2]$), increasing the demands on the stabilization system. Over longer distances these drifts will matter less, as the stabilization system controls the outgoing angles of each bright beacon; however diffractive effects and small motions of the gimbal and drone bodies will contribute to greater downstream oscillations in these cases, so the scaling of the performance of the stabilization system with distance is not clear. Additionally, background light from the dusk sun can cause the quadrant-cell detectors to misread the magnitude and position of the incoming beacon, resulting in locking loss and incorrect error estimations. To address this issue we have incorporated 520-nm and 705-nm bandpass filters (10-nm bandwidth) in front of each quadrant-cell detector for outdoor testing.

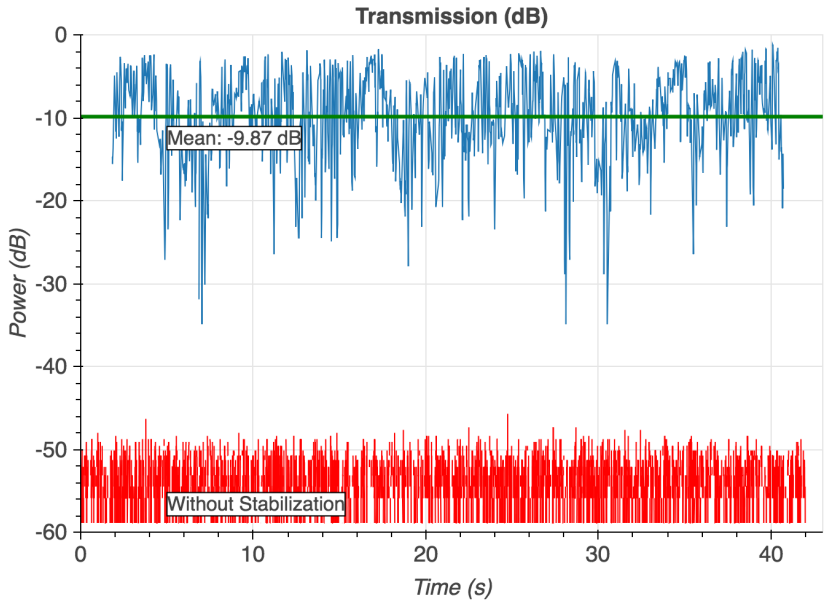

Figure 5.9: Preliminary transmission data from a sample flight in an indoor environment between one flying drone and one stationary drone. Here the pointing system remained locked for over 40 s, with a mean transmission of a bright "signal" beam of -9.87 dB and a peak transmission of ~ -2 dB. No light was received when the stabilization system was not running.

5.6 Extensions and Conclusions

We have taken the first few steps toward a complete, end-to-end, free-space quantum key distribution system between multi-rotor drones. Though our current progress has been promising, many challenges remain with integrating the QKD signal and analyzer optics benches and automatically establishing a robust link after the drones begin flight.

Regarding the latter point, we are developing systems for automatic acquisition and locking of the stabilization beacons. This requires the development of two subsystems: one for setting the proper initial position of the drone and gimbal, and one for adaptively refocusing the beacons from a large beam diameter at the receiver to a tight, approximately collimated beam. This would allow the drones to first acquire a rough alignment using large beacons; as the beacons focus down after initial acquisition, the precision of the tracking system would increase until the two beams are overlaid onto one another.

To achieve a rough initial acquisition, we are investigating the following protocol: an ultra-bright 465-nm LED is powered on each drone; the partner drone's onboard computer uses an inexpensive integrated camera (Raspberry Pi Camera Module V2) and inspects only the blue field of the measured RGB image. Provided the drones are oriented toward one another to within 60 degrees, the camera will show the distant drone's target LED as a bright spot. The drone will then reorient its gimbal and body so that the centroid of the

Item	Original Weight (g)	Reduced Weight (g)
16000 mAh Battery	1932	1932
Stabilization Optics/3D Printed Parts	1200	657
Motor Drivers (ea)	370	370
Detectors	924	924
Support Plate	202	58
OSU Tx/Rx Bench (Est)	600	600
Total Rx Payload	5228	4541
Total Rx Weight on Gimbal	1800	1257
Total Tx Payload	4304	3617
Total Tx Weight on Gimbal	1800	1257

Table 5.1: Weight budget as a result of moving to printed plastics

target LED is located in the center of the camera visual field.

After rough alignment has been achieved, each stabilization beacon will automatically adjust to have an opening angle between $0.5 - 1.0°$. This will be accomplished using a high-resolution piezoelectric lens translation stage intended for high-performance micro-camera applications (New Scale M3-FS). At the target drone the beacon will then form a much larger spot than can be usefully detected by the quadrant-cell detector; after both drones detect enlarged beacons they will attempt bidirectional locking while the beacon lenses are slowly translated, reducing the beam divergence. When the beacons are roughly collimated, the procedure will stop; at this point the stabilization system will be locked and the drones may begin sharing single photons for QKD.

Our collaborators at The Ohio State University are exploring the use of single-mode fiber-resonant-cavity LEDs (RC-LEDs), operating at 656 nm with a $\sim$ 6-nm bandwidth, for the quantum signal. Each of 4 LEDs is used to generate one of the BB84 signal states and generate pulses of multiple mean photon numbers for the purposes of decoy-state QKD. The wide-band signal photons are then filtered using a sub-nanometer bandwidth band-pass filter. The filter bandwidth is much narrower than the spectra of each LED, rendering each LED spectrally almost indistinguishable. This drastically reduces the efficacy of any potential attacks attempting to exploit spectral differences between different polarization states. Though LEDs are not a coherent source of photons, but a thermal source, many individual photons are able to couple into the fiber pigtail affixed to each LED with an efficiency of 0.009%. The single-mode fiber pigtailing of the diodes allows for strict spatial-mode filtering, thereby preventing side-channel attacks on the spatial mode of the photons. Work on integrating these sources into a compact, 3-D printed optics bench producing a high state quality and that can seamlessly integrate with our stabilization system is ongoing.

To achieve longer channel distances (> 20 m) we must alter our wireless infrastructure to allow for longer-range communication of PAT data between the two drones. Our current AC1750 WiFi-based communication system, while pragmatic, is severely limited in range due to its 5-GHz operating frequency. By moving to 915

MHz or even 2.4 GHz – both in the unlicensed scientific band – we can improve range without significantly compromising throughput. We are currently investigating Software-Defined Radio (SDR) as a replacement technology. SDR allows us to use high-speed, application-agnostic radio hardware in any way we choose by writing encoding and decoding software that interacts with the data streams read from the radio ADCs directly. As our communication needs are much simpler and require less performance than standard WiFi is configured to deliver, SDR offers an a means of increasing our signal reliability at the cost of programming complexity. Current results suggest future latencies of 1 ms at 100 Hz update rates.

To be a truly useful addition to the bestiary of viable QKD systems, it is not enough to have a single transmitter drone and a single receiver drone; in a future network, each drone must be able to act both as a transmitter and a receiver, as this allows them to act as a node between any two other QKD sources in the network (i.e., a trusted relay). Future work will involve expanding our payload to allow for bi-directional QKD data by incorporating both transmitter and receiver optics and electronics. To avoid interference, each channel will operate at a slightly different frequency. In its final form, the payload will be able to transmit to one partner and receive from another simultaneously, allowing the drone to establish real-time, quantum-secured key distribution network. As the system is designed to preserve the states used in QKD, the drone-to-drone link will be naturally suited for more exotic tasks such as the distribution of entanglement for future hybrid classical/quantum networking.

Chapter 6

Summary and Conclusions

Though quantum communication offers many theoretical advantages over classical communication, the future viability of quantum networks relies on practical solutions to significant experimental challenges. In this thesis we have demonstrated a variety of techniques for improving the functionality of quantum networking – specifically, quantum cryptography – over fiber-based and free-space channels. We have shown that it is possible to optimize the wavefront of single-photon signals (e.g., from a quantum dot or trapped ion) for coupling into fiber optics without the use of traditional imaging sensors – useful for coupling into and out of future quantum memories that could extend the range of quantum communication networks; we have also developed controllable techniques for simulating turbulence in the laboratory and used these techniques to demonstrate a protocol, selective deactivation, for improving secure key distribution rates over turbulent atmospheric channels. We have concluded by discussing our efforts to develop a pragmatic QKD system aboard multi-rotor drones, which could function as trusted nodes in a future quantum network and be used to distribute entanglement over such a network.

Applications for these techniques are plentiful. Though some quantum communication architectures have been demonstrated over active fiber optical telecommunications networks (i.e., not dark fiber) [94], these solutions rely on the existence of physical infrastructure at the source and destination nodes. For many applications of interest, e.g., rural or mountainous environments, the existence of suitable communication hardware is not guaranteed. This so-called "last-mile problem" is the subject of classical optical communication research efforts at a number of large internet companies, for example, Google X labs (Project Loon). These solutions are not necessarily intended to supplant fiber solutions, which, in some cases, can provide superior rates and speeds over free-space solutions, but are intended for the creation of short-term, low-cost networking nodes that can operate without requiring fiber infrastructure to be strung through difficult or rural terrain. A future quantum network interconnecting quantum computers and users must also address this last-mile problem if quantum-enhanced technologies are to be accessible by a wide range of users. Research on the last-mile problem is just as critical for quantum wide-area networks as it is for classical networks. The interim solution for areas which do not have access to robust, fiber-based quantum networking is over free

space, as it is in the classical case. A system that can bridge the gap between users and quantum computational resources located away from quantum-enabled communication infrastructure will be necessary until such infrastructure becomes commonplace. Our efforts toward enabling free-space quantum communication (e.g., entanglement distribution, quantum cryptography, etc.) using selective deactivation and multirotor drones are small steps along the path to solving this last-mile problem.

Many significant and interesting challenges remain; for example, how should point-to-point quantum protocols be best implemented for large-area networks with many users? Some protocols, for example, measurement device-independent quantum key distribution (MDI-QKD) [95], allow a point-to-point network to use an untrusted node, e.g., a (possibly malicious) satellite, to distribute a key securely. Future work may require the adoption of MDI-QKD onboard mobile free-space platforms, such as sea- and aircraft, demanding the development of high-precision clock distribution systems and low-SWAP optical payloads that enable projection of QKD photons onto the Bell states needed for the protocol. This allows the channel to be extended using an untrusted node; however, MDI-QKD performs very poorly in the presence of loss, and multi-user analogs of MDI-QKD have yet to be demonstrated.

Turbulence and accurate PAT remain significant hurdles and areas of interest for enabling practical free-space quantum key distribution [96, 97]. These issues become more and more relevant daily: laboratory QKD increasingly moves toward the use of higher-dimensional Hilbert spaces to increase channel capacity [98, 99], and small errors due to turbulence and PAT latency/accuracy can dramatically influence the receiver's ability to project into the correct basis elements. Despite the challenges, solutions to these practical issues facing free-space quantum communication are critical for increasing the practical utility of quantum protocols.

Appendix A

Polarization-Maintaining Fibers

A.1 Introduction

A prerequisite for a BB84-based quantum cryptographic link is the ability to faithfully send information in two (or more) mutually unbiased bases (MUBs) stably; that is, the ability to send the quantum states $|1\rangle$, $|0\rangle$, $|1\rangle + |0\rangle$, and $|1\rangle - |0\rangle$. For encoding in some degrees of freedom, e.g., time bins, single mode fibers (SMF) can serve as reliable quantum channels; however, if $|1\rangle$ and $|0\rangle$ are orthogonal polarization states, they are unstable to perturbations of the fiber, as the fiber applies a well-defined but variable unitary transformation to the polarization states propagating in the fiber. Traditional polarization-maintaining fibers (PMF) can stably send $|1\rangle$ and $|0\rangle$, but superpositions are not preserved (assuming these are aligned with the fast and flow axes of the fiber), as $|1\rangle$ and $|0\rangle$ correspond to nondegenerate fiber modes. It would be ideal to find a fiber where all 4 states of the two MUBs are stable; i.e. a fiber in which a 2-fold set of degenerate modes and an arbitrary linear combination of said modes are stable to external perturbations. Moreover, such a thing would constitute an arbitrary polarization maintaining fiber (PMF), which would likely be useful in both the classical and quantum domains; for example, a truly polarization-maintaining fiber would allow for the offloading of heavy and sophisticated single-photon detection electronics and optics for QKD to a ground station, while a tethered mobile platform (e.g., drone or balloon) establishes a long-distance optical link.

A.2 OAM-Based PMF

In collaboration with Boston University (PI: Siddharth Ramachandran), we performed a performance analysis of a new type of fiber which demonstrates polarization-maintaining capabilities. The fiber they have developed is able to preserve the phase relationship between the $L = \pm 7$ orbital angular momentum (OAM) states in fiber [100]. This is accomplished in part by introducing a hollow core to the fiber, thereby making OAM states into fiber eigenmodes.

By mapping the right- and left-hand circular polarizations to complementary high-L orbital angular mo-

mentum states (e.g., using a q-plate [101], an electro-optic device capable of transforming optical polarization into orbital angular momentum), it is possible to couple the polarization of a photonic state of interest into the orbital angular momentum eigenmodes of the specialty fiber. If such a fiber truly preserves the phase between these two components, it must also preserve all other polarizations (H, V, for example), which are linear superpositions of the circular states.

We studied the entanglement and polarization preservation behavior of high-l OAM states in 10m of this novel air-core fiber using both ancilla-assisted process tomography (AAPT) [102] and single quantum process tomography (SQPT). We observed preservation of entanglement and polarization (up to a correctable rotation) using AAPT. Using SQPT, we observed polarization invariance with changing temperature up to 50 degrees C; however, we also observed that standard SMF was temperature-invariant over the same range. We studied polarization rotation as a function of twisting paddles of a polarization controller, and found that the polarization state leaving the OAM fiber is relatively well-maintained (before/after process fidelity of $98.8 \pm 0.2\%$), while the polarization state leaving the SMF is scrambled (before/after process fidelity of $6 \pm 1\%$), as expected. Our AAPT experiments were limited by a low signal to noise ratio due both to detector limitations and a relatively high (9.6dB) total coupling loss in the OAM input/output coupling system, in addition to practical issues in achieving a high-purity entangled state to use as a source. Our SQPT measurements were limited by the high sensitivity of the fiber to out-of-plane motions, which made studying single-plane perturbations practically challenging.

A.3 AAPT with entangled photons

We performed ancilla-assisted quantum process tomography on the OAM-PMF. AAPT functions like a standard quantum process tomography (SQPT), but requires entanglement between the probe photon and a reference photon. State tomography, in this case realized by polarization tomography, is performed on both the probe and reference photons simultaneously, and coincidence counts are measured. For example, suppose the maximally entangled state $(|HH\rangle + |VV\rangle)/\sqrt{2}$ is created, and the probe photon is sent through a system whose response can be modeled as the identity matrix in state space. When both the probe and reference photons are measured in the $|H\rangle$ basis, coincidences will occur, while no coincidences will occur when the probe photon is measured in $|H\rangle$ and the reference is measured in $|V\rangle$. In this way, both the degree of entanglement and the influence of the system through which the probe photon propagates can be determined.

The source for this experiment is detailed in Fig. A.1. Photon pairs were generated at 1550 nm and

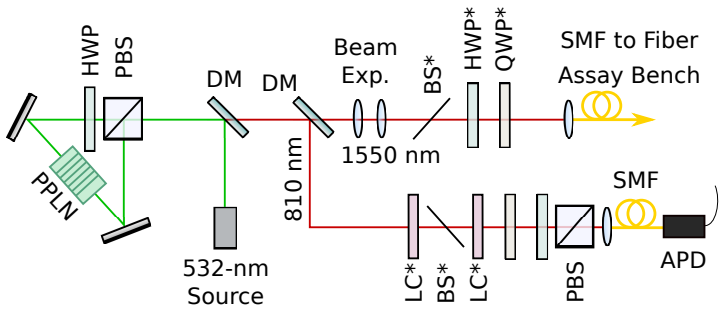

Figure A.1: Light from a 532-nm free space laser source generates entangled photon pairs at 810 nm and 1550 nm in a Periodically Poled Lithium Niobate (PPLN) crystal. The components denoted with a * are in the setup but not used for these experiments. The two wavelengths are separated with a dichroic mirror (DM). The 1550-nm photons are sent through a 10-m SMF patch cord to the OAM coupling system, while 810-nm photons are sent though a series of waveplates and a liquid crystal (LC) cell to ensure the desired polarization state in the reference arm, before it is passed through a polarization tomography setup, coupled into an SMF, and directed to an avalanche photodetector (APD).

810 nm via spontaneous parametric downconversion of a 532-nm pump in periodically-poled lithium niobate (PPLN). By tuning the pump polarization we were able to create pairs of photons in the maximally entangled polarization state $(|HH\rangle + |VV\rangle)/\sqrt{2}$. The 810 nm photon in each pair was routed directly to a standard automated tomography setup consisting of a quarter-wave plate (QWP) and half-wave plate (HWP) on rotation stages, and a static polarizing beamsplitter (PBS). The 1550-nm photon was coupled into a 10-m SMF and injected into the OAM-PMF testbed.

The system for OAM coupling and probe photon tomography is shown in Fig. A.2. For initial calibration and testing, the OAM-PMF was bypassed completely (here called the 'input' state), and a complete tomography was performed immediately after the 10-m SMF from the source (the 'output', labeled 'Tomo. for Bob ref.' in Fig. A.2). This allowed us to characterize the source and any deleterious effects of the long input fiber to the testbed. As with the 810-nm photon, the 1550-nm tomography consisted of a QWP/HWP/PBS combination.

To characterize the behavior of the OAM-PMF, the flipper mirror was lowered, thus bypassing the reference tomography system, and the 1550-nm photons were passed through a $q = 7/2$ plate to produce $l = 7$ OAM modes. The free-space OAM beams were coupled into the 10-m OAM fiber with an 8-mm coupling lens. The fiber was placed on a cardboard slab at approximately the same height as the input coupling stage. The output was collimated and passed through a second $q = 7/2$ plate, converting the beam back to $l = 0$; this then passed through an output tomography setup and coupled into a 10-m SMF using a 6.25-mm aspheric lens, and subsequently directed to a single-photon detector (NuCrypt CPDS-1000, cooled InGaS). The total loss experienced by the probe photon (as measured by a higher-power classical beam) was approximately 9.6 dB. Because a difference in angle of orientation of the input and output q-

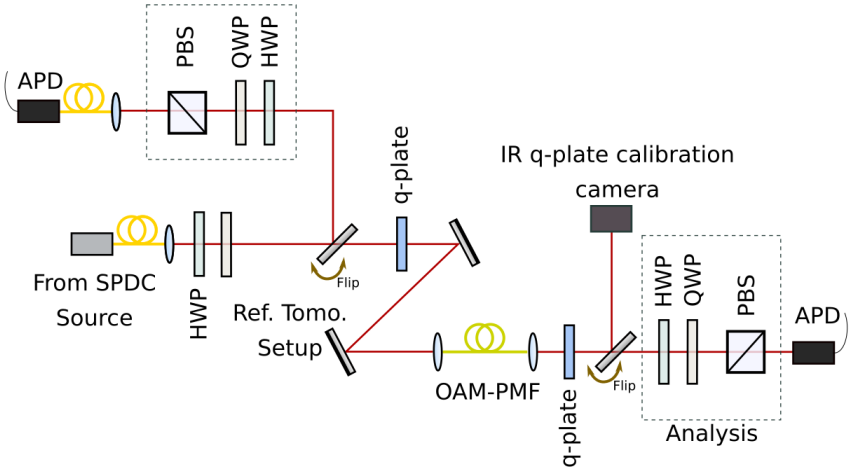

Figure A.2: OAM coupling and output tomography system. Light from the source is passed through two polarization-compensation waveplates which are used to correct for polarization rotation in the 10-m SMF link. To characterize the source photons before entering the OAM coupling system, a reference measurement is made by raising a flipping mirror that directs the injected photons to a polarization tomography setup, after which the light is coupled into an SMF for detection. For PMF measurements, the tomography system is removed, and the light passes through a $q = 7/2$ plate before being coupled into 10 m of OAM fiber. The fiber output is directed through a second q-plate and either observed visually (for alignment), or directed to the output tomography setup and coupled into an SMF for detection.

plates can cause polarization rotation, the output q-plate was kept on a rotation stage. An $|H\rangle$ state was sent through the input q-plate, fiber, and output q-plate, and then passed through a polarization beam displacing prism. The output q-plate was rotated by a few degrees until the output polarization is $|H\rangle$. This also compensates for slight changes in height of the input/output/arrangement of fiber, which could induce slight linear polarization rotation.

Each tomography setup measures 6 polarization projections: $|H\rangle$ and $|V\rangle$, $|D\rangle$ and $|A\rangle$ (linear polarizations at 45 degrees to H and V), and $|R\rangle$ and $|L\rangle$ (the two circular polarizations). Thus, each tomographic measurement uses 36 projective polarization state measurements, from which the behavior of the fiber can be determined (the minimal number of measurements required to perform a 2-photon state tomography is only 16 (e.g., H/V/D/R $\times$ H/V/D/R), but to improve the quality of our tomographic estimates we included all 36 measurements).

A.3.1 AAPT results and discussion

Pure quantum states may be represented as kets; however, an arbitrary quantum state (including mixed states) must be represented as a density matrix, denoted ρ. For a pure state $|\psi\rangle$, the density matrix $\rho \equiv |\psi\rangle\langle\psi|$. For a mixture of states $|\psi_i\rangle$ occurring with probability p_i the density matrix is defined as

$\rho \equiv \sum_i p_i |\psi\rangle\langle\psi|$. As density matrices are a generalized representation of quantum states, they allow for the expression of more general state operations – such as decoherence – which may not be unitary. The goal of quantum state tomography is to estimate the density matrix of a quantum state itself, e.g., after it has passed through some system; in contrast, the goal of AAPT is to determine how the process describing a quantum communication channel (e.g., transmission through an optical fiber) would affect any input state (obviously, the effect will often depend on the particular quantum state input).

The results of the AAPT measurements can be quantified by purity and concurrence. Purity, defined as $\mathrm{Tr}\,\rho^2$, is equal to 1 for a pure state (i.e., any state representable using a single ket); a completely mixed state (such as un-polarized light) will have purity $1/d$, where d is the dimension of the state space. The entanglement present between two particles may be quantified using the concurrence[1]. For maximally entangled states, the concurrence is 1; for example, for the state $(|HH\rangle + |VV\rangle)/\sqrt{2}$, both the purity and concurrence are equal to 1, but, for the state $|HH\rangle$, the purity is 1, while the concurrence is 0.

We demonstrated that the two-photon state was preserved up to a correctable rotation after the 1550-nm daughter photon was transmitted through the OAM-PMF. (Fig. A.3) indicates the best experimental result we obtained using the 10-m OAM-PMF in an "unperturbed" state. The density matrix ρ is plotted in both absolute value (left) and phase (right) for the input (top) and throughput (bottom) cases. The purity degrades slightly, from 0.88 to 0.81, as does the concurrence, from 0.86 to 0.79, which may be due to increased noise in the measurement due to the input/output coupling loss in the OAM coupling system. Nevertheless the results demonstrate that the fiber preserves entanglement, as expected.

We encountered some difficulties generating enough photons to perform the two-photon tomography with an adequate signal-to-noise ratio after passing through the lossy OAM-PMF testbed. This particular iteration of the PPLN-based entangled-photon source used in the AAPT assay has a number of compromises that make it appropriate for the generation of particular two-photon states entangled in both polarization and time mode. It is, however, not well-suited for generating high-purity, maximally entangled polarization states.

While our 1550-nm detectors are reasonably efficient (20%), they are gated at 40 MHz - half of our pump repetition rate. These detectors are quite noisy as well: while gating at 40 MHz we observe a dark count rate of almost 3000 cps (standard silicon avalanche photodiodes, by comparison, typically operate with below 500 cps dark counts when completely ungated – below 2 cps given our gating scheme. These detectors, however, cannot detect 1550-nm photons). This is nevertheless sufficient (though not preferable) for the experiment using this particular source. However, it imposes significant limitations on the maximum

[1]Defined by $C = \max(0, \lambda_1 - \lambda_2 - \lambda_3 - \lambda_4)$, where the λ_i are the ordered eigenvalues of the matrix $R = \sqrt{\sqrt{\rho}\tilde{\rho}\sqrt{\rho}}$, with $\tilde{\rho} = (\sigma_y \otimes \sigma_y)\rho^*(\sigma_y \otimes \sigma_y)$ [103].

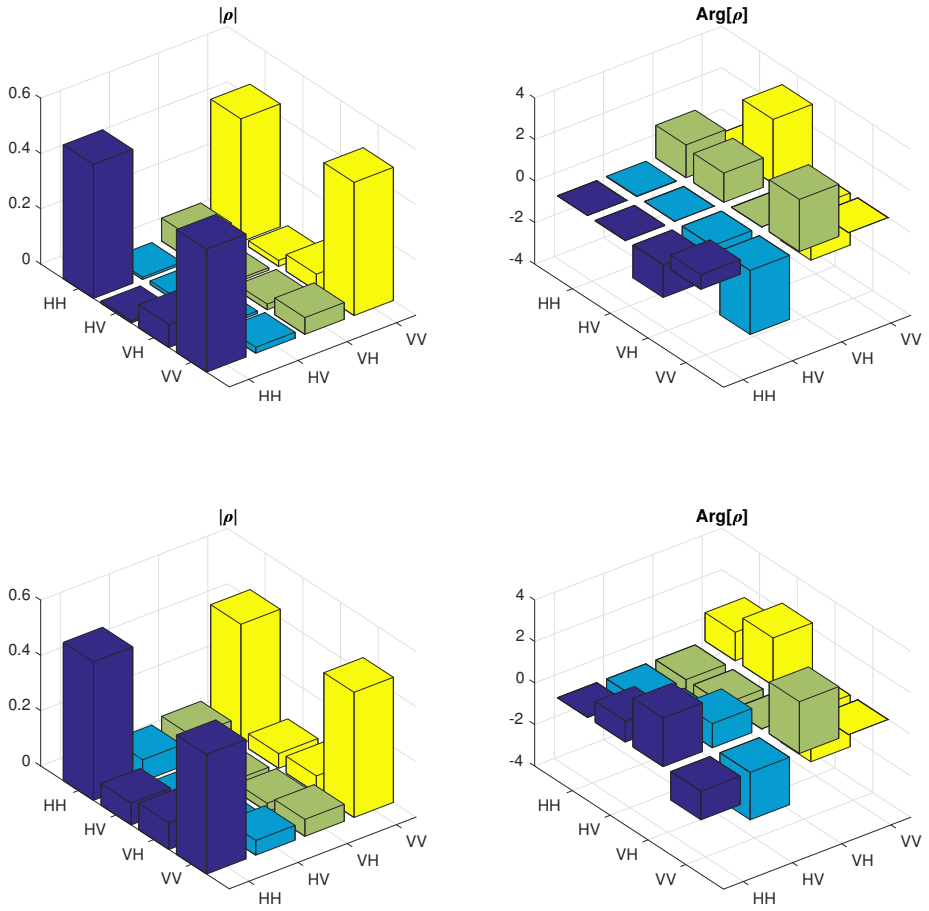

Figure A.3: **OAM-PMF** with two-photon state $(|HH\rangle + |VV\rangle)/\sqrt{2}$. Input (top) Purity $= 0.88 \pm 0.02$, Concurrence $= 0.86 \pm 0.02$. Output (bottom) Purity $= 0.81 \pm 0.02$, Concurrence $= 0.79 \pm 0.2$. While there is some apparent loss in purity (likely due to increased noise in the measurement - the OAM-fiber input/output coupling loss is high), the concurrence tracks close to purity in both cases. From this we may conclude that the OAM-PMF does not significantly decohere the two-photon state. Aside from a small reduction in purity, the difference between the input and output density matrices is a fixed, correctable rotation.

observable signal-to-noise (and therefore estimated purity) for polarization Bell states.

There is also significant spatial walkoff between different downconversion polarizations inside the PPLN, resulting in the H/V downconverted beams having slightly displaced beam paths. Generating the high-purity maximally entangled state requires careful alignment and balancing of these beam paths. Misalignment can compromises the number of photons of either polarization that can be coupled and detected. This walkoff also introduces phases between the polarization components, which can appear as a degradation in the polarization-entangled state purity (when there is a different phase for each frequency, e.g., due to dispersion).

The combination of polarization walkoff in the downconversion crystal and detection inefficiency made it prohibitive to generate high-quality polarization-entangled two-photon states with sufficient brightness to overcome our high 1550-nm detector noise. We were able to adapt and realign the source to produce maximally entangled polarization states appropriate for AAPT; however, the alignment was unstable and resulted in few viable tests using the entangled source. We elected to proceed with simpler classical tests.

A.4 Standard Process Tomography

Standard quantum process tomography (SQPT) allows one to determine the effect of some unknown operation on a quantum state. Typically these sorts of tomographies are limited to measurements at the single-particle level; however, we may apply some of the same techniques to describe operations on classical light fields as well. Here we implement a classical analog of SQPT in order to determine the action of the OAM-PMF on known polarizations of classical (i.e., bright) light.

As in the quantum case, one performs a standard tomography on a set of input polarization states (here, H, V, D, R). By the prescription given in Ref. [104] one may calculate the process matrix χ which describes the behavior of the unknown process. As a brief summary, we may represent any operation on a quantum system ρ as a map $\rho \to \mathcal{E}(\rho)$. $\mathcal{E}$ is a general operator that has a particular matrix representation given by χ_{mn} defined by

$$\mathcal{E}(\rho) = \sum_{mn} A_m \rho A_n^\dagger \chi_{mn}, \tag{A.1}$$

where the A_n are a set of measurement operators (e.g., relating to polarization) which span the space of operators for a given Hilbert space. In this way, χ is a particular representation of $\mathcal{E}$ with respect to a particular set of measurements (e.g., H, V, D, R). One way that χ differs from ordinary unitary operations on states is that it contains information about decoherence (reductions in purity). The "identity" for the χ

representation – no rotations, no mixture, etc, – is given by

$$\chi = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{A.2}$$

Substitutions of $H \leftrightarrow V$ (below, left), $R \leftrightarrow L$ (below, center), and complete loss of V (below, right) result in the following forms of χ:

$$\underset{H\leftrightarrow V}{\chi} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \underset{R\leftrightarrow L}{\chi} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \underset{V\text{lost}}{\chi} = \begin{pmatrix} 1/2 & -i/4 & -1/4 & 0 \\ i/4 & 0 & 0 & 1/4 \\ -1/4 & 0 & 0 & -i/4 \\ 0 & 1/4 & i/4 & 1/2 \end{pmatrix}. \tag{A.3}$$

Even small measurement fluctuations (due to photodetector noise, for example) may produce unphysical χ matrices when channel inversion techniques are used due to small errors between measurements whose results should be equal. For this reason, we first estimate χ using the naïve inversion (e.g., Ref.[104]) and then use an adaptive Monte Carlo method to find the physical (positive definite, hermitian) matrix closest to the estimated χ in terms of the process fidelity.

We can define the process fidelity $\mathbb{F} \equiv \left(\mathrm{Tr}[\sqrt{\sqrt{\chi_{\mathrm{meas}}}\chi_{\mathrm{target}}\sqrt{\chi_{\mathrm{meas}}}}]\right)^2 / (\mathrm{Tr}[\chi_{\mathrm{meas}}]\,\mathrm{Tr}[\chi_{\mathrm{target}}])$, which measures the similarity between two processes [105]. The fidelity is 1 for perfect agreement and 0 for complete orthogonality. For example, quantum channels which do not alter the state in any way will have fidelity $\mathbb{F} = 1$ with the identity χ matrix, given above (one can simply read off the χ_{00} term to find the fidelity with the identity channel). Typically process fidelities are measured with respect to known or desired channels, but they may also be used to evaluate differences between two instances of the same channel as a distance measure. We apply this interpretation of the fidelity to various fiber channels to compare their resilience to different perturbations.

A.4.1 Measurement

The system for the SPT measurements is depicted in Fig. A.4. Light at 1550 nm from a diode laser (average power $\sim$ 10 mW) is passed through a polarizing beamsplitter (PBS) to set the input polarization state. The horizontally polarized photons are then passed through a combination of automatic waveplates (QWP,

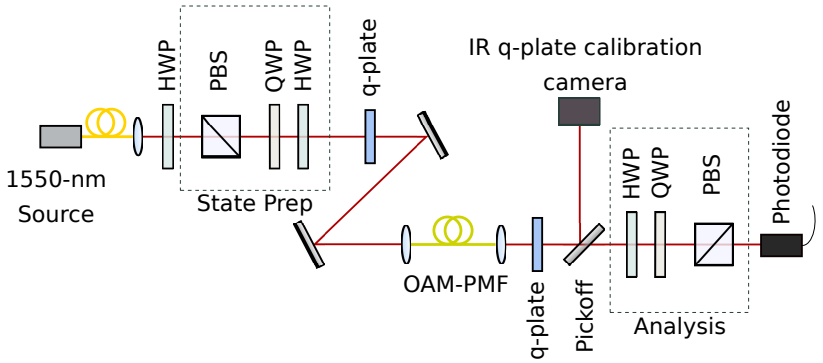

Figure A.4: Schematic for single process tomography (classical) measurements. The first half waveplate (HWP) is used to optimize power passed through the polarizing beamsplitter (PBS). Some measurements are made while the OAM-PMF fiber is heated, or the fiber is put through a large polarization controller in a 1-2-1 loop configuration

then HWP) to launch a polarization of our choosing into the q-plate and fiber. The fiber output is passed through a q-plate to transform back to polarization and analyzed by a HWP/QWP/PBS setup that mirrors the input state preparation. The photons are then coupled into SMF. In some measurements, the fiber was heated while coiled on a metal tray by a hotplate. In other measurements, the fiber was placed onto a fiber polarization controller (Thorlabs FPC560) in a 1-2-1 loop configuration. For SMF measurements, the q-plates and fiber are replaced by a 10-m sample of SMF.

Several sample measurements of the OAM-PMF and standard SMF are shown in Figs. A.5-A.8. Figs. A.5 and A.6 indicate reconstructed process matrices for the OAM fiber and the SMF under heating. Both cases show nearly identical χ during heating – even while stressed, in the SMF case ($F = 98.8 \pm 0.2\%$ for the OAM fiber, $96 \pm 1\%$ for the SMF). We did not perform any polarization correction for the SMF case, which results in the apparent rotation.

When the polarization controller is inserted, measurements are taken with all polarization controller paddles oriented vertically, and then (for reproducibility) with the paddles oriented at 45 degrees from the vertical, with a direction of displacement alternating with each paddle. These results are shown in Figs. A.7 and A.8. Evidently the polarization mapping of the OAM fiber is mostly unchanged ($F = 98 \pm 2\%$), while the polarization mapping of the SMF is scrambled to some extent ($F = 6 \pm 1\%$).

A.5 Results and Conclusion

Initially, it was our intention to measure the OAM fiber in several different configurations without heating or using a polarization controller. For instance, the simplest intended measurement was to take the loops of OAM fiber laying on the cardboard shelf and twist and fold them into twice the number of loops of half the size. When we attempted this measurement, we observed that the R and L polarizations were maintained,

but the rest (H,V,D,A) were scrambled, as we might expect from geometrical effects. Given that the $l = 7$ spin-orbit anti-aligned modes are 6 times more susceptible to geometric effects compared with SMF, being able to position the fiber exactly is critical, but proved prohibitively difficult in practice. We also attempted stretching the fiber loops into an elliptical shape and observed the same sensitivity to movement.

We were most surprised by the heat-resistance of the SMF, even when several meters of the standard SMF were wrapped around a 3.2-cm diameter disc and heated. We were able to demonstrate that the OAM-PMF preserves entanglement and polarization for one position in the holding tray, though we were unable to systematically study the effect of fiber perturbations on either of these properties. Experimentally, this was due to issues with the source and detector, but also, as evidenced by classical domain tests, the extreme sensitivity of the polarization to the geometric phase, which amplified any out-of-plane perturbations in the fiber. We conclude that, though the fiber does preserve polarization, the 7-fold increase in sensitivity to geometric effects garnered by mapping the input polarization to the $L = 7$ OAM modes rendered it impractical for tethering applications.

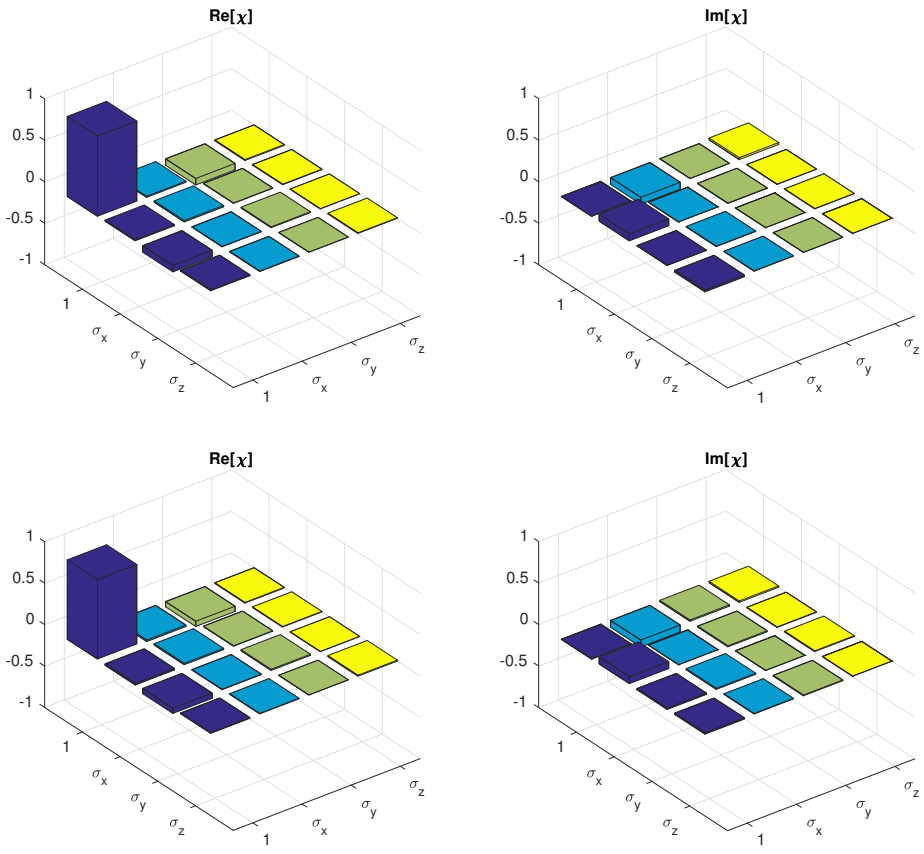

Figure A.5: **OAM-PMF** Reconstructed process matrix χ for the unheated fiber (top), and fiber at 50°C (bottom). The two processes are nearly identical within measurement noise (fidelity of $98.8 \pm 0.2\%$).

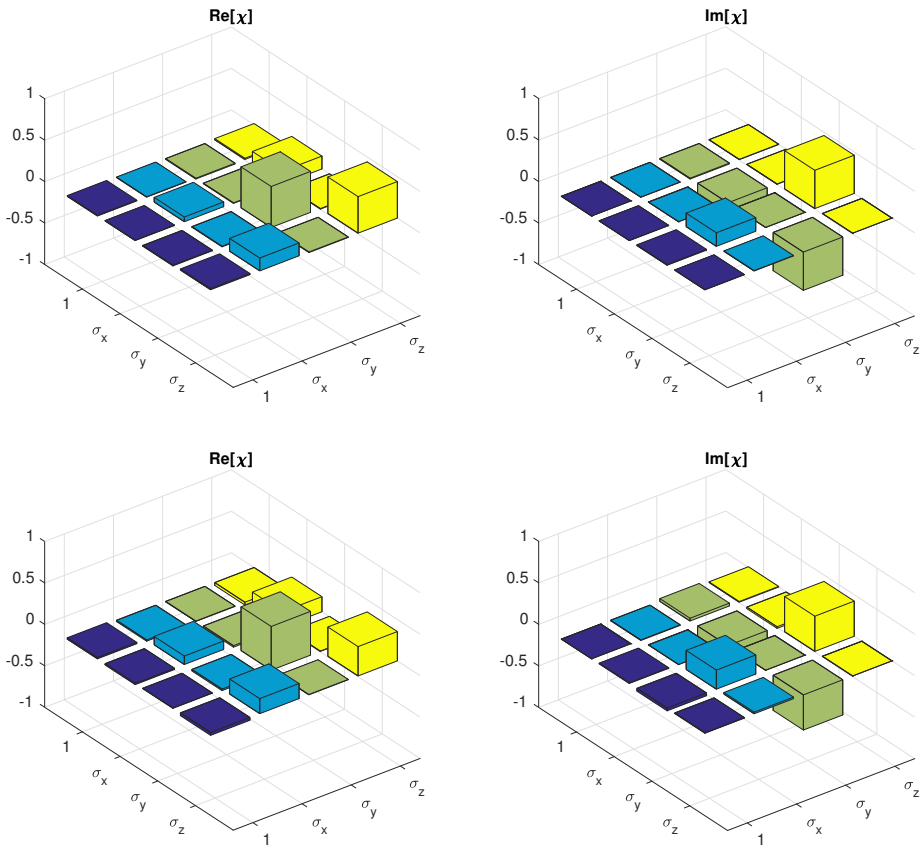

Figure A.6: **Standard SMF** Reconstructed process matrix χ for the unheated, stressed fiber (top), and stressed fiber at 60°C (bottom). The two processes are nearly identical, with measured before/after process fidelity of $96 \pm 1\%$.

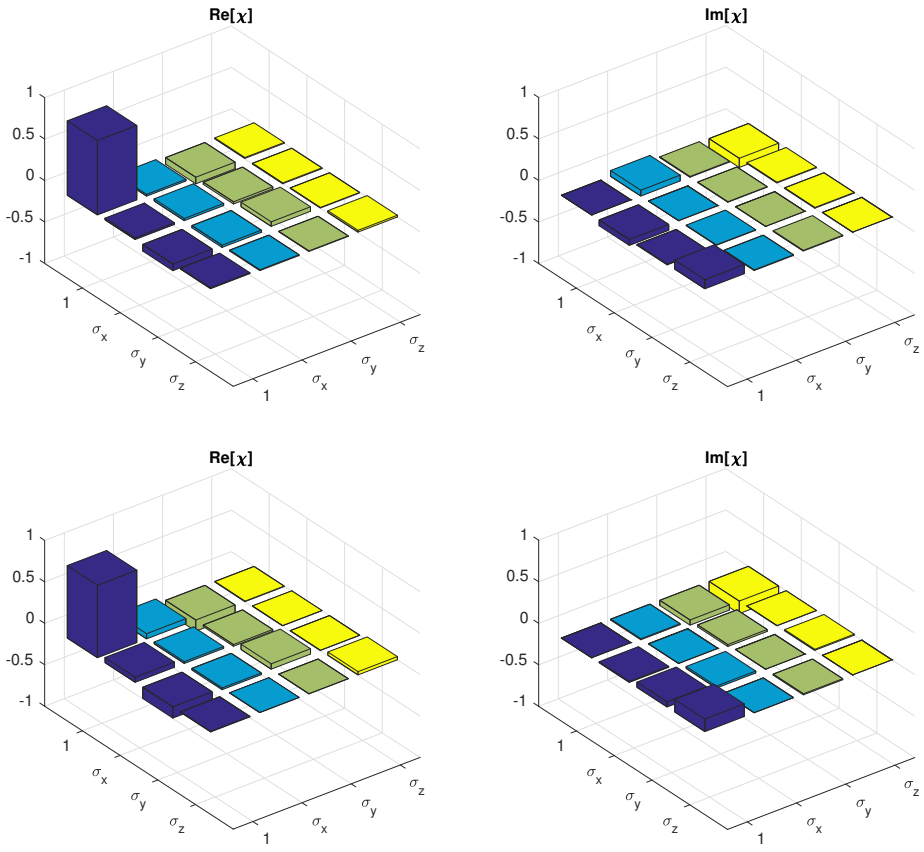

Figure A.7: **Polcon Test: OAM-PMF** (above) Vertical paddles (below) paddles at $\pm 45°$. Fidelity between measured χ matrices is $98 \pm 2\%$.

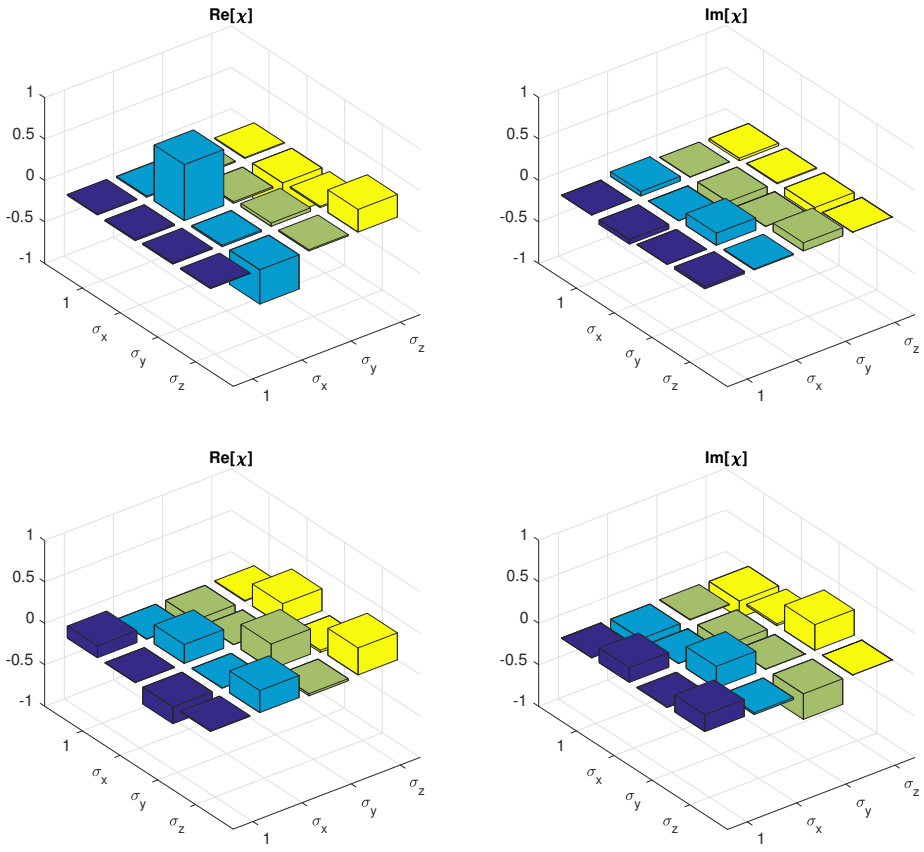

Figure A.8: **Polcon Test: Standard SMF** (above) Vertical paddles (below) paddles at 45°. Fidelity between measured χ matrices is $6 \pm 1\%$.

Appendix B

Propagate: A Program for Visualizing Turbulence-Induced Aberrations

B.1 Motivation

For free-space communication, quantum or otherwise, it is often critical to evaluate the characteristics of the scintillations introduced into the signal beam. Because turbulence in the atmosphere is random and evolves in time, it is typical to study free-space communications using the probability distribution of quantities of interest, e.g., the probability of fade (how often a beam fails to arrive within a certain aperture due to turbulence-induced beam wander), or the ability to couple a beam into a single- or multi-mode optical fiber [67]. It is often difficult, however, to gain a qualitative intuition about how scintillations change as the channel length increases or as the beam changes height as it propagates (e.g., satellite uplinks/downlinks). For example, it is certainly possible to derive the desired distribution from known parameter distributions, but for an intuitive understanding this can be too much of an undertaking, as the distributions are typically quite complex and situation-dependent. To help sample and visualize turbulence-induced scintillations in communication beams we have developed a simple program – called *Propagate* – that allows a user to change channel parameters and observe sample beam outputs from the chosen channel.

A screenshot of the program is shown in Figure. B.1. The user selects the resolution and number of screens to use (model precision), and the channel height profile: Flat, Curved (as over long distances over the Earth's surface), or satellite downlink/uplink. Each profile uses a model of how the atmospheric turbulence scales with height from sea level, and for models where it is relevant, the user can select either the traditional Hufnagel-Valley model [91] or a model based on data provided to us by the Space and Naval Warfare Systems Command (SPAWAR). The user may also select the beam launch height (valid for curved and flat profiles), which can drastically influence the structure of scintillations over long channels, as the beam may skim very close to sea level where turbulence is strongest. The most important parameters are the C_n^2 value, indicating the local turbulence strength in units of $m^{2/3}$, and the channel length in meters. The program then randomly generates phase screens from the appropriate distribution according to the techniques outlined in Chapter 3 (Fig. B.2) [71]. The program then simulates propagation of the beam using

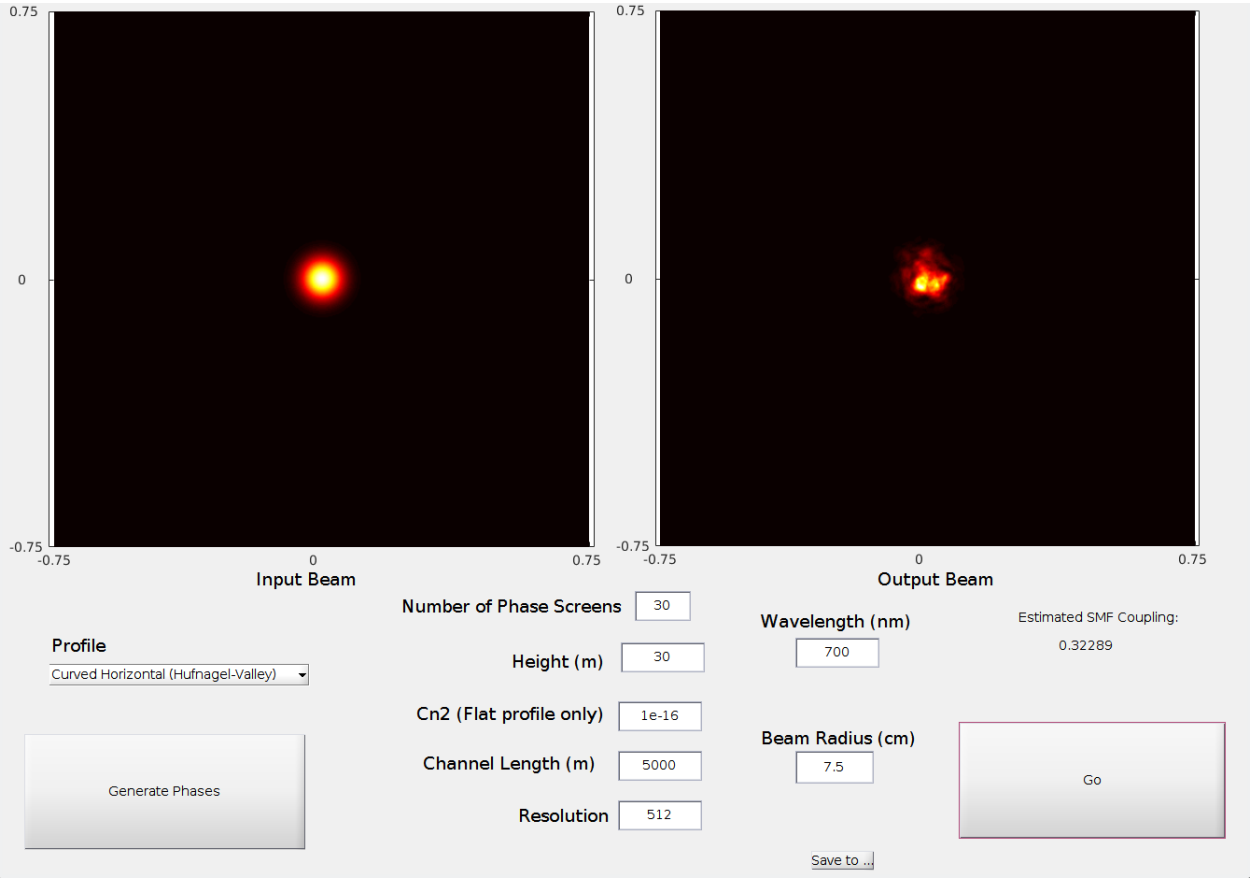

Figure B.1: Screenshot of Propagate

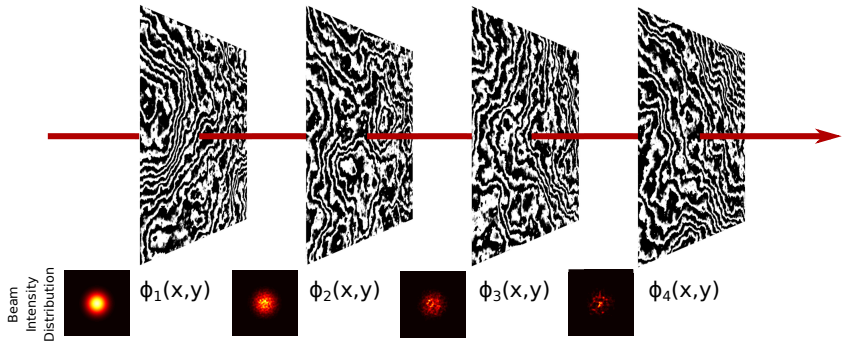

Figure B.2: Example 4-screen propagation over a curved horizontal profile, with a launch height of 30 m, propagation distance of 10 km, and a wavelength of 700 nm. As discussed in Chapter 3 and in Ref. [71], each phase screen is generated to mimic the accumulation of phase across the beam within each chunk of the channel. The harsh black and white trash featured in each screen is due to setting the phase to be modulo 2π for visualization.

numerical Fresnel integration between each phase screen. More phase screens result in a more accurate simulation because it is more similar to the actual propagation of light through infinitely many thin, weak turbulent layers. The program also outputs an estimated single-mode fiber coupling efficiency based on the mode defined by the input beam, and is able to stitch together multiple turbulence snapshots to form live video by interpolating between the phase fronts introduced in each layer, allowing users to gain an intuition for how beams may move over time in a non-frozen seeing scenario.

B.2 Example

Quantum communication to and from satellites has already been demonstrated [29] and is approaching viability for advanced entanglement-based protocols [106]. Gaining an understanding of the turbulence-induced losses for satellite to ground communication (and vice-versa) is critical for determining if certain noise-sensitive protocols are even possible. In Figure. B.3 we compare two different snapshots of the different in turbulence characteristics for downlinks versus uplinks. As turbulence is strongest near the surface of the Earth, the downlink case experiences much better overall scintillation strength as compared to the uplink case (the phases accumulated from the strongest turbulent layers have much less propagation distance to result in scintillations). Using 20 layers, the software estimates single-mode fiber coupling efficiencies (assuming perfect beam tracking and no atmospheric absorption or aerosol scattering) of around 0.001 - 0.1 for the downlink case, and around 10^{-5} for the uplink case. The program can also be used to generate statistics on these parameters, but it is intended to be a tool for visualization rather than establishing *a priori* estimable distributions.

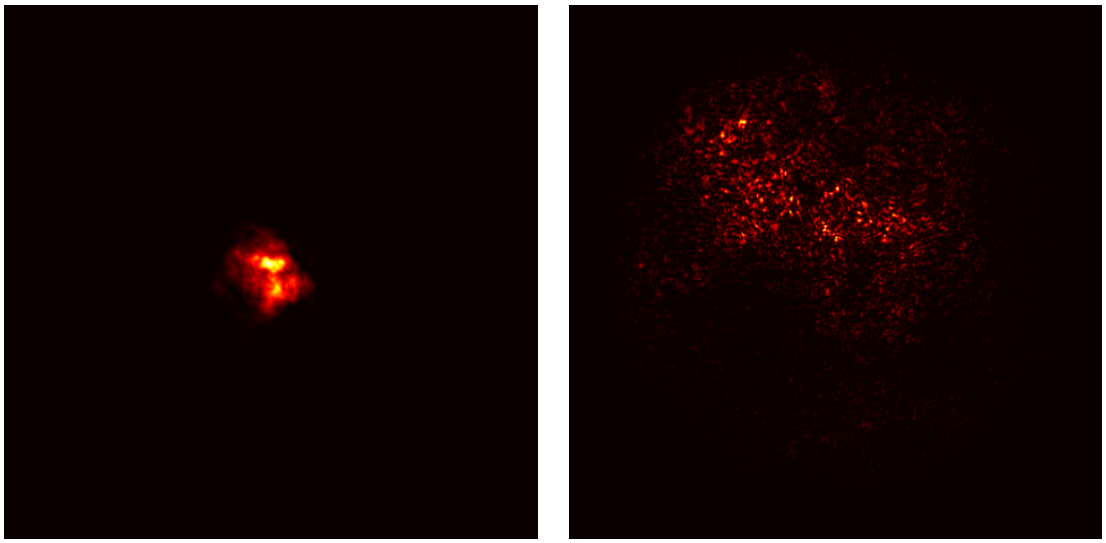

(a) (b)

Figure B.3: Example comparison of snapshots from (a) downlink and (b) uplink channel outputs (40-cm diameter beam at 700 nm over 120 km). Scintillations in the uplink are significantly worse than for the downlink.

Appendix C

Optical Stabilization Code (c_node)

While a daunting engineering challenge, achieving a robust lock between our transmitter QKD drone and receiver drone (Chapter 5) has required careful programming to ensure a high-speed, low-latency connection between all the interlocking systems in use. The software must 1) read local quadrant-cell position data, 2) transmit this data to the partner drone, 3) read data transmitted by the partner drone, and 4) execute the steering mirror corrections required to maintain lock. This requires a reasonably performance-minded approach as all code must run on a Raspberry PI microcomputer with finite resources. Here we include our stabilization code and annotate a few portions of its functionality.

```
// Preliminaries
#include <chrono>
#include <cstdlib>
#include <signal.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <thread>
#include <atomic>
#include <netdb.h>
#include <vector>
#include <atomic>
#include <string>
#include <sstream>
#include <iterator>
#include <fstream>
#include <sys/stat.h>
#include <fcntl.h>
#include <cmath>
#include "serialib.h"
#include "adc_lib.h"

#define  SER1_PORT   "/dev/ttyUSB1" // Open serial ports to fast steering mirrors
#define SER2_PORT "/dev/ttyUSB0"
std::atomic<char*> data;
char pos_buffer[256];

template <typename T> int sgn(T val) { // A function for returning the sign of numerical input
    return (T(0) < val) - (val < T(0));
}
```

```
int portno_tx, portno;
char UDP_IP[256];

bool terminateProgram = false;
void CtrlHandler(int signum) {
    terminateProgram = true;
}
```

The following method is run as a thread; it measures the three quadrant-cell detector voltages (V_{lr}, V_{tb}, and V_{tot}) and transmits them over a UDP connection over AC1750 WiFi to the partner drone.

```
void transmit_qc() {
    signal (SIGINT, CtrlHandler);
    signal (SIGQUIT, CtrlHandler);

    // Set up UDP connection
    int fd;
    if ((fd = socket(AF_INET, SOCK_DGRAM, IPPROTO_UDP)) < 0) {
        printf("cannot create socket");
        return;
    }
    struct hostent *hp;
    struct sockaddr_in servaddr;

    memset((char*)&servaddr, 0, sizeof(servaddr));
    servaddr.sin_family = AF_INET;
    servaddr.sin_port = htons(portno_tx);
    int size = 8192;
    setsockopt(fd, SOL_SOCKET, SO_SNDBUF,  &size, sizeof(int));

    hp = gethostbyname(UDP_IP);
    if (!hp) {
        fprintf(stderr, "could not obtain address of %s\n", UDP_IP);
        return;
    }
    char buffer[256];
    memcpy((void *)&servaddr.sin_addr, hp->h_addr_list[0], hp->h_length);

    // Begin TX operation
    printf("Tx Running\n");
    std::chrono::time_point<std::chrono::high_resolution_clock> t_now;
    long l_time;

    uint8_t id, ch_num;
    int32_t adc[8];
    float volt[8];
    int32_t iTemp;
    uint8_t buf[3];
    int error_count = 0;

    // Set up  ADC
    if (!bcm2835_init())
        return;
    bcm2835_spi_begin();
    bcm2835_spi_setBitOrder(BCM2835_SPI_BIT_ORDER_LSBFIRST );
    bcm2835_spi_setDataMode(BCM2835_SPI_MODE1);
    bcm2835_spi_setClockDivider(BCM2835_SPI_CLOCK_DIVIDER_1024);
    bcm2835_gpio_fsel(SPICS, BCM2835_GPIO_FSEL_OUTP);
    bcm2835_gpio_write(SPICS, HIGH);
    bcm2835_gpio_fsel(DRDY, BCM2835_GPIO_FSEL_INPT);
    bcm2835_gpio_set_pud(DRDY, BCM2835_GPIO_PUD_UP);
```

```
    ADS1256_CfgADC(ADS1256_GAIN_1, ADS1256_7500SPS);
    ADS1256_StartScan(0);
    ch_num = 8;
    int sendbuf_n;
    ADS1256_VAR_T* g_tADS1256 = get_state();
    float v0,v1,v2;

    // Begin reading and transmitting data
    while(!terminateProgram) {
        for( int i = 0; i < 3; i++) {
            while((DRDY_IS_LOW() == 0));
            ADS1256_SetChannel(i);        /*Switch channel mode */
                    bsp_DelayUS(25);
            adc[i] = ADS1256_ReadData();
            volt[i] = (adc[i] * 100.) / 167.;
            std::this_thread::sleep_for (std::chrono::microseconds(200));

        }

        v0 = volt[0]/1000000.;
        v1 = volt[1]/1000000.;
        v2 = volt[2]/1000000.;

        bzero(buffer,256);
        sendbuf_n = sendto(fd, buffer, strlen(buffer), 0, (struct sockaddr *)&servaddr,
        ↪  sizeof(servaddr));

        std::this_thread::sleep_for (std::chrono::milliseconds(7));
    }
    bcm2835_spi_end();
    bcm2835_close();

    return;

}
```

Below we reproduce the method, also run as a thread, that receives the data transmitted by the above method, `transmit_qc`, running on the partner drone. The data is stored in a thread-safe atomic variable.

```
void receive_data() {
    signal (SIGINT, CtrlHandler);
    signal (SIGQUIT, CtrlHandler);

    // Set up UDP connection
    struct sockaddr_in myaddr;
    struct sockaddr_in remaddr;
    socklen_t addrlen = sizeof(remaddr);
    int recvlen;
    int fd;
    char buf[256];
    if ((fd = socket(AF_INET, SOCK_DGRAM, IPPROTO_UDP)) < 0) {
        perror("cannot create socket\n");
        return;
    }
    memset((char *)&myaddr, 0, sizeof(myaddr));
    myaddr.sin_family = AF_INET;
    myaddr.sin_addr.s_addr = htonl(INADDR_ANY);
    myaddr.sin_port = htons(portno);

    if (bind(fd, (struct sockaddr *)&myaddr, sizeof(myaddr)) < 0) {
        perror("bind failed");
        return;
    }
    printf("Rx running...\n");
```

```
    // Receive data
    while(!terminateProgram) {
        recvlen = recvfrom(fd, buf, 256, 0, (struct sockaddr *)&remaddr, &addrlen);
        buf[recvlen] = 0;
        data.store(buf, std::memory_order_relaxed);
    }
    close(fd);

    return;
}
```

Below we reproduce the feedback thread. The receive thread updates then atomic variable `data`. `feedback` reads from `data` when it is able to gain thread-safe access. The thread then calculates an error and applies a correction by sending a serial message with the estimated mirror slew amount to the fast steering mirrors.

```
void feedback() {
    signal (SIGINT, CtrlHandler);
    signal (SIGQUIT, CtrlHandler);

    // Read proportional gain values from config file
    float KpX = 0.0;
    float KpY = 0.0;
    std::ifstream configFile;
    configFile.open("./pidconfig.txt");
    if(!configFile) {
        printf("Unable to open config file\n");
        exit(0);
    }
    int j = 0;
    configFile >> KpX;
    configFile >> KpY;

    configFile.close();
    printf("||||| K: %f %f\n",KpX,KpY);
    printf("Feedback loop\n");

    // Initialize serial interfaces and variables
    serialib ser1, ser2;
    int ret;
    char ser_buf[128];
    if (ret != 1) {                                                                    /
        printf ("Error while opening port. Permission problem ?\n");
        return;
    }

    printf ("Serial port opened successfully !\n");

    ret = ser2.Open(SER2_PORT, 57600);
    if (ret != 1) {
        printf ("Error while opening port. Permission problem ?\n");
        return;
    }
    printf ("Serial port opened successfully !\n");

    ser1.WriteString("C33H1024c;");
    ser2.WriteString("C33H1024c;");

    char* buf;
```

```
int errorX,errorY;

float KdX = 0.0;
float KdY = 0.0;
float DARK_LEVEL = 0.1;
float value1,value2,value3,v1norm,v2norm,v1_v,v2_v,v1norm_old,v2norm_old,v3_old;
std::string errXstr, errYstr;
std::vector<float> v;
printf("Loop running...\n");
std::string bufstr;
int d = 0;

float Ix = 0.0;
float Iy = 0.0;

timespec deadline;
deadline.tv_sec = 0;
deadline.tv_nsec = 10000000;
bool signal_lost = true;

// Begin feedback loop
while(!terminateProgram) {
        // Read latest data updated by receive_data()
    buf = data.load(std::memory_order_relaxed);
    if(buf == NULL) {
        clock_nanosleep(CLOCK_REALTIME,0,&deadline,NULL);
        continue;
    }
    bufstr = std::string(buf);
    if(bufstr == "") {

    }

    std::istringstream processbuf(bufstr);

    std::copy(std::istream_iterator<float>(processbuf),
    std::istream_iterator<float>(),
    std::back_inserter(v));

    value1 = v[0];
    value2 = v[1];
    value3 = v[2];
    v.clear();
    value1 = float(value1)-2.385; // Convert voltages to errors
    value2 = float(value2)-2.385;
    value3 = float(value3);

    v1norm = value1/value3; // Normalize to total incident power
    v2norm = value2/value3;

    v1norm = sgn(v1norm)*pow(fabs(v1norm),7.0/5.); // Compute inflected error
    v2norm = sgn(v2norm)*pow(fabs(v2norm),7.0/5.);

    errorX = int(round(v1norm * KpX  )); // Scale by X gain
    errorY = int(round(v2norm * KpY ));  // Scale by Y gain

    v1norm_old = v1norm;
    v2norm_old = v2norm;

    if(value3 < DARK_LEVEL) {
        printf("LOST SIGNAL\n");
        clock_nanosleep(CLOCK_REALTIME, 0, &deadline, NULL);
        signal_lost = true;
        continue;
```

```
        }
        if(signal_lost) {
                printf("LOCKED...\n");
        }
        signal_lost = false;
        // Format error as mirror command
        errXstr = std::to_string(errorX);
        errYstr = std::to_string(errorY);

        if(errorX > 0) {
            errXstr = "+"+errXstr;
        } else {
            errXstr = errXstr;
        }
        if(errorY > 0) {
            errYstr = "+"+errYstr;
        } else {
            errYstr = errYstr;
        }


        // Write mirror command over serial
        if(errorX != 0) {
            ser2.WriteString((errXstr+";").c_str());
        }
        if(errorY != 0) {
            ser1.WriteString((errYstr+";").c_str());
        }

        v1norm_old = v1norm;
        v2norm_old = v2norm;
        v3_old = value3;

        clock_nanosleep(CLOCK_REALTIME, 0, &deadline, NULL); // Wait
    }
    ser1.close();
    ser2.close();
    return;
}
```

The following `main` method is the entry point to the program and instantiates all threads.

```
int main(int argc, char** argv) {
    signal (SIGINT, CtrlHandler);
    signal (SIGQUIT, CtrlHandler);

    memcpy(UDP_IP, argv[1],256);
    portno = atoi(argv[2]);
    portno_tx = atoi(argv[3]);


    printf("%s\t%i\t%i\n",UDP_IP,portno,portno_tx);

    std::thread t_transmit_qc(transmit_qc);
    std::thread t_receive(receive_data);
    std::thread t_feedback(feedback);


    t_transmit_qc.join();
    t_receive.join();
    t_feedback.join();

    return 0;
}
```

Appendix D

Tuned interference for long-distance optical communication

D.1 Background & Motivation

Diffraction of optical beams over the extremely long communication channel lengths encountered in space communications (e.g., from Earth to a deep-space satellite) can result in significant loss, as the collection area required for a beam launched from a distant satellite may reach into the scale of kilometers squared. Furthermore, the beam-forming optics on the transmitter are necessarily small to conserve size, weight, and power, which makes diffractive losses even worse.

One solution to this problem is to introduce a constellation of beams of a small diameter that sum coherently at the receiver to produce a spot much smaller than the diffraction-limited beam size of a single beam at the receiver. This requires the positions and relative phases of the source constellation to be carefully chosen to focus as much light as possible into a small area at the receiver plane. The effectiveness of this strategy depends on the number of apertures used and the degree to which their phases can be controlled precisely. This process is analogous to long-baseline interferometry [107, 108], whereby a telescope consisting of a network of small, phase-coherent telescopes can function as a single telescope with an aperture size proportional to the size of the array.

We propose the following optimization problem: given a set of N co-propagating Gaussian beams of width σ^2, what are the optimal positions $(\delta_x^{(i)}, \delta_y^{(i)})$ and phases $\phi^{(i)}$ of the beams at the transmitter plane such that the power at the receiver plane is maximized within a circular aperture of radius R_f?

The problem of the diffraction pattern produced from an arbitrary source distribution at some measurement plane is, in general, intractable. Numerical evaluations of the Fresnel integral suffice, but only when the simulation can be performed with sufficient resolution over the entire path length. For space applications, the relative sizes of the source beam radii, their distances from one another, and the propagation path length are all different enough that producing an accurate numerical approximation of Fresnel diffraction is not a reasonable approach. For source distributions that are sums of functions with analytically calculable Fresnel diffraction patters, however, the receiver-plane pattern can be taken as a coherent sum of the electric fields

produced by each source individually.

Assume each source is a Gaussian beam with spatial distribution

$$E(\delta_x, \delta_y) \propto \exp\left\{\frac{-((\delta_x - x')^2 + (\delta_y - y')^2)}{2\sigma^2}\right\}. \tag{D.1}$$

Then the Frenel diffraction pattern at a plane a distance z from the origin is given by

$$F = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} A \exp\left\{\frac{-((\delta_x - x')^2 + (\delta_y - y')^2)}{2\sigma^2}\right\} \exp\left\{\left(\frac{ik}{2z}\right)((x - x')^2 + (y - y')^2)\right\} dx'dy', \tag{D.2}$$

which has the closed-form solution

$$F(x, \delta_x, y, \delta_y, k, \sigma^2, z) = A\frac{2\pi\sigma z}{(z - ik\sigma^2)} \exp\left\{\frac{-k[(\delta_x - x)^2 + (\delta_y - y)^2)]}{2k\sigma^2 + 2iz}\right\}. \tag{D.3}$$

The complete field is then given by

$$E(x, y, z) = \sum_{i=1}^{N} e^{i\phi^{(i)}} F(x, \delta_x^{(i)}, y, \delta_y^{(i)}, k, \sigma^{2(i)}, z), \tag{D.4}$$

making the power into the final aperture

$$P(A) \sim \int\int_A \left|\sum_{i=1}^{N} e^{i\phi^{(i)}} F(x, \delta_x^{(i)}, y, \delta_y^{(i)}, k, \sigma^{2(i)}, z)\right|^2. \tag{D.5}$$

The goal of the optimization is to choose $\delta_x^{(i)}, \delta_y^{(i)}, \phi^{(i)}$, and $\sigma^{2(i)}$ so that P is maximized for a chosen z and receiver aperture size A.

D.2 Preliminary Testing

We applied this idea using the built-in genetic algorithm that ships with MATLAB. Only the positions and phases of the sources were optimized in our test, in order to reduce the size of the search space. The inputs/outputs for one such run are shown in Figure. D.1. We expect that the ideal solution should possess some form of radial symmetry – or at least X/Y mirror symmetry – as the output aperture is circular. This algorithm, however, was unable to converge to such a solution without forcing that form of symmetry as a precondition for optimization. By imposing such symmetry, a radially-symmetric output may be obtained; a future deep search would involve optimizing over all parameters, however, to find the optimal transmitter aperture number, positions, and launch phases.

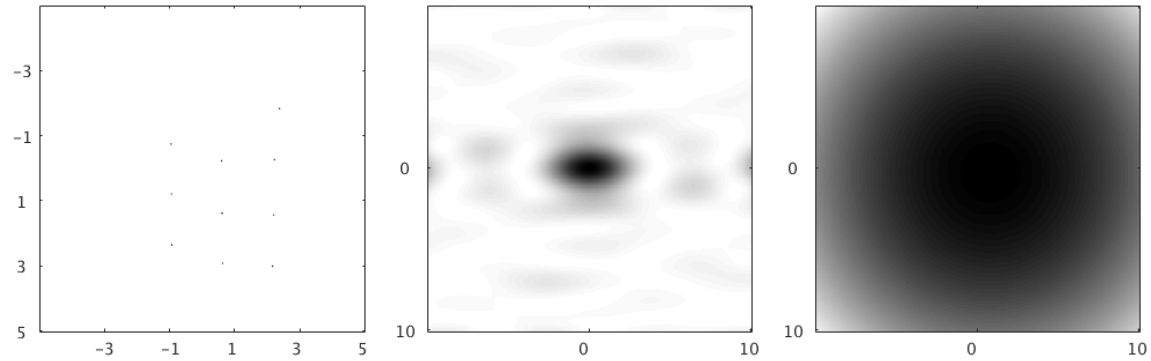

(a) (b) (c)

Figure D.1: Sample input (a) and output (b) at 20,000 km, compared with the incoherent output (c) at the same plane. Phases not plotted. The horizontal and vertical scales are in meters. We would expect some form of radial symmetry, but the genetic algorithm included in MATLAB was unable to converge to such a solution.

References

[1] M. Mosca, "Cybersecurity in an era with quantum computers: will we be ready?." Cryptology ePrint Archive, Report 2015/1075, 2015.

[2] "Google Quantum AI Team (`https://ai.google/research/teams/applied-science/quantum-ai/`)."

[3] "IBM Quantum Computing (`https://www.research.ibm.com/ibm-q/`)."

[4] "Microsoft Quantum Computing (`https://www.microsoft.com/en-us/quantum/`)."

[5] L. K. Grover, "A fast quantum mechanical algorithm for database search." `arXiv:quant-ph/9605043`, 1996.

[6] P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *SIAM Journal on Scientific and Statistical Computing*, vol. 26, no. 2, p. 1484, 1995.

[7] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, "Cloud quantum computing of an atomic nucleus," *Phys. Rev. Lett.*, vol. 120, p. 210501, May 2018.

[8] D. S. Abrams and S. Lloyd, "Simulation of many-body fermi systems on a universal quantum computer," *Phys. Rev. Lett.*, vol. 79, pp. 2586–2589, Sep 1997.

[9] C. C. McGeoch and C. Wang, "Experimental evaluation of an adiabiatic quantum system for combinatorial optimization," in *Proceedings of the ACM International Conference on Computing Frontiers*, CF '13, (New York, NY, USA), pp. 23:1–23:11, ACM, 2013.

[10] H. J. Kimble, "The quantum internet," *Nature*, vol. 453, pp. 1023 EP –, 06 2008.

[11] J. L. Park, "The concept of transition in quantum mechanics," *Foundations of Physics*, vol. 1, pp. 23–33, Mar 1970.

[12] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," *Nature*, vol. 299, pp. 802 EP –, 10 1982.

[13] J. F. Fitzsimons, "Private quantum computation: an introduction to blind quantum computing and related protocols," *npj Quantum Information*, vol. 3, no. 1, p. 23, 2017.

[14] D. Adrian, K. Bhargavan, Z. Durumeric, P. Gaudry, M. Green, J. A. Halderman, N. Heninger, D. Springall, E. Thomé, L. Valenta, B. Vandersloot, E. Wustrow, and S. Z.-b. Paul, "Imperfect Forward Secrecy : How Diffie-Hellman Fails in Practice," 2015.

[15] D. Bernstein, J. Buchmann, and E. Dahmen, *Post-Quantum Cryptography*, vol. 1. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009.

[16] F. Miller, *Telegraphic Code to Insure Privacy and Secrecy in the Transmission of Telegrams.* C.M. Cornwell, 1882.

[17] G. Vernam, "Secret signaling system," 1919. US Patent 1,310,719.

[18] C. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pp. 175–179, 1984.

[19] H.-K. Lo, X. Ma, and K. Chen, "Decoy state quantum key distribution," *Phys. Rev. Lett.*, vol. 94, p. 230504, Jun 2005.

[20] C. H. Bennett, G. Brassard, and N. D. Mermin, "Quantum cryptography without bell's theorem," *Phys. Rev. Lett.*, vol. 68, pp. 557–559, Feb 1992.

[21] J. S. Bell, "On the Einstein Podolsky Rosen paradox," *Physics*, vol. 1, no. 3, pp. 195–200, 1964.

[22] A. K. Ekert, "Quantum cryptography based on bell's theorem," *Phys. Rev. Lett.*, vol. 67, pp. 661–663, Aug 1991.

[23] C. E. Shannon, "A Mathematical Theory of Communication," vol. 27, no. July 1928, pp. 379–423, 1948.

[24] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, "The security of practical quantum key distribution," *Reviews of Modern Physics*, vol. 81, no. 3, pp. 1301–1350, 2009.

[25] P. W. Shor and J. Preskill, "Simple proof of security of the BB84 quantum key distribution protocol," *Physical Review Letters*, vol. 85, no. 2, pp. 441–444, 2000.

[26] H.-K. Lo, "Proof of unconditional security of six-state quantum key distribution scheme." `arXiv:quant-ph/0102138`, 2001.

[27] I. Cisco Systems, "Introduction to optical fibers, dB, attenuation and measurements," July 2018.

[28] J.-P. Bourgoin, B. L. Higgins, N. Gigov, C. Holloway, C. J. Pugh, S. Kaiser, M. Cranmer, and T. Jennewein, "Free-space quantum key distribution to a moving receiver," *Opt. Express*, vol. 23, pp. 33437–33447, Dec 2015.

[29] S.-K. Liao, W.-Q. Cai, W.-Y. Liu, L. Zhang, Y. Li, J.-G. Ren, J. Yin, Q. Shen, Y. Cao, Z.-P. Li, F.-Z. Li, X.-W. Chen, L.-H. Sun, J.-J. Jia, J.-C. Wu, X.-J. Jiang, J.-F. Wang, Y.-M. Huang, Q. Wang, Y.-L. Zhou, L. Deng, T. Xi, L. Ma, T. Hu, Q. Zhang, Y.-A. Chen, N.-L. Liu, X.-B. Wang, Z.-C. Zhu, C.-Y. Lu, R. Shu, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan, "Satellite-to-ground quantum key distribution," *Nature*, vol. 549, pp. 43 EP –, 08 2017.

[30] H. Ko, B.-S. Choi, J.-S. Choe, K.-J. Kim, J.-H. Kim, and C. J. Youn, "Critical side channel effects in random bit generation with multiple semiconductor lasers in a polarization-based quantum key distribution system," *Opt. Express*, vol. 25, pp. 20045–20055, Aug 2017.

[31] M. S. Lee, M. K. Woo, J. Jung, Y.-S. Kim, S.-W. Han, and S. Moon, "Free-space qkd system hacking by wavelength control using an external laser," *Opt. Express*, vol. 25, pp. 11124–11131, May 2017.

[32] M. Tomamichel, C. C. W. Lim, N. Gisin, and R. Renner, "Tight finite-key analysis for quantum cryptography," *Nature Communications*, vol. 3, pp. 634 EP –, 01 2012.

[33] N. T. Islam, C. C. W. Lim, C. Cahall, J. Kim, and D. J. Gauthier, "Provably secure and high-rate quantum key distribution with time-bin qudits," *Science Advances*, vol. 3, no. 11, 2017.

[34] H.-L. Yin, T.-Y. Chen, Z.-W. Yu, H. Liu, L.-X. You, Y.-H. Zhou, S.-J. Chen, Y. Mao, M.-Q. Huang, W.-J. Zhang, H. Chen, M. J. Li, D. Nolan, F. Zhou, X. Jiang, Z. Wang, Q. Zhang, X.-B. Wang, and J.-W. Pan, "Measurement-device-independent quantum key distribution over a 404 km optical fiber," *Phys. Rev. Lett.*, vol. 117, p. 190501, Nov 2016.

[35] P. L. McMahon and K. De Greve, *Towards Quantum Repeaters with Solid-State Qubits: Spin-Photon Entanglement Generation Using Self-assembled Quantum Dots*, pp. 365–402. Cham: Springer International Publishing, 2015.

[36] C. Laplane, P. Jobez, J. Etesse, N. Gisin, and M. Afzelius, "Multimode and long-lived quantum correlations between photons and spins in a crystal," *Phys. Rev. Lett.*, vol. 118, p. 210501, May 2017.

[37] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin, "Quantum repeaters based on atomic ensembles and linear optics," *Rev. Mod. Phys.*, vol. 83, pp. 33–80, Mar 2011.

[38] L. M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics," *Nature*, vol. 414, pp. 413 EP –, 11 2001.

[39] K. Kutluer, M. Mazzera, and H. de Riedmatten, "Solid-state source of nonclassical photon pairs with embedded multimode quantum memory," *Phys. Rev. Lett.*, vol. 118, p. 210502, May 2017.

[40] A. Seri, A. Lenhard, D. Rieländer, M. Gündoğan, P. M. Ledingham, M. Mazzera, and H. de Riedmatten, "Quantum correlations between single telecom photons and a multimode on-demand solid-state quantum memory," *Phys. Rev. X*, vol. 7, p. 021028, May 2017.

[41] J. Simon, H. Tanji, J. K. Thompson, and V. Vuletić, "Interfacing collective atomic excitations and single photons," *Phys. Rev. Lett.*, vol. 98, p. 183601, May 2007.

[42] A. D. Hill, D. Hervas, J. Nash, M. Graham, A. Burgers, U. Paudel, D. Steel, C. Schneider, M. Kamp, S. Höfling, J. Wang, J. Lin, W. Zhao, and P. G. Kwiat, "Optimizing single-mode collection from pointlike sources of single photons with adaptive optics," *Opt. Express*, vol. 25, pp. 18629–18642, Aug 2017.

[43] J. D. Sterk, L. Luo, T. a. Manning, P. Maunz, and C. Monroe, "Photon collection from a trapped ion-cavity system," *Phys. Rev. A*, vol. 85, p. 062308, jun 2012.

[44] R. Noek, G. Vrijsen, D. Gaultney, E. Mount, T. Kim, P. Maunz, and J. Kim, "High Speed, High Fidelity Detection of an Atomic Hyperfine Qubit," *Opt. Lett.*, vol. 38, no. 22, p. 4, 2013.

[45] W. B. Gao, a. Imamoglu, H. Bernien, and R. Hanson, "Coherent manipulation, measurement and entanglement of individual solid-state spins using optical fields," *Nat. Photonics*, vol. 9, no. 6, pp. 363–373, 2015.

[46] K. H. Madsen, S. Ates, J. Liu, A. Javadi, S. M. Albrecht, I. Yeo, S. Stobbe, and P. Lodahl, "Efficient out-coupling of high-purity single photons from a coherent quantum dot in a photonic-crystal cavity," *Phys. Rev. B*, vol. 90, no. 15, p. 155303, 2014.

[47] M. Arcari, I. Söllner, A. Javadi, S. Lindskov Hansen, S. Mahmoodian, J. Liu, H. Thyrrestrup, E. H. Lee, J. D. Song, S. Stobbe, and P. Lodahl, "Near-Unity Coupling Efficiency of a Quantum Emitter to a Photonic Crystal Waveguide," *Phys. Rev. Lett.*, vol. 113, no. 9, pp. 1–5, 2014.

[48] G. M. Akselrod, C. Argyropoulos, T. B. Hoang, C. Ciracì, C. Fang, J. Huang, D. R. Smith, and M. H. Mikkelsen, "Probing the mechanisms of large Purcell enhancement in plasmonic nanoantennas," *Nat. Photonics*, vol. 8, no. 11, pp. 835–840, 2014.

[49] N. Somaschi, V. Giesz, L. De Santis, J. C. Loredo, M. P. Almeida, G. Hornecker, S. L. Portalupi, T. Grange, C. Anton, J. Demory, C. Gomez, I. Sagnes, N. D. L. Kimura, A. Lemaitre, A. Auffeves, A. G. White, L. Lanco, and P. Senellart, "Near optimal single photon sources in the solid state," *Nat. Photonics*, vol. 10, no. 2, pp. 1–6, 2015.

[50] J. D. Wong-Campos, K. G. Johnson, B. Neyenhuis, J. Mizrahi, and C. Monroe, "High-resolution adaptive imaging of a single atom," *Nat. Photonics*, vol. 10, pp. 606–610, jul 2016.

[51] A. Tiranov, P. C. Strassmann, J. Lavoie, N. Brunner, M. Huber, V. B. Verma, S. W. Nam, R. P. Mirin, A. E. Lita, F. Marsili, M. Afzelius, F. Bussières, and N. Gisin, "Temporal multimode storage of entangled photon pairs," *Phys. Rev. Lett.*, vol. 117, p. 240506, Dec 2016.

[52] G. N. M. Tabia, "Recursive multiport schemes for implementing quantum algorithms with photonic integrated circuits," *Phys. Rev. A*, vol. 93, p. 012323, Jan 2016.

[53] M. L. Plett, P. R. Barbier, and D. W. Rush, "Compact adaptive optical system based on blind optimization and a micromachined membrane deformable mirror.," *Appl. Optics*, vol. 40, pp. 327–330, jan 2001.

[54] A. Courteville, "Optimization of single-mode fiber coupling efficiency with an adaptive membrane mirror," *Opt. Eng.*, vol. 41, p. 1073, may 2002.

[55] P. Villoresi, S. Bonora, M. Pascolini, L. Poletto, G. Tondello, C. Vozzi, M. Nisoli, G. Sansone, S. Stagira, and S. De Silvestri, "Optimization of high-order harmonic generation by adaptive control of a sub-10-fs pulse wave front.," *Opt. Lett.*, vol. 29, pp. 207–209, jan 2004.

[56] C. Bonato, S. Bonora, A. Chiuri, P. Mataloni, G. Milani, G. Vallone, and P. Villoresi, "Phase control of a path-entangled photon state by a deformable membrane mirror," *J. Opt. Soc. Am. B*, vol. 27, no. 6, 2010.

[57] M. Minozzi, S. Bonora, a. V. Sergienko, G. Vallone, and P. Villoresi, "Optimization of two-photon wave function in parametric down conversion by adaptive optics control of the pump radiation.," *Opt. Lett.*, vol. 38, pp. 489–491, feb 2013.

[58] V. N. Mahajan, "Strehl ratio of a Gaussian beam.," *J. Opt. Soc. Am. A*, vol. 22, no. 9, pp. 1824–1833, 2005.

[59] A. P. Burgers, J. R. Schaibley, and D. G. Steel, "Entanglement and Quantum Optics with Quantum Dots," in *From Atomic to Mesoscale*, pp. 103–120, World Scientific, aug 2015.

[60] W. B. Gao, P. Fallahi, E. Togan, J. Miguel-Sanchez, and A. Imamoglu, "Observation of entanglement between a quantum dot spin and a single photon," *Nature*, vol. 491, no. 7424, pp. 426–430, 2012.

[61] K. De Greve, L. Yu, P. L. McMahon, J. S. Pelc, C. M. Natarajan, N. Y. Kim, E. Abe, S. Maier, C. Schneider, M. Kamp, S. Höfling, R. H. Hadfield, A. Forchel, M. M. Fejer, and Y. Yamamoto, "Quantum-dot spin–photon entanglement via frequency downconversion to telecom wavelength," *Nature*, vol. 491, pp. 421–425, nov 2012.

[62] J. R. Schaibley, A. P. Burgers, G. A. McCracken, L.-M. Duan, P. R. Berman, D. G. Steel, A. S. Bracker, D. Gammon, and L. J. Sham, "Demonstration of quantum entanglement between a single electron spin confined to an inas quantum dot and a photon," *Phys. Rev. Lett.*, vol. 110, p. 167401, Apr 2013.

[63] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres," *Nature*, vol. 526, pp. 682–686, oct 2015.

[64] J. Enderlein, "Theoretical study of detection of a dipole emitter through an objective with high numerical aperture," *Opt. Lett.*, vol. 25, pp. 634–636, May 2000.

[65] R. Maiwald, A. Golla, M. Fischer, M. Bader, S. Heugel, B. Chalopin, M. Sondermann, and G. Leuchs, "Collecting more than half the fluorescence photons from a single ion," *Phys. Rev. A*, vol. 86, p. 043431, Oct 2012.

[66] A. Kolmogorov, "The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers," *Dokl. Akad. Nauk SSSR*, vol. 434, no. 1890, pp. 9–13, 1941.

[67] L. C. Andrews and R. L. Phillips, *Laser Beam Propagation through Random Media*. 1000 20th Street, Bellingham, WA 98227-0010 USA: SPIE, 2005.

[68] D. L. Fried, "Spectral and angular covariance of scintillation for propagation in a randomly inhomogeneous medium.," *Applied optics*, vol. 10, no. 4, pp. 721–731, 1971.

[69] W. B. Miller, J. C. Ricklin, and L. C. Andrews, "Effects of the refractive index spectral model on the irradiance variance of a Gaussian beam," *Journal of the Optical Society of America A*, vol. 11, p. 2719, October 1994.

[70] S. S. Chesnokov, V. P. Kandidov, V. I. Shmalhausen, and V. V. Shuvalov, "Numerical/optical simulation of laser beam propagation through atmospheric turbulence.," tech. rep., DTIC Document, 1995.

[71] J. D. Schmidt, *Numerical Simulation of Optical Wave Propagation with Examples in* ***MATLAB***. SPIE, 2010.

[72] L. Burger, I. A. Litvin, and A. Forbes, "Simulating atmospheric turbulence using a phase-only spatial light modulator," *South African Journal of Science*, vol. 104, pp. 129 – 134, 04 2008.

[73] G. R. Boyer, B. Lamouroux, and B. S. Prade, "Atmospheric birefringence under wind speed gradient shear∗," *J. Opt. Soc. Am.*, vol. 68, pp. 471–474, Apr 1978.

[74] W.-H. Lee, "Iii computer-generated holograms: Techniques and applications," vol. 16 of *Progress in Optics*, pp. 119 – 232, Elsevier, 1978.

[75] "How to use a binary amplitude deformable miror device (dmd) as a phase modulator: Lee hologram method."

[76] E. Anzuola and A. Belmonte, "Generation of atmospheric wavefronts using binary micromirror arrays," *Appl. Opt.*, vol. 55, pp. 3039–3044, Apr 2016.

[77] R. Rampy, D. Gavel, D. Dillon, and S. Thomas, "Production of phase screens for simulation of atmospheric turbulence," *Appl. Opt.*, vol. 51, pp. 8769–8778, Dec 2012.

[78] G. Vallone, D. G. Marangon, M. Canale, I. Savorgnan, D. Bacco, M. Barbieri, S. Calimani, C. Barbieri, N. Laurenti, and P. Villoresi, "Adaptive real time selection for quantum key distribution in lossy and turbulent free-space channels," *Phys. Rev. A*, vol. 91, p. 042320, Apr 2015.

[79] G. Vallone, D. G. Marangon, M. Canale, I. Savorgnan, D. Bacco, M. Barbieri, S. Calimani, C. Barbieri, N. Laurenti, and P. Villoresi, "Adaptive real time selection for quantum key distribution in lossy and turbulent free-space channels," *Phys. Rev. A*, vol. 91, p. 042320, Apr 2015.

[80] W. Wang, F. Xu, and H.-K. Lo, "Prefixed-threshold real-time selection method in free-space quantum key distribution," *Phys. Rev. A*, vol. 97, p. 032337, Mar 2018.

[81] B. G. Christensen, *Advanced tests of nonlocality with entangled photons*. PhD thesis, University of Illinois at Urbana-Champaign, 2016.

[82] P. G. Kwiat, K. Mattle, H. Weinfurter, A. Zeilinger, A. V. Sergienko, and Y. Shih, "New high-intensity source of polarization-entangled photon pairs," *Phys. Rev. Lett.*, vol. 75, pp. 4337–4341, Dec 1995.

[83] M. Wayne, *Optical Quantum Random Number Generation: Applications of Single-Photon Event Timing*. PhD thesis, University of Illinois at Urbana-Champaign, 2017.

[84] R. J. McIntyre, "Multiplication noise in uniform avalanche diodes," *IEEE Transactions on Electron Devices*, vol. ED-13, pp. 164–168, Jan 1966.

[85] M. C. T. Bahaa E. A. Saleh, *Fundamentals of Photons*. Wiley, 2nd edition ed., 2007.

[86] R. K. Franz Aurenhammer, "Chapter 5 - voronoi diagrams," in *Handbook of Computational Geometry* (J.-R. Sack and J. Urrutia, eds.), pp. 201 – 290, Amsterdam: North-Holland, 2000.

[87] C. J. Pugh, S. Kaiser, J.-P. Bourgoin, J. Jin, N. Sultana, S. Agne, E. Anisimova, V. Makarov, E. Choi, B. L. Higgins, and T. Jennewein, "Airborne demonstration of a quantum key distribution receiver payload," *Quantum Science and Technology*, vol. 2, no. 2, p. 024009, 2017.

[88] H. Chun, I. Choi, G. Faulkner, L. Clarke, B. Barber, G. George, C. Capon, A. Niskanen, J. Wabnig, D. O'Brien, and D. Bitauld, "Handheld free space quantum key distribution with dynamic motion compensation," *Opt. Express*, vol. 25, pp. 6784–6795, Mar 2017.

[89] G. Mélen, P. Freiwang, J. Luhn, T. Vogl, M. Rau, W. Rosenfeld, and H. Weinfurter, "Handheld quantum key distribution," in *Conference on Lasers and Electro-Optics*, p. FTu3G.1, Optical Society of America, 2018.

[90] G. Vest, M. Rau, L. Fuchs, G. Corrielli, H. Weier, S. Nauerth, A. Crespi, R. Osellame, and H. Weinfurter, "Design and evaluation of a handheld quantum key distribution sender module," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 21, pp. 131–137, May 2015.

[91] D. Tofsted, S. O'Brien, and G. Vaucher, "An atmospheric turbulence profile model for use in army wargaming applications I." Army Research Laboratory, February 2006.

[92] P. Marks, "Quantum cryptography to protect swiss election." New Scientist. `https://www.newscientist.com/article/dn12786-quantum-cryptography-to-protect-swiss-election/`.

[93] H.-K. Lo, M. Curty, and K. Tamaki, "Secure quantum key distribution," *Nature Photonics*, vol. 8, pp. 595 EP –, 07 2014.

[94] J. F. Dynes, W. W.-S. Tam, A. Plews, B. Fröhlich, A. W. Sharpe, M. Lucamarini, Z. Yuan, C. Radig, A. Straw, T. Edwards, and A. J. Shields, "Ultra-high bandwidth quantum secured data transmission," *Scientific Reports*, vol. 6, p. 35149, oct 2016.

[95] H.-K. Lo, M. Curty, and B. Qi, "Measurement-device-independent quantum key distribution," *Phys. Rev. Lett.*, vol. 108, p. 130503, Mar 2012.

[96] D. Vasylyev, W. Vogel, and A. A. Semenov, "Theory of atmospheric quantum channels based on the law of total probability," *Phys. Rev. A*, vol. 97, p. 063852, Jun 2018.

[97] P. Papanastasiou, C. Weedbrook, and S. Pirandola, "Continuous-variable quantum key distribution in uniform fast-fading channels," *Phys. Rev. A*, vol. 97, p. 032311, Mar 2018.

[98] A. Sit, F. Bouchard, R. Fickler, J. Gagnon-Bischoff, H. Larocque, K. Heshami, D. Elser, C. Peuntinger, K. Günthner, B. Heim, C. Marquardt, G. Leuchs, R. W. Boyd, and E. Karimi, "High-dimensional intracity quantum cryptography with structured photons," *Optica*, vol. 4, pp. 1006–1010, Sep 2017.

[99] M. Mirhosseini, O. S. Magaña-Loaiza, M. N. O'Sullivan, B. Rodenburg, M. Malik, M. P. J. Lavery, M. J. Padgett, D. J. Gauthier, and R. W. Boyd, "High-dimensional quantum cryptography with twisted light," *New Journal of Physics*, vol. 17, no. 3, p. 033033, 2015.

[100] P. Gregg, P. Kristensen, S. Golowich, and S. Ramachandran, "Demonstration of a Thin-Ring Air Core Fiber Supporting 22 Stable Angular Momentum Modes," *42nd European Conference on Optical Communication*, no. 1, pp. 653–655, 2016.

[101] L. Marrucci, "The q-plate and its future," *Journal of Nanophotonics*, vol. 7, pp. 7 – 7 – 5, 2013.

[102] J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, P. G. Kwiat, R. T. Thew, J. L. O'Brien, M. A. Nielsen, and A. G. White, "Ancilla-assisted quantum process tomography," *Phys. Rev. Lett.*, vol. 90, p. 193601, May 2003.

[103] S. Hill and W. K. Wootters, "Entanglement of a pair of quantum bits," *Phys. Rev. Lett.*, vol. 78, pp. 5022–5025, Jun 1997.

[104] I. L. Chuang and M. A. Nielsen, "Prescription for experimental determination of the dynamics of a quantum black box," *Journal of Modern Optics*, vol. 44, no. 11-12, pp. 2455–2467, 1997.

[105] I. Bongioanni, L. Sansoni, F. Sciarrino, G. Vallone, and P. Mataloni, "Experimental quantum process tomography of non-trace-preserving maps," *Phys. Rev. A*, vol. 82, p. 042307, Oct 2010.

[106] J. Chapman, C. Zeitler, H. Bernstein, K. Meier, and P. Kwiat, "Progress towards implementing superdense teleportation in Space," in *Advances in Photonics of Quantum Computing, Memory, and Communication XI* (Z. U. Hasan, P. R. Hemmer, A. L. Migdall, and A. E. Craig, eds.), p. 13, SPIE, feb 2018.

[107] J. D. Monnier, "Optical interferometry in astronomy," *Reports on Progress in Physics*, vol. 66, pp. 789–857, may 2003.

[108] D. Gottesman, T. Jennewein, and S. Croke, "Longer-baseline telescopes using quantum repeaters," *Physical Review Letters*, vol. 109, p. 070503, aug 2012.