ADVANCED QUANTUM COMMUNICATION SYSTEMS

BY

EVAN ROBERT JEFFREY

B.A., Washington University, 1999

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, 2007

Urbana, Illinois

Abstract

Quantum communication provides several examples of communication protocols which cannot be implemented securely using only classical communication. Currently, the most widely known of these is quantum cryptography, which allows secure key exchange between parties sharing a quantum channel subject to an eavesdropper. This thesis explores and extends the realm of quantum communication.

Two new quantum communication protocols are described. The first is a new form of quantum cryptography—*relativistic* quantum cryptography—which increases communication efficiency by exploiting a relativistic bound on the power of an eavesdropper, in addition to the usual quantum mechanical restrictions intrinsic to quantum cryptography. By doing so, we have observed over 170% improvement in communication efficiency over a similar protocol not utilizing relativity.

A second protocol, Quantum Orienteering, allows two cooperating parties to communicate a specific direction in space. This application shows the possibility of using joint measurements, or projections onto an entangled state, in order to extract the maximum useful information from quantum bits. For two-qubit communication, the maximal fidelity of communication using only separable operations is 73.6%, while joint measurements can improve the efficiency to 78.9%.

In addition to implementing these protocols, we have improved several resources for quantum communication and quantum computing. Specifically, we have developed improved sources of polarization-entangled photons, a low-loss quantum memory for polarization qubits, and a quantum random number generator. These tools may be applied to a wide variety of future quantum and classical information systems.

To Mom and Dad

Acknowledgments

This thesis is possible because of the help and support of countless people. A few I wish to thank specifically. Every teacher makes a mark on their students, but I want to particularly thank Mrs. Robinson for giving me a chance in 3rd grade, Pastor Gary Boe, Mr. Daddow, and Patrick and everyone else from the OPTAG program at ISU.

Teachers taught me many things, but my friends and family made me who I am. Thanks to my friends from high school, Dion, Dan, Matt, Mark, and Alicia, and those from college Bryan, Eric, Krystal, and Adam. Thanks to Jim Oliver and everyone I worked with at EAI, Matt, Bryan, and Ace. Thanks to my family for their relentless encouragement and support, Mom, Dad, Ryan, and Liz.

Thanks to everyone at UIUC who helped me through graduate school, and to learn more than I ever thought possible. Thanks to the instructors who helped prepare me for the qual, and thanks to Nigel Goldenfeld and the rest of qual committee, who forced me to learn things I would not have any other way. Thanks to my adviser Paul Kwiat and everyone else who worked in the Quantum Information group. It was an honor and a privilege to work will all of them. Also, thanks to my friends from graduate school, Joe, Evan, Micah, Dave, and Brian.

This work and my graduate career were financially supported by the ARDA/DTO-funded ARO research grant #DAAD19-03-1-0282, as well as a DCI Fellowship grant, and the University of Illinois.

Finally, thanks to my thesis committee: Paul Goldbart, Brian DeMarco, John Stack, and Paul Kwiat.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Quantum information is the study of communication and computation systems that utilize quantum resources such as entanglement (joint states of separated particles) and superposition states (states such as $|0\rangle + |1\rangle$ which have an amplitude to be in one of several states, and with a well defined phase between them). Since the beginning of (classical) information theory, it has been known that the number of steps to complete a computation depends on the model of a computer used. However, the strong form of the Church-Turing [3] thesis suggests that the classical Turing machine is the universal computer, able to compute any decidable function in the most efficient manner possible, up to polynomial time transformation. While there is no known machine that outperforms a Turing machine using only classical physics, computers that store their state in quantum bits (qubits) and perform gate operations using unitary operators may be able to efficiently solve problems that Turing machines require exponentially more time to solve. For instance, Shor's algorithm [4] allows a quantum computer to factor large numbers in polynomial time, while the fastest known classical algorithms are super-polynomial [5]; however, the true classical complexity of factoring is not known. Grover's algorithm [6] allows a quantum computer to search an unordered list in $O(\sqrt{n})$ operations, asymptotically faster than any classical computer (which requires $O(n)$ operations).

While quantum computation allows some functions to be computed faster than is possible with classical computers, the class of "decidable" functions is the same. Certain problems such as the halting problem (determining whether or not a given problem will ever finish executing) are known to be unsolvable by classical computers, quantum computers are also unable to solve these problems. In principle, given enough time and memory, a classical computer can simulate a quantum computer, allowing it to evaluate any quantum algorithm. Quantum communication, the branch of quantum information involving more than one party, actually allows protocols that cannot be implemented using purely classical resources. The best known protocol of this type is quantum key distribution (QKD). This protocol allows two parties, Alice and Bob, to generate a shared secret key — a random string of bits on which they perfectly agree, yet which an eavesdropper has essentially zero knowledge of. Classical, this can only be done with a physically secure communication medium, usually regarded as requiring Alice and Bob to meet in person to exchange keys.

QKD allows Alice and Bob to make this exchange over a quantum channel, even if it is controlled by the eavesdropper.

QKD is currently the most widely known and implemented quantum communication protocol; however, there are other classical communication protocols that may be implemented or improved upon by allowing quantum resources, including bit-string commitment [7], the Byzantine General's problem [8], fingerprinting [9, 10], and others. Beyond these fundamentally classical protocols, there is another type of quantum communication, that is, using protocols that directly operate on quantum information. In particular, quantum teleportation [11] and quantum repeaters [12, 13] may become important in the future for connecting quantum computers.

Here we discuss the implementation of novel two quantum communication protocols. The first is relativistic quantum cryptography (Chapter 5). This protocol uses constraints from relativity in order to improve the efficiency of communication by excluding eavesdropping strategies that violate causality. Ordinarily, Alice and Bob must randomly select a communication basis for each photon transmitted. The selection process must be random, since it must be impossible for an eavesdropper to predict which basis they will choose. Unfortunately, the cases where Alice and Bob do not randomly select the same basis must be discarded, in a process known as sifting. In our relativistic protocol, Bob does not have to select his basis until after Eve can no longer access the signal, allowing him to eliminate the sifting step and improving the total communication efficiency.

The second quantum communication protocol is quantum orienteering (Chapter 6). This protocol allows the communication of a direction in space without reference to an external reference frame, a process impossible with "pure" information. Instead, the protocol relies on the specific physical embodiment of the information. In particular, we show that when a direction is encoded into multiple spin-1/2 particles, or in our case, two photons treated as effective spin-1/2 particles, the measurement which gives the most information is a joint measurement. Such measurements project onto an entangled state, rather than independently measuring each copy of the state. Perhaps more surprising is that specific encoding matters: two parallel spins provide less information than a pair of anti-parallel spins, though classically these two encodings are equivalent.

We used several tools developed to allow these and other quantum information protocols to be implemented, including a bright, pure source of polarization entangled photons (Chapter 2) used by both protocols. This source uses a pair of type-I non-linear crystals. When pumped with an appropriate UV laser, one crystal generates pairs of horizontally polarized photons ($|HH\rangle$), while the other generates pairs of vertically polarized photons ($|VV\rangle$). If the light from the these two crystals is combined such that the processes are

quantum mechanically indistinguishable (i.e., it is not possible to tell, even in principle, from which crystal a photon is emitted), then the resulting state is an entangled state of the form $\cos(\theta)|HH\rangle + e^{i\phi}\sin(\theta)|VV\rangle$, which can be converted into many desired entangled states.

In Chapter 3 we describe a low-loss quantum storage system based on an optical delay. We use this delay line to implement relativistic quantum cryptography. The delay is created by reflecting photons between a pair of mirrors many times before the photon exits a small hole in one of the mirrors. Our design uses cylindrical mirrors in order to generate very long delays (up to 7 μs in a 2-meter long cavity). While this sort of delay lacks some features of potential systems that transfer an optical state onto an atomic system and back, such as continuously controllable delay and very long storage times, our delay line technique also has considerable advantages. In addition to its relative simplicity, it is broad band, polarization preserving, and can store many photons at once in different time slots or spatial modes.

Finally, we have developed a high-speed random number generator (Chapter 4) suitable for use in quantum cryptography applications. We generate random bits using the inherent unpredictability of photon detection as a source of randomness by measuring the arrival time to several digits of precision. The major advance with our design is that we can generate multiple random bits per detected photon, in contrast to previous designs which generate at most a single random bit per photon. The disadvantage to this approach is that the output shows bias – some values are more likely than others; however, we eliminate this with post-processing by using hash functions, functions that map an arbitrary length input string to a fixed length output string, to compress the data to an essentially perfectly random sequence. We can generate random bits after post-processing as fast as 20 Mbit/s out of a single photon-counting detector running at 4 MHz average count rate. The best available true random number generators generate 12 Mbit/s using a high-speed photo-multiplier tube.

These tools help us to implement multiple quantum communication protocols, and will likely have application to future quantum information applications as well. In particular, quantum repeaters, an essential component of long-range quantum communication, require entangled photons and quantum storage, while our high-speed quantum random number generator is suitable for many cryptographic protocols and other applications which consume random numbers at a high rate.

Chapter 2

Entangled Photon Source

2.1 Introduction

Perhaps the most important resource for quantum communication is the production of entangled photons. While massive qubits such as quantum dots or trapped ions may in the end be the medium of choice for quantum computing, quantum communication will probably always rely on optical qubits, which can be easily transmitted over long distances with minimal interaction with the environment. Optical qubits are useful not only for long distance communication, but also possibly for interconnect busses within a quantum computer. Photons have several accessible degrees of freedom; however, we use polarization qubits for our quantum communication protocols. Photon polarization (spin) forms an effective 2-level system (it is a massless spin-1 particle, therefore the $m_z = 0$ state is excluded), and it is usual to identify the computational basis states as horizontally polarized ($|H\rangle$) and vertically polarized ($|V\rangle$), relative to an arbitrary standard. Several of the most useful states are show in Table 2.1. Multi-qubit states can be represented as tensor products of these sing-qubit states. For instance, two photons, both horizontally polarized can be written as $|HH\rangle \equiv |H\rangle \otimes |H\rangle$, and an entangled Bell state as $|\phi^+\rangle \equiv (|HH\rangle + |VV\rangle)/\sqrt{2}$.

Currently, the workhorse entanglement source is spontaneous parametric downconversion (SPDC). In this process, a UV pump beam acts on a material having a second-order nonlinearity, such as β-barium borate (BBO). With some small ($\sim 10^{-10}$) probability, individual photons from the pump beam will down-

Polarization	Notation	Equiv. Representation
Horizontal (0°)	$\|H\rangle$	—
Vertical (90°)	$\|V\rangle$	—
Diagonal (45°)	$\|D\rangle$	$\|H\rangle + \|V\rangle)/\sqrt{2}$
Anti-diagonal (−45°)	$\|A\rangle$	$\|H\rangle - \|V\rangle)/\sqrt{2}$
Right-circular	$\|R\rangle$	$\|H\rangle + i\|V\rangle)/\sqrt{2}$
Left-circular	$\|L\rangle$	$\|H\rangle - i\|V\rangle)/\sqrt{2}$

Table 2.1: All polarization states can be written as a superposition of the computational basis states $|H\rangle$ and $|V\rangle$.

convert into pairs of lower energy photons. Due to the requirement that the process conserve energy and momentum (also known as phase matching, the requirement that amplitudes for downconversion through the crystal add constructively in phase), these pairs of photons are potentially entangled in every possible degree of freedom [14]. This process is particularly useful compared with older entangled photon sources such as atomic cascades [15, 16], because the phase matching conditions allow the source to be tuned in frequency and direction by varying the orientation of the non-linear crystal and the pump beam. Thus, the process can be tuned for a specific application. In later chapters, we will show how this source is used to implement two new quantum information protocols, Relativistic Quantum Cryptography (Chapter 5) and Quantum Orienteering (Chapter 6).

2.2 Downconversion

Downconversion is a form of three-wave mixing in which pairs of low-frequency photons are created from individual high-frequency photons. The high-frequency mode is referred to as the *pump*, usually provided by a UV laser, while the two lower-frequency modes are called the *signal* and the *idler*. This process requires a material with a non-linear optical response. In such materials, the electric displacement vector is not simply proportional to the electric field, but also depends on higher-order terms such as $E_i E_j$. These higher-order terms allow light from different modes (frequency, momentum, and polarization) to interact. Classically, three-wave mixing requires at least 2 of the modes to be pumped externally. However, quantum mechanics provides a mechanism that allows the process to occur spontaneously, essentially amounting to stimulation from the vacuum field in a manner analogous to spontaneous decay in atomic systems.

Non-linear optical processes are classified by their input and output polarizations. Non-linear crystals must lack inversion symmetry in order to posses a 2nd-order (χ_2) non-linearity, and are birefringent (they have different indices of refraction for different polarization states). In order to generate downconversion from a given system, the specific non-linear coefficient must be non-zero and the phase matching constraints discussed below must be satisfied. We use Type-I downconversion, characterized as an *eoo* process. This notation indicates the pump is extraordinarily polarized compared to the optic axis of the downconversion material, while both the signal and idler have ordinary polarization.

2.3 Phase Matching

Each portion of the downconversion crystal illuminated by the pump laser has an amplitude to emit down-converted photons. However, for most possible wavelength and direction combinations, contributions from

each part of the crystal cancel. Only when the process (approximately) conserves energy and momentum within the crystal will the amplitudes combine constructively to give a large signal. This is known as the phase matching condition, and the maximum phase mismatch depends on the length of the crystals and the bandwidth of the pump. For Type-I phase-matching, the down-converted photons form a cone centered on the pump beam. Conditional on a photon being detected on a specific part of that cone, the other is known to be on the opposite side of the cone. The cone's apex can be adjusted by tilting the crystal's optic axis. This changes the extraordinary index of refraction seen by the pump beam, and therefore the momentum-conservation condition.

2.4 Entanglement Generation

Photon pairs emitted from Type-I downconversion are automatically entangled in the continuous variables of frequency and momentum, as a straightforward consequence of the conservation laws. However, the polarization output from a single crystal is a pure product state: both photons have ordinary polarization. In order to generate polarization entangled photons, we use two abutting crystals: one with its optic axis in the vertical plane, which generates horizontally polarized photon pairs from vertical pump photons, and the other with its optic axis in the horizontal plane, which generates vertically polarized photon pairs from horizontal pump photons. We pump both crystals with a diagonally polarized pump laser. If the crystals are sufficiently thin and close together, then due to diffraction it is impossible to tell *even in principle* which crystal a given photon pair came from. In this case, to find the output polarization state, the amplitudes for emission from the first and second crystals are added coherently, to give an entangled output state:

$$|\psi\rangle = \cos(\theta)|HH\rangle + e^{i\phi}\sin(\theta)|VV\rangle. \tag{2.1}$$

The parameter θ depends on the relative amplitude for downconversion from the two crystals, and can be adjusted by rotating the pump polarization, and ϕ is the phase between the two processes, dependent on the exact thickness of the crystals and the birefringent phase accumulated as the light propagates through them. If necessary, the phase can be adjusted by adding a variable waveplate to one of the output beams. However, some applications only require a (nearly) maximally-entangled state for which the phase between the $|HH\rangle$ and $|VV\rangle$ terms does not matter. A schematic depiction of this sort of entanglement source is shown in Fig. 2.1.

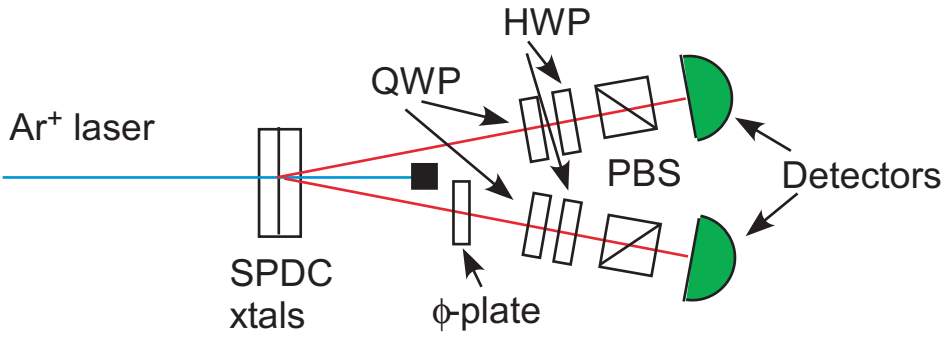

Figure 2.1: Figure showing an example entangled photon source using parametric downconversion. A vertically polarized pump laser (351 nm Ar^+ line) is rotated to diagonal polarization by a half waveplate HWP. It then traverses two 0.6-mm BBO crystals, with their optic axes in orthogonal planes. The crystals are cut such that the downconversion emerges in a cone with a 3° opening angle. An optional "ϕ plate" in one arm is a quarter waveplate (QWP) which is tilted to adjust the phase of the output entangled state. A QWP, HWP and polarizer in each arm allow state tomography to characterize the emitted entangled state.

2.5 Compensation

In the two-crystal Type-I downconversion system described above, the state emitted by the crystals varies slightly with direction. The primary difference is in the phase ϕ in the state $|HH\rangle + e^{i\phi}|VV\rangle$. The photons emitted in the first crystal travel through the second crystal with extraordinary polarization, and acquire a phase associated with the path length through the crystal and its index of refraction, which in turn depends on the direction of propagation due to the birefringent nature of the crystal. When the collection irises subtend a large solid angle, this phase variation across the collection aperture effectively causes decoherence, degrading the quality of the state. In the BBO source described above, we typically use irises with a diameter of 5 mm, at a distance of 1200 mm from the crystals (14 arcminutes). In this arrangement, the phase of the emitted entangled state varies by about 70 degrees across the iris, causing significant loss of state quality. In order to improve the quality of the state, we add an extra birefringent element into each arm of the downconversion, thereby canceling this phase variation over the width of the iris to first order [1]. The result is a considerably brighter, more pure source of polarization entangled photons than was previously available.

2.6 Fiber Coupling

For some applications, including our quantum key distribution (QKD) experiment described in chapter 5, it is more desirable to have entangled photons in fiber optics, at least for part of the transmission channel. The basic operation of the source is the same; however, the signal and idler have different wavelengths, one at 670 nm (to match laser diodes also used in the experiment) and the other at 737 nm (required by energy

conservation). We also used a different non-linear crystal, BiB_3O_6 (BiBO), which has a higher non-linear interaction coefficient, making it a brighter source. In order to preserve entanglement through the fiber optics, we must use single-mode fibers, otherwise the phase shift between different modes in a multi-mode fiber will degrade entanglement. A schematic of our fiber-coupled entanglement source is shown in Fig. 2.2. In order to get a high coupling efficiency for the fiber-coupled source, we must focus the pump to a small spot in the downconversion crystals [17]. Since this source is non-degenerate — the emitted downconversion photons are at different wavelengths — we can use collinear downconversion. In that configuration, the "cone" collapses to a beam, with both photons emitted in the same direction as the pump. Beam-like modes are easier to couple into fiber, but first we must use a laser cutoff filter to block the pump, and a dichroic mirror to separate the two colors of downconversion. Each color is then collected into a single-mode fiber using a microscope objective. We can perform a quantum state tomography on the output from the fiber optics to verify that entanglement is preserved. However, bends in fiber optics do induce polarization rotations due to geometric effect (e.g., Berry's phase [18]) and stress-induced birefringence. Thus, the output state will not be the clean $|HH\rangle + |VV\rangle$ state emitted by the crystals, but one which has undergone some local-unitary rotation on each polarization qubit. This rotation can be undone using waveplates if necessary; however, a more difficult problem is that the rotation, and therefore the output state, will drift as the fiber moves and the temperature changes. Such drift is a serious problem for our applications, so we reduced it by taping the fiber to our optical tables to keep them held in place and connected to a thermal reservoir. The resulting entangled state can be stable for days at a time; we simply perform a state tomography at the beginning of data collection, and the drift throughout the course of an experiment is minimal.

It is also necessary to use fibers without a large protective jacket, as the fiber is usually able to twist inside the jacket. Over time, the fiber will gradually relax within the jacket after any movement, requiring hours or days to reach a stable configuration. Fiber patch cords are commonly available with 3-mm and 0.9-mm jackets, and we have observed that the latter seem to work considerably better for entanglement distribution.

2.7 Tomography

To show that the source of entangled photon pairs is indeed generating entanglement, it is sufficient to measure a violation of Bell's inequality [19, 20]. However, for most applications of entanglement it is important to know the exact state produced. To do this, we use a process known as quantum state tomography [2, 21]. Tomography is the mathematical reconstruction of the complete density matrix from a series of projective

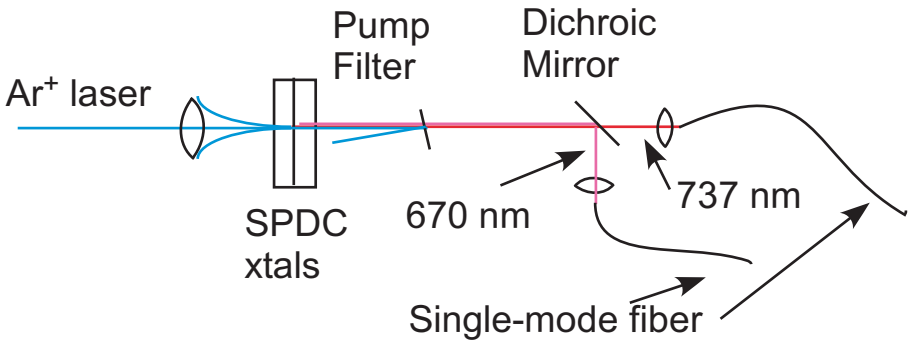

Figure 2.2: Fiber-coupled source of polarization entangled photons. Here, the downconversion is collinear with the pump, and the signal and idler are separated by wavelength using a dichroic mirror before coupling into single-mode fibers. To achieve high coupling efficiency, we focus the pump beam to a small spot in the crystals.

measurements, each of which provides partial information about the quantum state.

The two-qubit state of photons emitted by our SPDC source is most completely represented by a 4×4 density matrix. This form can not only represent standard pure quantum states, but also mixture caused by e.g., entanglement to other degrees of freedom. The expectation value for a specific projection ($|\psi\rangle$), given the density matrix representation (ρ) of a state is given by:

$$P(\psi) = \langle\psi|\rho|\psi\rangle. \tag{2.2}$$

Since the 4×4 density matrix is Hermitian and has unit trace, there are 15 free parameters, so we need to perform at least 15 different measurements. In practice, we require 16 unique measurements in order to normalize the intensity of our source. The usual choice for a 16-measurement state tomography is to measure every combination of $|H\rangle$, $|C\rangle$, $|D\rangle$, and $|R\rangle$ on each qubit, for a total of 16 measurements. These data are then used to predict a density matrix that would generate those outcomes. In practice, the results are subject to statistical fluctuations of a Poissonian nature, and may not always exactly reflect a legal (i.e., positive semi-definite) density matrix. Therefore, we use a maximum likelihood optimization [2] to find the best fit density matrix.

While it is most obvious to use the 16 measurements listed above to do state tomography, we have shown that it is not optimal in the sense of minimizing the uncertainty with a given number of state copies. We do considerably better by taking 36 separate measurements, the combinations of projections onto $|H\rangle$, $|V\rangle$, $|D\rangle$, $|A\rangle$, $|R\rangle$, $|L\rangle$ for each photon [21]. All 36 measurements are then used in the maximum likelihood optimization. The reason for this advantage is that the density matrix is a representation of a probability of a given outcome. To get a probability from the data we take, each measurement, consisting of a number of

coincident detections, must be divided by the total rate of production. In the 16-measurement tomography, the intensity is given by the sum of four measurements: $HH + HV + VH + VV$. Thus, those 4 measurements of the 16 strongly affect every element of the density matrix, while the other 12 only affect a single parameter. By using the more balanced set of 36 measurements, each measurement contributes to our knowledge of the intensity, and thus decreases the uncertainty of the reconstructed density matrix.

Chapter 3

Optical Storage Cavity

Optical qubits have excellent properties for quantum communication, since they travel at the speed of light and are essentially impervious to electric and magnetic fields. Unfortunately, this makes them hard to store for processing, and they are susceptible to absorption in normal media. For many quantum computation applications, it is desirable to convert the "flying" optical qubit into a stationary (e.g., atomic) qubit. For some applications, however, only a brief storage time is necessary. This is the case for relativistic quantum cryptography, as well as some implementations of quantum repeaters and feed-forward control in optical quantum computing. In these applications, it is more convenient to store the photon directly rather than converting it to a stationary qubit and back. The simplest approach to store a photon for a brief period is to reflect it back and forth between mirrors in a cavity configuration. Two meters is a reasonable size for a tabletop cavity, giving a base delay time of $t_0 = 6$ ns per reflection. Even the finest mirrors will have some loss on each reflection, with the total storage efficiency for time T given by $\eta \approx R^{T/t_0}$ where R is the reflectivity of the mirror. Assuming an arbitrary ruler of useful storage as 50% efficient and a commercially available mirror with 99.99% reflectivity, we can tolerate nearly 7000 reflections yielding greater than 40 μs of storage time, a virtual eternity for modern computer systems. This is equivalent to 0.25 dB/km in vacuum, or 0.4 dB/km through glass, which is competitive with the best single-mode optical fibers, and considerably better than fibers can achieve for visible light.

There are a number of approaches to folding an optical path into a small volume. Conceptually the simplest approach is to build a single loop with a switching element, such as a Pockels cell. Figure 3.1 shows a schematic of this type of storage, in which a photon or pulse can be stored for an integer number of loop times, then switched out by an electrical signal. This approach is not useful for long storage times, however, since the loss in even a very good switch will be $\sim 1\%$. This particular implementation also suffers from the drawback that the polarization degree of freedom is used to control the switching behavior, meaning only a single polarization can be stored. More complicated switches, such as interferometric switches, allow storing arbitrary polarizations, typically with higher loss. Nevertheless, this simple delay can be combined with the approaches below to create a low-loss storage loop that can reach very large, adjustable delays.

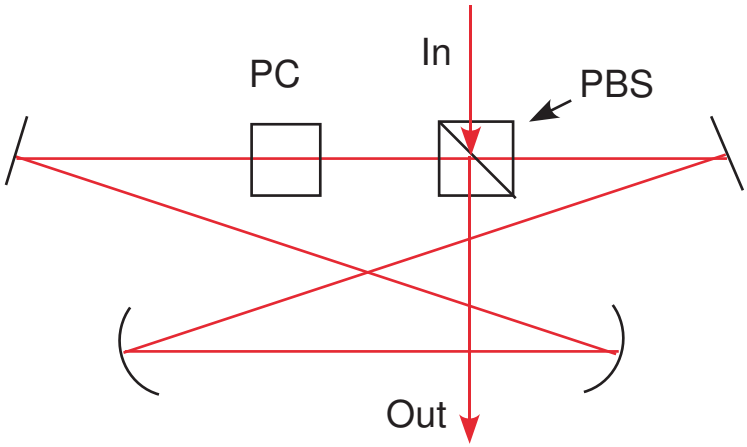

Figure 3.1: A simple switchable delay line. Photons are switched into and out of the cavity loop by means of a Pockels cell that rotates the polarization between horizontal polarization, which is transmitted by the polarizing beam-splitter (PBS), and vertical polarization, which is reflected out of the cavity by the PBS.

A better approach to optical storage is to use a pair of spherical mirrors in a cavity configuration as shown in Fig. 3.2. A small hole is drilled in one of mirrors, and they are arranged so that light entering the hole can bounce back and forth many times before exiting the same coupling hole. Once inside the cavity, the light only bounces off of mirrors, which can be made highly reflective, and avoids transmissive elements, which typically have higher loss. The curvature of the mirror surfaces acts both to keep the beam confined within the cavity volume and to periodically focus the beam to keep diffraction losses from becoming too large. In order to have high efficiency, some parameters must be optimized. First, the coupling hole must be large enough to pass the incident beam without noticeable clipping. The "πw_0" criteria [22] states that a hard aperture must be approximately three times the full-width $1/e^2$ of a Gaussian beam in order to pass 99% of the beam (a 1% loss can be tolerated here as it is only encountered twice). While the input beam can be made small by focusing, it will change due to diffraction and periodic focusing, with the optimal configuration usually being an input beam with a waist near the center of the cavity and a Rayleigh range comparable to the length of the cavity. Second, the nearest intermediate reflection must be far enough from the coupling hole that a negligible amount of light leaks out of the cavity early. Thus, as the number of passes are increased, the size of the mirror must be increased to allow the pattern of spots to be spread out enough to avoid leakage losses. Below we will see another approach to deal with this problem.

Optical cavities such as this are analyzed using ray matrices [22] (or ABCD matrices) which describe the evolution of a paraxial light beam as it traverses the elements of an optical system. Rotationally symmetric

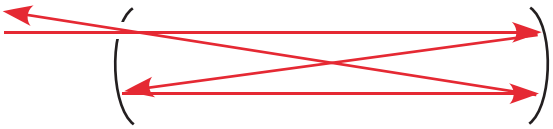

Figure 3.2: A multi-pass delay line made from spherical mirrors. A photon entering through the hole in one mirror will bounce several times before exiting through the same hole, tracing out the spot pattern shown on the right. This delay can have significantly lower loss due to the absence of a switch.

systems have simple behavior, and the action of an optical element is represented as

$$\begin{pmatrix} x' \\ \frac{dx}{dz}' \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \cdot \begin{pmatrix} x \\ \frac{dx}{dz} \end{pmatrix} \tag{3.1}$$

with the same ABCD matrix governing evolution in both the $\hat{x}$ and $\hat{y}$ directions. A ray is described by two parameters, x, the the displacement from the central ray, and dx/dz, the slope of the ray in the $\hat{x}$ direction. The change in those parameters due to a specific element is given by multiplying by an ABCD matrix. For a rotationally symmetric optical system, the evolution in the $\hat{x}$ and $\hat{y}$ directions is independent and described by identical equations (though possibly with different initial values for y and dy/dz).

The two important optical elements, a curved mirror of radius R and a free path of length d are given by:

$$M(R) = \begin{pmatrix} 1 & 0 \\ -2/R & 1 \end{pmatrix}, \quad L(d) = \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix}. \tag{3.2}$$

Since this evolution is linear, ray matrices from multiple elements are simply multiplied to obtain the ray matrix for the entire system. By multiplying all the elements encountered in the cavity (two mirror reflections and two free-space propagators), we can obtain a single ABCD matrix that describes how the beam evolves through each round trip in the cavity. If after N cycles through the whole cavity the beam returns to the input hole, the cavity is reentrant with period N, and light entering the cavity will be delayed by $t = 2Nd/c$.

For moderate delays, it is typical to use spherical mirrors for this purpose. Spherical mirrors are easy to fabricate, and can be easily coated with very high reflectivity coatings ($R > 99.99\%$ high-reflectivity coatings are available). The ABCD matrix for this cavity is given by:

$$A = L(d) \cdot M(R) \cdot L(d) \cdot M(R) = \begin{pmatrix} 1 & d \\ -2/R & 1 - \frac{2d}{R} \end{pmatrix}^2, \tag{3.3}$$

and the asymptotic behavior of a beam reflected many times may be understood by examining the eigenvalues

of A. The solutions to the characteristic polynomial can be written as

$$\lambda = \left(1 - x \pm \sqrt{x^2 - 2x}\right)^2, \tag{3.4}$$

where the only parameter is a dimensionless quantity $x = d/R$, the separation of the mirrors divided by their radius of curvature. If $\lambda > 2$, the eigenvalues will be real and one of them will have magnitude greater than 1. In this case, a beam passing into the cavity will eventually escape, as the eigenvalue of the iterated ray matrix A^N becomes exponentially large. On the other hand, if $\lambda < 2$, the solutions for complex conjugate pairs of the form $e^{\pm i\phi}$. In this case, it is possible for the evolution of a ray through the cavity to be periodic, returning to its initial condition after an integer number of cycles. This is the desired behavior for our storage cavity. The cavity will be periodic if the ray matrix satisfies

$$\begin{aligned} A^N &= I, \\ e^{iN\phi} &= 1, \\ N\phi &= 2\pi m. \end{aligned} \tag{3.5}$$

These solutions correspond to the beam spots lying on a circular or elliptical path on the surface of each mirror. The maximum delay is given with a circular pattern and $N_{max} = 2\pi r/\delta$, where δ is the minimum spot spacing and r is the radius of the mirror face. Thus, the maximum delay time scales only linearly with the radius of the mirror used. A better solution would use the entire area of the mirror to achieve longer storage in the same volume; in this case the number of reflections scales are the *square* of the mirror size.

To use the entire mirror area, it necessary to break the $x - y$ symmetry. Then, the beam propagation in the x and y directions will be represented by different ABCD matrices having different eigenvalues $e^{\pm i\phi_1}$ and $e^{\pm i\phi_2}$, and the entire matrix will only be cyclic when both

$$N\phi_1 = \pi m_x \quad \text{and} \quad N\phi_2 = \pi m_y. \tag{3.6}$$

The reason for only requiring a multiple of π rather than 2π is that for these systems we usually place the coupling hole in the center ($x = y = 0$) rather than the edge, so the total ray matrix may have either $+1$ or -1 on the diagonals. For these configurations, the beam will trace out a Lissajous pattern on the mirror faces, covering much more of the area before exiting, and allowing much longer delay times.

It is much more difficult to construct mirrors having different radii of curvature in two directions, since they are not surfaces of rotation. In addition, the ratio of R_x and R_y is a critical alignment parameter, and

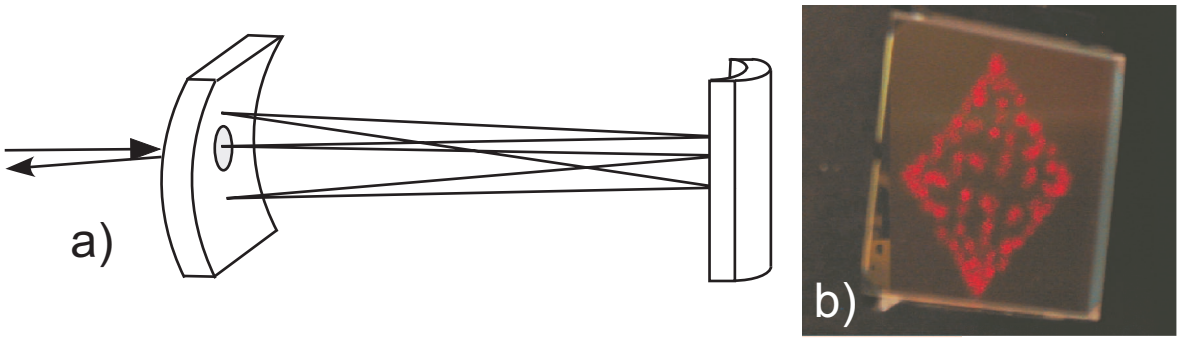

Figure 3.3: A multi-pass delay line using cylindrical mirrors to achieve very long delay times. The number of bounces before exiting the coupling hole is adjusted by varying both the mirror separation *and* the relative twist angle. The beam traces out a Lissajous pattern on each mirror, shown to the right, allowing most of the surface of the mirror to be used, rather than only the perimeter.

manufacturing tolerances are usually not good enough to get reliable results. One common approach is to bend spherical mirrors into the desired astigmatic shape. This works acceptably for metal mirrors used in infrared applications, but is not practical for glass mirrors with dielectric mirror coatings. Furthermore, this generates systems with a relatively small astigmatism, which leads to many near approaches to the coupling hole. More desirable configurations require greater differences in curvature. A second approach is to use mirrors with an x vs. y curvature difference greater than desired, then twist one of the mirrors slightly to compensate. This allows the more desirable configurations, and solves the manufacturing tolerance problems; however, it does not address the difficulties in creating astigmatic mirrors of sufficient quality. We designed a more effective system using a twisted cylindrical mirror cavity (Fig. 3.3), which allows the entire range of cavity configurations and uses mirrors that are somewhat easier to fabricate. If a cavity of this type is built with two cylindrical mirrors whose axes of curvature are aligned in orthogonal directions, we obtain a configuration very similar to that with spherical mirrors, since the cavity looks similar (it has the same total focusing power) in the x and y directions. As one mirror is twisted around the cavity axis, this symmetry is broken and the eigenvalues split, leaving a system very similar to that with astigmatic mirrors.

Since this twisted cylindrical cavity not only has different behavior in the x and y directions, but coupling between them, it cannot be described by either a single 2x2 ABCD matrix, or separate matrices for each direction. Instead, we must use a 4x4 matrix which can represent an arbitrary collinear system. The basic ray matrices we need to describe our cylindrical cavity are: a propagator $L(d)$, which gives the evolution for propagation through free space; a mirror $M(r_x, r_y)$, for reflection off a mirror with radii of curvature r_x and r_y; and a rotator, $R(\theta)$ which gives a coordinate transform used to allow orientation of mirrors with

non-orthogonal axes. The formulas for these matrices are:

$$L(d) = \begin{pmatrix} 1 & d & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & d \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

$$M(r_x, r_y) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -2/r_x & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -2/r_y & 1 \end{pmatrix}, \tag{3.7}$$

$$R(\theta) = \begin{pmatrix} \cos(\theta) & 0 & \sin(\theta) & 0 \\ 0 & \cos(\theta) & 0 & \sin(\theta) \\ -\sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & -\sin(\theta) & 0 & \cos(\theta) \end{pmatrix},$$

giving a total ABCD matrix equation of

$$\begin{pmatrix} x' \\ \frac{dx}{dz}' \\ y' \\ \frac{dy}{dz}' \end{pmatrix} = L(D) \cdot R(\theta) \cdot M(0, r) \cdot R(-\theta) \cdot L(D) \cdot M(r, 0) \begin{pmatrix} x \\ \frac{dx}{dz} \\ y \\ \frac{dy}{dz} \end{pmatrix}. \tag{3.8}$$

We used numerical methods to solve the eigenvalue equation. Again, since this is a real matrix, the eigenvalues for stable solutions come in complex conjugate pairs, and we search for solutions with θ_1 and θ_2 satisfying Eqn. (3.6). A plot of all solutions for $40 \leq N \leq 120$ is shown in Fig. 3.4.

It is simple to perform a numerical search for all solutions of a given N, then simply select the one that has the best properties. The primary consideration when determining the best configuration is to look for the nearest reflection to the coupling hole to be as far as possible from it. This not only prevents loss from either leakage through the coupling hole or from imperfections in the optical coating near the hole, but is typically less sensitive to alignment. If the nearest approach is relatively far, small perturbations in the mirror alignment will not move to a lower-pass configuration.

We use cavities of this sort to implement relativistic quantum cryptography. However, they may have other applications in quantum information, including feed-forward error correction [23] in optical quantum computing and quantum repeaters [12, 13] for entanglement distribution, which require switchable storage.

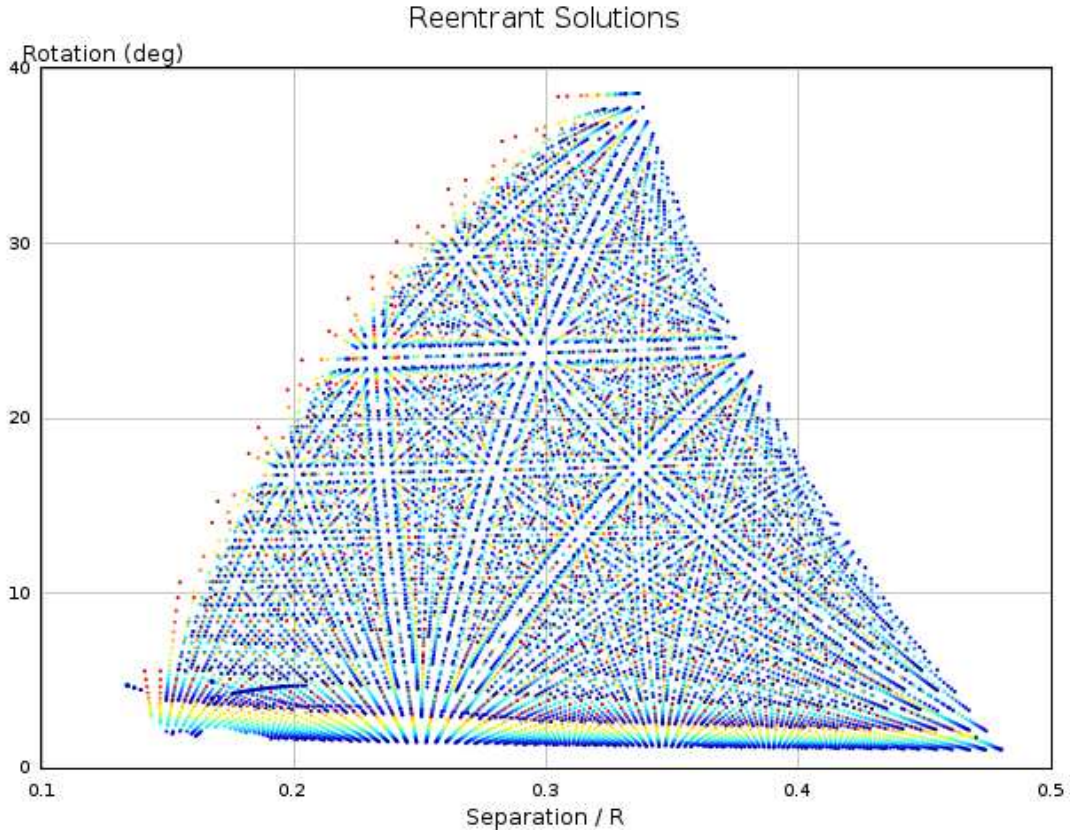

Figure 3.4: A plot of stable, periodic cavity solutions. The adjustable parameters are the relative twist of the second mirror and the ratio of the separation of the mirrors (in units of the radius of curvature). If the two mirrors have slightly different curvatures (e.g., due to manufacturing variation), most of the solutions will still work with slight adjustment, although the degenerate case, where $\theta = 0$ and $m_x = m_y$, is no longer accessible.

This system has many advantages over more complicated systems. It is inherently broad-band, since high reflectivity (HR) mirror coatings typically operate well over tens of nanometers, i.e., very broad compared to atomic line widths. Furthermore, the center wavelength can easily be tuned to match any desired center wavelength; for instance, a cavity could be constructed to match the entire 1.5 μm telecom band, or whatever atomic lines are needed for other parts of a communication system. Moreover, when properly aligned, the cavity is reentrant for any input k-vector that does not immediately escape confinement, giving an almost equal delay. Thus, a single cavity can store many photons multiplexed in space and frequency, and pipelined in time.

High quality mirrors with reflectivity greater than 99.999% are available, and are used by projects such as LIGO [24] and high finesse cavities for atomic quantum information [25]. However, for our cavity the total storage efficiency as a function of time delay shows a mirror efficiency of only 99.85%. While the difference is inconsequential for a single reflection, the last 0.1% makes an enormous difference with the large number of reflections in long storage-time cavities. The primary problem appears to be surface scatter. There are several possible causes of this, the first of which is buildup of contaminants on the surface of the mirrors. This is evidently part of the problem, as over time the reflectivity dropped to 99.7%. Cleaning the surface with standard solvents did not improve the performance; however using OptiClean [26], a polymer product, which is applied to the surface in liquid form, then peeled off once it solidifies, improved the performance back to 99.8%, nearly its initial value. This suggests that there is a more fundamental problem with the surface quality of the mirror substrate. Further investigations are warranted.

Despite the difficulties in obtaining mirrors with the required efficiency, it is possible to create relatively long photon storage using this technique, and the delay-time is adjustable over a wide range. If the difficulty in manufacturing cylindrical mirrors of the required reflectivity can be overcome, this could allow for delay times exceeding 10 μs, which can be useful for a number of quantum information protocols.

State-of-the-art mirrors can have reflectivity as high as 99.9996% [25]. If we could obtain cylindrical mirrors with this reflectivity, reflection loss would no longer be limiting the storage times, but the density to which spots can be packed onto the mirror faces would be. For instance, at this level we could tolerate 12000 reflections with only 5% loss and 80 μs storage time. However, that would place 6000 reflections on each mirror with an average spacing of only 1 mm, which is too close to be useful. If the mirror size were increased to allow this many spots, it is likely that several other factors would come into play. Contamination of the surface and scattering from air can be solved by placing the cavity in a vacuum. However, other problems are harder to solve. In particular, the difficulty of aligning for a particular cavity solution increases with the number of spots. Alignment and vibration sensitivity would likely be the limiting factors when trying

to achieve such a long delay. A better way to take advantage of such an efficient storage cell is a hybrid switched system, using a long storage time that periodically passes through a switch. We have constructed cells with delay times of up to 5 μs, so that is within the range of achievable alignment, although probably near the limit, given the particular mirrors we use. If we had mirrors with such a high efficiency, a single trip through that cavity would have a loss of about 0.3%. Coupling this to an optical switch with a 1% loss would allow delays switched in 5 μs increments with quite high efficiency. Another way to exploit high reflectivity mirrors would be to use multiple k-vectors. If the exiting beam is simply redirected back to the cavity with a slightly different k-vector, the delay time can be extended without using switches. This might allow long delay times while completely avoiding the losses associated with switches.

Chapter 4

Quantum Random Number Generator

4.1 Introduction

Many applications require random input. Sources of random numbers can be broadly divided into two classes. The first of these is the pseudo-random number generators (PRNGs). These sources maintain an internal state, and from that state use algorithmic means to generate a series of numbers that behave for some application the same as random numbers. The initial value of the internal state is called a seed, and for the same seed a given generator will always produce the same pseudo-random sequence. This is a particularly desirable feature for some applications, such as simulation of stochastic processes, since a simulation can easily be rerun with the same input to verify its operation. A simple and very common PRNG used in non-cryptographic applications is the linear congruential generator, defined by the sequence

$$\begin{aligned} x_0 &= \text{seed}, \\ x_{i+1} &= m \cdot x_i + b \pmod{2^N}, \end{aligned} \tag{4.1}$$

which with a carefully chosen m and b will cycle through all of the N-bit values in a random-looking fashion. Other PRNGs are based on encryption algorithms, making it hard to deduce from the output what the internal state is, and therefore make it difficult for an attacker to predict the output.

However, by their very nature pseudo-random generators are predictable and contain patterns. For instance, the linear congruential generator above has a well-known defect when used for many numerical simulations. successive random numbers are organized into k-tuples, those points will lie on at most $2^{\frac{N}{k}}$ hyperplanes [27]. Therefore, for other applications, especially cryptographic applications where a malicious agent may try to guess the output of the generator, it is desirable to use the second class, the true random number generators. True random number generators do not have internal state, but instead sample some inherently random physical process to generate their output.

True random numbers are particularly desirable for security applications, as the inherent predictability

in PRNGs can be exploited to circumvent the security protocol. For instance, in secure web transactions a public-key encryption algorithm is used to exchange a session key, a single-use random string used as a key for a symmetric encryption algorithm. A potential attacker might try to crack the public-key algorithm to obtain the session key and decode the message. However, if the session key is predictable because it comes from a pseudo-random source, the attacker might be able to simply guess the key directly and avoid the difficult problem of breaking the public-key algorithm.

Quantum cryptography consumes random numbers at a far greater rate. Each cycle requires a two-bit random number to select from four basis states, (or an average of at least 2.6 bits/cycle to select from six basis states). Predicting these values would make direct eavesdropping superfluous and undermine the protocol. Quantum cryptography systems have been demonstrated operating at speeds of up to 1.25 GHz. True random number generators are not available that operate at this speed, so these systems must use pseudo-random generators. Proper choice of an algorithm can make discovering the internal state difficult, and the state can be periodically reseeded by a slower random number source. However, this still presents a potential vulnerability that is undesirable. Thus there is demand from both classical and quantum communication protocols to design true random number generators that operate at high speeds.

True random number generators can also be divided into at least two classes depending, on the nature of the physical system. They can be thermodynamic or quantum. Thermodynamic sources are typified by Johnson noise sources. In these random number generators, the voltage on a resistor at finite temperature is measured periodically with positive voltages assigned a "1" and negative voltages are "0". If the polling period is sufficiently long, as determined by the circuit characteristics, sequential measurements will be essentially uncorrelated, and the output can be used as a source of random numbers. Practical considerations, such as stray capacitance and inductance, limit the frequency of operation to a few kHz, nowhere near the speed required. Furthermore, there will usually be some bias in the output – either 0 or 1 will be slightly more likely. Therefore, some processing is needed to remove the bias, further reducing the rate of random number generation. Several Intel microprocessor chipsets include true random number generators based on thermal voltage noise [28], while some VIA microprocessors use frequency noise on free-running oscillators [29] to provide random numbers in the range of 1 Mbit/s. Operating systems such as Linux combine random data from these sources along with other unpredictable data [30], such as interrupt timing from disks and keyboards. All of these systems are presumed to be inherently random, but this can be hard to verify.

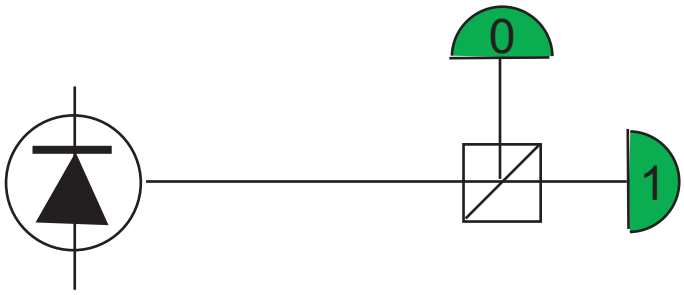

Figure 4.1: Random number generator using a pulse of light to generate a single random bit from detection of a photon. A diagonally polarized photon will be reflected or transmitted by a PBS with 50% probability. These two outcomes are detected by single-photon counters and assigned values of 0 and 1.

4.2 Optical Random Number Generator

Quantum random number generators use the inherent randomness in a quantum process as their source of entropy. One of the oldest quantum sources of randomness is radioactive decay [31]. However, this requires a radioactive source, and higher data rates require more active sources. Sources with short half-lives will also decay with time, an additional complication. Optical sources can provide very fast source of randomness. A small LED or laser diode can emit 10^{16} photons/s, and is small, inexpensive, and safe. Just as the number of radioactive nuclei decaying in a given second is inherently random, a diagonally polarized photon incident on a polarizing beam splitter will be reflected or transmitted each with 50% probability in a completely unpredictable fashion. This suggests the very simple photonic random number generator shown in Fig. 4.1 [32, 33].

Due to the high rate of photon emission by LEDs, the rate of random number generation from optical techniques is limited not by this emission rate, but by the capabilities of the available single-photon detectors. A typical avalanche photo-diode (APD) can run in Geiger mode with an efficiency of 60%, a dead time between detections of 50 ns, and a maximum count rate of 5 MHz, above which the dead time increases and the diode risks destruction by overheating. Photomultiplier tubes (PMT) have lower efficiencies, in the range of 20%, but can have shorter dead times and higher maximum count rates. These parameters limit the speed of the random number generator described above. First, the maximum rate of operation is 10 Mbit/s, with each of the two detectors running at the maximum speed of 5 MHz. However, because of the inefficiency in the detectors, sometimes neither will fire, other times the LED will emit more than one photon and both detectors will fire. Two-photon and zero-photon events, which provide no random bits reduce the theoretical maximum random number generation to 7.5 MBit/s. This is achieved with an optical pulse train at 20 MHz (limited by the dead time of 50 ns) bright enough to cause each detector to fire 25% of the time (to keep the average count rate at 5 MHz). A single bit will be generated when exactly one of

the two detectors clicks, which will happen a fraction $P_{bit} = 2 \cdot 0.25 \cdot 0.75 = 37.5\%$ of the time. As in all physical random number generation, imperfections in the apparatus, such as bias and after-pulsing, require some post-processing to eliminate the concomitant non-randomness, reducing the final generation rate.

However, there are a number of other sources of randomness available from photon detection. First, the 0-photon and 2-photon events also occur randomly, and if properly treated can increase the random bit rate. Furthermore, while the detector count rate is limited, the time resolution can be extremely high, anywhere from 200 ps down to 50 ps [34]. Our random number generator exploits timing information to achieve random bit rates > 20 MBit/s from a single APD. It does this by using as an entropy source the time between photon arrivals on a single detector. A CW LED illuminates a single APD, and the brightness is adjusted to maintain a safe counting rate. Then, the interval between successive clicks is measured by a high speed counter. This interval, measured in units of the high speed clock period, is emitted as a random value at each photon detection.

While the simple beam splitter-based RNG generates 0s and 1s with approximately equal probability, this time-based system generates multi-bit numbers that are not evenly distributed. Because photon detection is a Poisson process, the counter values follow the corresponding waiting-time distribution, which is exponential decay. This requires post-processing to generate a string of uncorrelated random bits. However, since as mentioned, imperfections in the beam-splitter RNG require post-processing anyway, the added burden is not particularly great.

4.3 Analysis

In order to quantify the degree of randomness in arrival time distribution, we use the Shannon entropy [35]

$$S = -\sum_{i=0}^{N} P_i \log_2 P_i, \tag{4.2}$$

where the P_is are the probability of each outcome. The Shannon entropy is measured in bits, and the goal is to generate output with very close to 1 bit of Shannon entropy per output bit. This is the condition where every possible output is equally likely. Other quantities of interest are the min-entropy

$$S_{min} = -\log_2 \max(\{P_i\}), \tag{4.3}$$

which conveys the likelihood of an attacker guessing an entire output string on the first try, and the guessing entropy

$$S_g = \log_2 2N_{50}, \tag{4.4}$$

based on the number of guesses an attacker must make in order to find the right answer 50% of the time. These will be discussed further below.

The post-processing to remove bias and compress the randomness from photon arrival time data into the desired high-density randomness is based on a hash function. Hash functions are a type of mathematical operation that map arbitrary length input strings onto fixed length output strings. The basic idea is to collect a series of clicks from a photon counter while keeping a running tally of the amount of randomness present. The arrival times are written into a buffer, and when enough entropy is present, the entire block is fed into the hash function, which compresses the data into a fixed length block intended to have near perfect entropy density.

To keep track of the entropy in the input buffer, we need to know the probability of each possible value. This can be measured fairly easily, or calculated from a physical model of detection. Here we will use a physical model for simplicity:

$$P_i = \begin{cases} 0 & i < \delta \\ e^{(\delta - i)/\lambda} & i \geq \delta. \end{cases} \tag{4.5}$$

This corresponds to a dead-time of δ and an average time between clicks of λ, both measured in units of the counter resolution. For large λ, the average entropy of this distribution is $\log_2(\lambda) + \log_2(e)$. To ensure that the randomness behavior is from the photon arrival times, the counter resolution should be no finer than the timing resolution of the detector. Higher-frequency counters will still generate random-looking data; however, the source of that randomness has to do with the less well understood properties of the APD. The probability for each measurement is multiplied together (or, their logarithms added) until the probability of a given buffer occurring is less than 2^{-B}, where B is the number of bits output by the hash function. This ensures that there are greater than 2^B possible inputs to the hash function, and if it operated ideally, the output would be perfectly random. This is complicated by collisions in the hash function: cases where different inputs map to the same output. Since it is impossible for a compressive function such as a hash function to operate without collisions, we need to supply some extra entropy to the input buffer to make sure every output from the hash block has approximately the same probability. To do this, we select an entropy threshold, T, which is the number of bits of entropy we need in the hash function's input in order for the output to be very nearly random. We will show below that for typical hash functions a threshold of $T = B + 10$ bits is sufficient to ensure nearly perfect random numbers.

To implement this system, we need to quantify how much excess entropy we must provide as input to the hash function in order that the output will be sufficiently unpredictable. In order to do this, we make some assumptions about the hash function. First, it must take on every output value with the same frequency; for instance, a hash function with a "stuck-bit" that is always either 0 or 1 would not be appropriate for this use. Second, it must not expose the bias inherent in our partially random input. For instance, the input to our hash function will be counter values, the high order bits of which will usually be zero. If some of the output bits depend only on those high order bits, they will be biased as well. One useful property for a cryptographic hash-function to have is the strict avalanche criterion and the bit independence criterion [36]. These properties state that whenever a single input bit is changed, each output bit flips with a 50% probability, and that whether two bits i and j change has no correlation. However, these properties are both very difficult to prove for real hash functions and not sufficient to guarantee correct operation of a true random number generator. Nevertheless, cryptographic hash functions work very well for this application, and we use the SHA family of hash functions [37] to do compression. It is important to note that this application of hashing is not vulnerable to collision attacks, so recently discovered vulnerabilities in SHA-1 [38] are not relevant.

Given the uniform distribution assumption of our hash function, it is easy to calculate the entropy and min-entropy of the output from our RNG. To do this, we need to calculate the distribution of outputs from our hash function, and we do so assuming that it is random, i.e., the patterns in our input data do not cause some hash outputs to be unnaturally preferred over others. In this case, if the hash function has size B bits (and therefore $2^B = Q$ possible outcomes), and the probability of each accepted input is $P(X) \leq 2^{-T} = 1/N$, then we can compute the output entropy as follows: we will conservatively approximate that we have exactly N possible inputs, each equally likely, with probability 2^{-T}; then the probability that m inputs map to a given output is given by the binomial distribution:

$$P = \binom{N}{m} \left(\frac{1}{Q}\right)^m \left(1 - \frac{1}{Q}\right)^{N-m}. \tag{4.6}$$

We can rearrange the sum in Eqn.(4.2) by grouping all terms having the same p_i giving:

$$S = \sum_m^N Q \binom{N}{m} \left(\frac{1}{Q}\right)^m \left(1 - \frac{1}{Q}\right)^{N-m} \left(\frac{m}{N}\right) \log_2 \left(\frac{m}{N}\right). \tag{4.7}$$

The min-entropy is a more conservative measure that only takes into account the most probable outcome of the hash function. Thus, it represents the probability that an adversary could guess an entire hash function register in a single guess. We estimate the min-entropy by looking at the long tail of the binomial distribution

Hash Size	Input entropy	Output entropy	Output min-entropy
128	128	124.8	122.9
128	133	127.98	125.9
128	138	127.999	127.46
256	261	255.98	253.4
256	266	255.999	255.28

Table 4.1: Shannon entropy and min-entropy emitted from the hash function, dependent on the input entropy in bits. The larger the hash block size, the less overhead is required.

and finding the threshold m such that there is roughly a 50% probability of at least one bin having at least m inputs map to it. Both of these quantities have been calculated numerically and are listed in Table 4.1 for several values of the input entropy S_{in} and hash function size B. For common hash function sizes, 10 extra bits of input entropy is sufficient to saturate the output of the hash function. For a 256-bit hash function, an entropy of 255.999 corresponds to an attacker guessing 50.1% of bits correctly, and the min-entropy of 255.28 corresponds to being no more than twice as likely to guess an entire 256-bit block than if it were perfectly random.

4.4 Tests of Randomness

Many applications desire proof that a given source generates random numbers. This is not possible. Testing can show the presence of patterns, but not their absence. Nevertheless, many people have devised tests of random behavior in an attempt to catch common weaknesses in random and pseudo-random generators [39, 40, 41]. However, these tests are essentially useless for cryptanalysis of our RNG. To demonstrate this, we constructed a random number generator in which we replaced the photon arrival times with an easily cryptanalyzed linear-congruential PRNG. The resultant generator passed all of those tests, despite being easily cracked and containing zero entropy (i.e., the output is entirely predictable, if you know the algorithm used). Unsurprisingly, when a correctly working generator is used, it also passes all of those tests.

Instead of testing for randomness, then, a high-quality random number generator must be built defensively to have confidence in a calculated lower bound for the entropy content of the output. This requires carefully characterizing the behavior of the input entropy source, measuring the extent of possible defects in its behavior, and accounting for them. In our generator, we look for several types of correlations, and use them to adjust the estimate of the entropy calculator which decides how many counter values go into each hash block.

The first and simplest bias to look for is deviation from an ideal Poisson distribution. A number of factors cause this deviation, the most important of which are after-pulsing and low efficiency at the end of

dead time. Due to trapped electrons in an APD, there is a small chance ($\sim 1\%$) of exhibiting breakdown without absorbing a photon, usually very shortly after the end of dead time. Furthermore, the measurement apparatus itself might have some bias, for instance, the low-order bit might be considerably more likely to be 0 or 1 if the clock driving it is not exactly a square wave. These types of bias are easy to account for—we simply measure the actual distribution of photon arrival times, and use the probability for each value to calculate its contribution towards the total entropy. If some outcomes are substantially more likely than they should be, their estimated entropy contribution will be reduced accordingly.

The second form of bias is more difficult to detect, and that is correlation between sequential or nearby counts. We expect these correlations to be relatively low, due to the asynchronous nature of our generator. A synchronous generator that measured the time between a starting clock and the next photon arrival time would potentially have strong correlations, since the detection time in one clock cycle would determine how likely an early detection in the next cycle was due to dead time, after-pulsing, and detector efficiency. However, in this asynchronous mode of operation, the counter always starts immediately after a click, which presumably erases most of the existing state of the detector. There are, however, still some sources of correlation possible, such as fluctuations in the LED brightness or detector efficiency due to, for example, power supply fluctuation that cover multiple detections. We search for these in the autocorrelation function of the input counter data. This is something of a brute-force approach, and cannot detect many forms of correlation; however, it is a reasonable test in the absence of a strong candidate for a source of correlations. As shown in Fig. 4.2, we find no significant correlation between nearby results. Multi-point correlation functions could be considered, however there is no basis to expect such behavior, and therefore no indication of how high-order would be the correlations needed to be examined.

Very long timescale correlations are difficult to test for by autocorrelation. However, the most likely source of correlations is slow drift in the count rate, or sensitivity to 60 Hz fluctuations from power lines. In order to combat these factors, we keep a running average of the mean count rate that feeds into the entropy calculator, allowing it to follow gradual changes in the system.

4.5 Conclusion

We have implemented this random number generator using an APD at a maximum count rate of 4 MHz, yielding random bits at > 20 MBit/s. This is considerably faster than any other quantum random number generator we know of, and uses only a *single* photon counter, rather than two, as in some other designs [32, 33]. The bit rate can be increased further by increasing the time resolution of the measurement system,

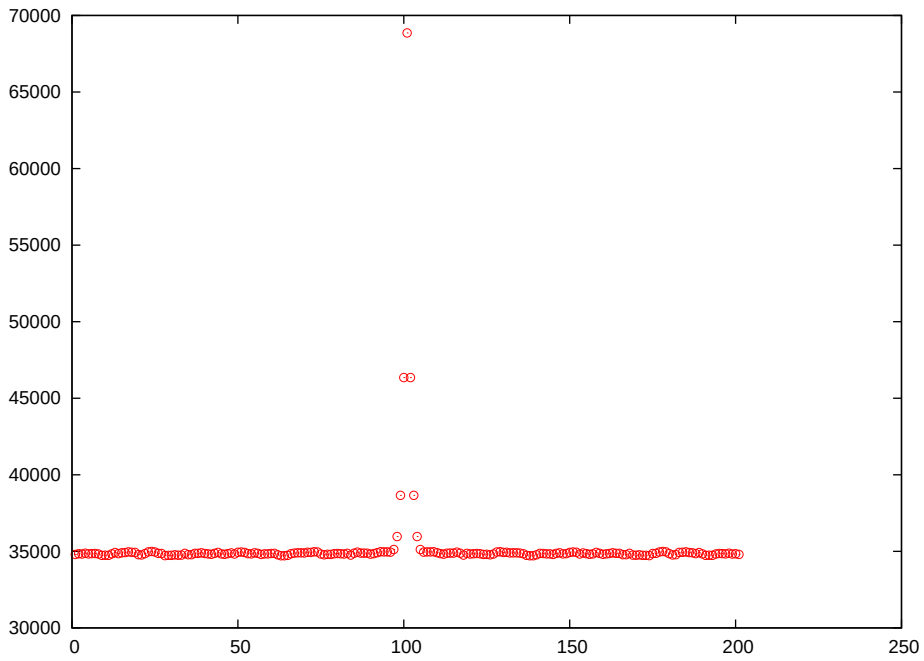

Figure 4.2: Autocorrelation on a sample of raw detection-interval measurements. The spike at $\tau = 0$ shows that the data is perfectly correlated with itself at zero lag. At other lags, the correlation drops to the baseline level.

which is 5 ns. The jitter of the photon detectors is approximately 200 ps, improving our time resolution to that level would increase the entropy per photon by approximately $\log_2(5ns/200ps) = 4.6$ bits, nearly doubling the theoretical bit rate. Increasing the resolution beyond this level is of questionable value, since some of the entropy will be from the jitter in the detection process, not the absorption of photons themselves. While we have no reason to believe this is anything but random, we prefer to gather entropy only from the relatively well understood behavior of photons. Furthermore, since the data rate increases as the log of the time-resolution, increasing the timing resolution without increasing the count rate can only give small gains at some point. Instead, to increase the random bit rate, we would need to use a detector such as a PMT, which can operate at higher count rates. A hypothetical PMT with a 50 MHz maximum count rate and the same 200 ps jitter could achieve approximately 400 Mbit/s random bit generation.

Our generator is ideal for use in quantum cryptography, since it can provide random numbers at such a high rate. For extremely high speed cryptography systems, it is still not nearly fast enough. These systems would require either several generators running in parallel, or to expand the data stream using a pseudo-random generator. This would, of course, somewhat compromise the security of the overall system; however, the more random input to the generator, the harder it will be for an eavesdropper to crack the RNG and intercept the secret communication.

Chapter 5

Relativistic Quantum Cryptography

5.1 Introduction

Quantum cryptography, also known as quantum key distribution (QKD) was the first, and is currently the only commercially implemented quantum information application. Within the validity of certain assumptions, it allows uncrackable encryption, an alluring prospect to organizations protecting secrets whose value is incalculable. This is achieved by exploiting the no-cloning property of quantum information [42], which implies that it is impossible to copy an unknown quantum state without disturbing the state of both the original and the copy. Thus, QKD provides secure communication based on the laws of quantum physics, rather than mathematical algorithms believed (but not proved) to be difficult to reverse.

The most common QKD protocol is BB84 [43]. In this protocol, the canonical parties, Alice and Bob, mutually agree on a shared, secret sequence of bits which can be used as a key for a conventional cryptographic protocol. This can be either a Vernam cipher, a provably secure classical encryption algorithm, or a traditional symmetric cipher such as AES [44]. While the use of AES or other classical encryption algorithm compromises the "perfect security" claim, high quality symmetric ciphers are generally regarded as less vulnerable than the public-key algorithms currently used for key exchange, such as RSA [45]. All known public-key algorithms are also vulnerable to quantum computers [4], while symmetric block ciphers are not known to be. Although a quantum search (e.g., Grover's algorithm) can speed up the search through a symmetric cipher's key space, it only reduces the search speed by a factor of $\sqrt{N}$. The time required to brute-force attack such a cipher is still exponential in the key size, so this threat can be protected against by simply doubling the key size. Doing so would usually only double encryption and decryption time, but would dramatically increases the time required to break the cipher.

Every QKD protocol has a maximum channel noise (the bit error rate, or BER)under which it can operate. For BB84, this is approximately 11%. More advanced protocols, such as the six-state protocol [46] (SSP) can tolerate higher error rates at the expense of lowered efficiency on low-noise channels. We have implemented a new protocol which uses additional constraints coming from special relativity to eliminate

one type of loss, so-called sifting loss. This relativistic quantum cryptography (RQC) protocol increases the efficiency of both BB84 and SSP, while making the SSP advantageous regardless of channel noise. RQC is one of only a small number of communication protocols that combine quantum mechanics and relativity to achieve something not available with either alone, and the only one that has been implemented.

5.2 Quantum Key Distribution

In BB84, Alice transmits single photons to Bob in one of four states of polarization, horizontal ($|H\rangle$), vertical ($|V\rangle$), diagonal (45°, or $|D\rangle$) or anti-diagonal (−45°, or $|A\rangle$). Bob then measures each photon in one of two bases, either the rectilinear basis (H/V) or the diagonal basis (D/A). Alice and Bob announce on a public channel which basis, but not which polarization, they used for each photon. In each case where they agree, they have perfectly correlated results, while in cases where they disagree their results are random. The cases where they agree form the sifted key, and are assigned bit values according the actual polarization sent or measured in each case. The presence of an eavesdropper, Eve, will introduce errors into the sifted key. Alice and Bob find and correct those errors using public communications and a classical error correction protocol, keeping track of how much information they release about the key as they do so. They assume that all of the errors they found were due to noise induced by an eavesdropper (a conservative assumption – there will always be some unavoidable errors due to imperfections in the apparatus or channel). Now, Alice and Bob have identical bit strings, while the hypothetical eavesdropper has less than perfect information about it. By computing hash functions [47] of their data and throwing some of it away, they can reduce the eavesdropper's knowledge such that she has an arbitrarily low chance of knowing even a single bit of the final key.

One important parameter for a QKD protocol is the robustness against errors. If Eve's potential knowledge of the sifted key is greater than Bob's, Alice and Bob will be left with no secret bits after error correction and privacy amplification. Effectively, the information revealed by Alice during error correction will allow Eve to completely correct her copy of the key. Privacy amplification is then impossible. It is possible for Alice and Bob to use two-way communication to implement a procedure known as advantage distillation to increase Bob's knowledge relative to Eve so that error correction and privacy amplification may be used. This procedure is usually very inefficient, however, and is rarely used in practice.

The mutual information, measured in bits per photon, between Alice and Bob (I_{AB}) is the same for all protocols, and reflects the bits sacrificed to error correction. This is given by

$$I_{AB} = 1 + (D\log_2(D) + (1-D)\log_2(1-D))\eta, \tag{5.1}$$

Where η is a parameter of the error correction algorithm. The Shannon limit corresponds to $\eta = 1$, the best possible error correction algorithm. For CASCADE [48], the most commonly used protocol, $\eta \approx 1.16$. For a given protocol, Alice and Bob can compute the maximum knowledge of an eavesdropper in terms of the disturbance caused by her tampering. For BB84, Eve's information on the sifted key is limited to [49]

$$I_{AE}^{BB84} \leq \frac{1}{2}\phi(2\sqrt{D(1-D)}) \tag{5.2}$$
$$\phi(x) \equiv x\log_2(x) + (1-x)\log_2(1-x).$$

In order to efficiently generate a secret key (i.e., without advantage distillation), Eve's information on the key must be less than Bob's. As the disturbance increases, and Eve's information approaches Bob's information, the efficiency of BB84 drops precipitously. One way to combat this is to use the six-state protocol [46]. This protocol adds an additional communication basis, the circular basis, consisting of right-circular polarization ($|R\rangle$) and left-circular polarization ($|L\rangle$). An optimal eavesdropper gains less information for a fixed disturbance [50]. This allows operation over channels having more noise and a corresponding higher efficiency in sifted-key to final-key processing for any disturbance. Unfortunately, since Alice and Bob only use the same basis 1/3 of the time, the sifted key rate is lower than it would be for BB84, and the overall efficiency is lower for all but the highest error rates. Maximal efficiency vs. disturbance is shown in the theory curves of Fig. 5.4 for each protocol discussed.

5.3 Delayed Choice Quantum Key Distribution

Our new QKD protocol allows Bob to measure each photon in the correct basis, without providing to Eve further useful information. This eliminates the sifting step, in principle increasing the efficiency of BB84 by 100% and SSP by 200%. In order to do this, Bob stores the photon he receives in a quantum storage cell until after Alice announces her basis. Bob can then measure with confidence, choosing the same basis as Alice each time and eliminating the inefficiency caused by random selection. It is possible to gain a similar advantage if Alice and Bob use one basis much more frequently than the other, while only occasionally checking in the other basis [51]. This can also dramatically reduce the sifting loss, though not eliminate the process entirely.

The primary advantage of RQC is the increased efficiency. If the efficiency of the quantum storage is close to 100%, RQC can exceed the efficiency of the traditional protocols by a considerable margin. A secondary advantage is that it eliminates one form of attack against QKD protocols in which Eve uses basis disagreement to hide the disturbance caused by her probe. It is possible for an eavesdropper to determine

from her probe whether or not she induced an error in Bob's measurement, then tamper with the classical channel to force Alice and Bob to discard those events, hiding her disturbance in the sifting process. Of course, proper authentication of the public channel protects against this form of attack; however, it is impossible for Eve to perpetrate this when there is no sifting.

There are a few limitations imposed by the RQC protocol as well. Most importantly, since Alice and Bob rely on preservation of causality to protect against Eve, they must ensure that the events corresponding to Bob receiving a photon and Alice transmitting the measurement basis are causally separated. A second inconvenience is that Bob cannot rely on passive basis selection using a beam-splitter, but must actively choose his analysis basis, using, e.g., a Pockels cell. Finally, it is possible to reduce the upper bound on Eve's information by assuming that she is not able to wait until the basis is announced to measure her ancilla. This is not done in practice, since one cannot ensure that Eve cannot indefinitely store the ancilla she uses to measure Alice and Bob's photons. However, in RQC it is much easier for her to actually implement such an attack with modern technology. Nevertheless, since Alice and Bob conservatively assume Eve has such technology already, this does not affect the final key rate.

The security condition for this relativistic protocol can be seen from a space-time diagram of the transactions involved, shown in Fig. 5.1. The two relevant events are Alice's transmission of the basis information, and Bob's receipt of the quantum signal. As long as the latter is not causally after the former, the protocol is equivalent (from an eavesdroppers perspective) to the standard protocol, either BB84 or SSP. The simplest way to understand this is to realize that if Bob measures his photon immediately, the protocol becomes the non-delayed version. Eve will, in general possess a qubit entangled to Bob's photon; however, since entanglement does not allow signaling, her reduced density matrix and the information she gains from measuring it cannot depend on what measurement Bob performs, or when he does so.

In order to implement RQC, Alice and Bob must verify that the causality constraint illustrated in Fig. 5.1 is satisfied. In order to do so, they must know their relative separation in space and time to relatively high accuracy – the uncertainty in their space-time separation must be added to the storage time. For typical storage times of a few hundred nanoseconds, this requires that their space-time coordinates be known to better than 100 ns, a reasonable figure. Furthermore, Bob must have a quantum storage capable of preserving the state of a photon with low noise and low loss. We implement this using the cavity-based storage described in Chapter 3.

The minimum storage time required depends on several factors. The first is a fixed interval, dependent on Bob's apparatus. If he receives both the quantum and classical signal at nearly the same time, he must still be able to store his photon for as long as it takes to decoded the classical basis signal and set his analysis

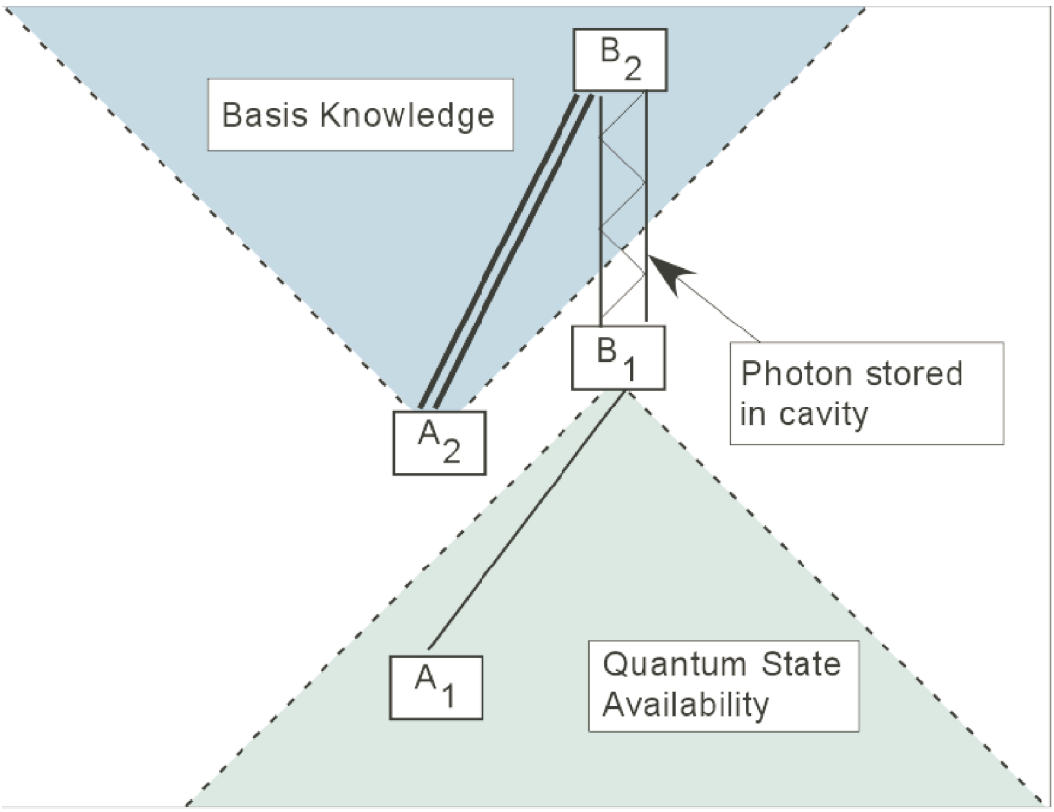

Figure 5.1: The required relative ordering of events in relativistic QKD, shown on a standard space-time diagram. If the two shaded cones overlap, the protocol is not secure, as an Eavesdropper in the intersecting coordinates could get the basis information and use it to measure the signal photon without disturbance.

Pockels cell to measure in the appropriate basis. For our apparatus, this is roughly 100 ns. The second is the uncertainty in the timing-positioning. We used a shared high-frequency clock to establish a common time base with an uncertainty of 20 ns. The third factor is the transmission speed of the classical basis signal. As shown in Fig. 5.1, the slower the classical signal, the longer until Bob receives it, and the more time he must store the quantum signal. If the signal is transmitted optically, as in our experiment, this factor is essentially zero; however, if the signal travels through coaxial cable or fiber, additional delay is required. Such delay is particularly important, since it is proportional to the distance between Alice and Bob, while the other delays are fixed. For free-space optical communication, this factor is negligible. Finally, in some cases, Alice may have to wait a period of time before transmitting her basis selection. This is particularly important when using a passive-choice entangled source, since Alice does not know the basis until after one member of the photon pair is detected. In this case, Bob must add the latency of Alice's detectors (13 ns for our APDs) and the time to encode the basis message to his storage time. This adds approximately another 150 ns to the total delay.

5.4 Classical Communication

In addition to the storage cavity discussed in Chapter 3, the other key element of an RQC implementation is the classical communication. The limited available storage time requires a very low-latency communication system. Most widely available communication systems trade latency for higher bandwidth and improved features, such as multiple access. For instance, 100 Mbit Ethernet has a minimum latency, required to detect collisions from multiple transmitters, of 6.4 μs, far too long to be useful. A single packet is at least 64 bytes long, whereas we only need 2 bits to identify the basis.

Our communication system is an optical serial link implemented on a programmable logic device (CPLD). Such devices allow logic gates to be connected in a variety of ways, and are reprogrammable using flash memory. We [1] implemented a finite-state-machine (FSM) which, upon recognizing an input, identifies which basis to send, encodes that into a 5-bit message (only 2 are needed, the others are used for framing and timing), and sends it to a laser diode. A similar CPLD on Bob's side recognizes the code from Alice and sets the Pockels cell to the appropriate value. The CPLDs also output data to a PC using a digital IO board, allowing us to record the key data.

With the CPLD we used (Xilinx 9572XL), the maximum frequency the state machines can run at is 55 MHz, and the total end-to-end latency of the system is approximately 120 ns, not counting propagation between the modulator and demodulator. Substantially faster devices are available, based on FPGA (field

[1] Thanks to Matt Brenner for the initial modulator design

programmable gate array) technology, which also allow much higher logic density. While the latency is already acceptably low, using a more powerful logic device would allow further stages of the QKD protocol, such as error reconciliation and privacy amplification, to be implemented in hardware, reducing demands on the communication speed to the PC. Currently, we simply poll the output of the CPLDs from the PC, consequentially generating large amounts of data which must simply be immediately discarded, since most of the time the apparatus is idle.

5.5 Implementation

We implemented RQC with both a classical faint-pulse laser and an entangled photon source. The classical laser source is BB84 only, while we implemented both BB84 and SSP with the entangled source. In all cases we find an improvement in the overall data rate for the relativistic version. As expected, the difference is most pronounced (as high as 167% improvement) for the SSP. We also simulated an eavesdropper, showing how the final key rate decreases as an eavesdropper induces noise.

5.5.1 Classical Source

The classical laser source is shown in Fig. 5.2. A separate laser is used for each polarization state, $|H\rangle$, $|V\rangle$, $|D\rangle$, and $|A\rangle$. Each laser is supplied a small bias current of approximately 5 mA through a bias tee. A 2 ns electrical pulse then travels through a high bandwidth 4-way MOSFET switch, which directs the pulse to an adjustable attenuator and then the modulation input of one of the bias tees. Each attenuator is adjusted to balance the relative brightness of each laser at the output of the source. The lasers are combined with polarizing beam splitters, a waveplate to produce $|D\rangle$ and $|A\rangle$, and beam splitter into a single-mode fiber, which eliminates distinguishing information from the spatial mode. Each laser is in a temperature-controlled mount, which is adjusted to match them to the same wavelength.

The first implementation of the classical source suffered from considerable variances in the laser diodes. Despite driving the lasers with the same pulse generator through a switch, the pulse width was not the same for the two lasers (3.8 ns versus 4.1 ns). This is a form of distinguishing information that an eavesdropper could use to determine which laser each pulse came from without measuring or affecting the polarization state – the only degree of freedom Bob can detect. Furthermore, the lasers could not be matched to the same wavelength within the achievable temperature range. To overcome this difficulty, we tested many laser diodes to find a set of four that have similar characteristics, allowing much better matching of pulses. This also highlights one of the problems with faint-laser BB84. While we reduced the distinguishing information

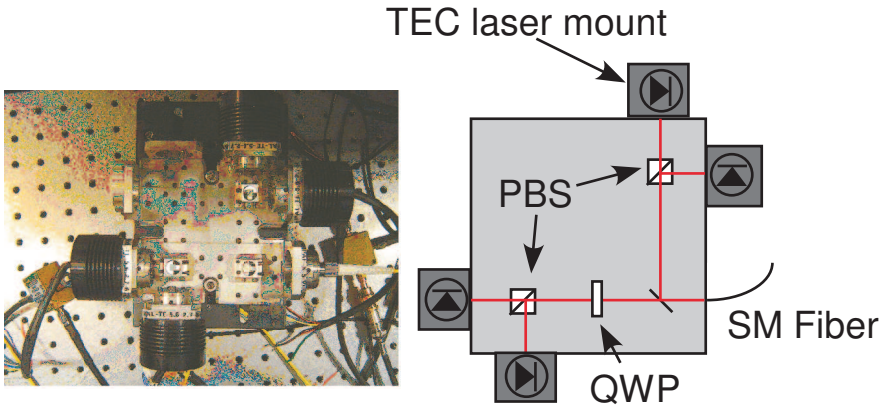

Figure 5.2: Left: Photo of the classical laser source. Right: Schematic of the same source. Light from four laser diodes is combined such that each laser generates one of the four BB84 polarization states. Pulses from these lasers are coupled into a single-mode fiber to clean up the spatial mode, then attenuated until the average photon number is less than 1/pulse.

we know about, it is always possible that some information could leak out. In order to actually show that two states are indistinguishable, we would need to show interference between them. For instance, if single pulses from two lasers were combined on a screen or camera, they would show a interference pattern. The visibility of the interference fringes depends on the distinguishability (D) of the two pulses, according to the formula

$$V^2 + D^2 \leq 1. \tag{5.3}$$

However, this is very difficult to do. The distinguishability must be on the order of 5%, leading to a minimum visibility of 99.9%. It would be very difficult to achieve a signal-to-noise ratio high enough to see this in a single-shot fringe pattern. A better way might be to use 2-photon interference with a Hong-Ou-Mandel interferometer, but that would also be a very difficult experiment.

5.5.2 Entangled Source

An alternative solution uses entangled photons. Entanglement is only possible in the absence of distinguishing information. Any variation in mode shape, frequency, or timing between the polarization states will cause decoherence that will introduce errors, exactly as if Eve had taken that information herself. It is for this reason that entangled QKD can be secure even if the source is controlled by a third party, even the eavesdropper – as long as Alice and Bob measure quantum correlations they can find a bound for how much information Eve can have about the results of their measurements.

Our entangled photon source for QKD was described in Chapter 2. We use a two-crystal Type-I source as described in Chapter 2. In this case, the crystals (BiBO) are tipped such that the cones collapse into beams

collinear with the pump. The pump is removed with a laser edge filter, leaving only the downconversion: Bob's photon has a wavelength of 670 nm, chosen to match the wavelength of the classical source described in the previous section, while Alice uses the conjugate wavelength, 737 nm. These two beams are split by a dichroic mirror, and coupled into single-mode fibers. By coupling Bob's photon into a fiber, we can connect either the classical source or the quantum source to Bob's apparatus. In order to get high quality states, we then have to couple Alice's photons into fiber as well. Back in free space, Alice's photon goes through beam splitters to randomly select which basis she will measure in. Each of the three beams goes to waveplates and a PBS to measure it in the appropriate basis. This measurement projects Bob's photon into a known polarization [52]. This beam-splitter based state selection is essentially a simple quantum random number generator, as described in Chapter 4. From this point, operation is exactly the same as with the classical laser source, with two major exceptions: as mentioned, any correlation between Alice's polarization measurement and external degrees of freedom will cause detectable decoherence, alerting Alice and Bob to its presence, and the multi-photon probability is essentially eliminated. Classical laser pulses have a significant probability of producing multiple photons in a single pulse, which can be a problem as the eavesdropper can remove a single photon without Bob noticing. However, decoy state protocols[2] [53] allow Bob to detect this attack, significantly reducing the importance of multi-photon events. In the entangled case, however, even in the rare multiple-pair event, there are no correlations *between* the pairs, so splitting off one photons is less helpful to Eve.

5.5.3 Polarization Analysis

One of the main advantages of RQC is that it increases the efficiency of the six-state protocol. However, since it requires Bob to perform active basis selection, he must be able to switch between three different bases quickly. We accomplish this with an electo-optic device (Pockels cell) which rotates the polarization in response to an applied voltage. A typical arrangement would include *two* Pockels cells, one set at 22.5° to select the $|D\rangle/|A\rangle$ basis, and the other at 45° to select $|R\rangle/|L\rangle$. In our implementation, we eliminate one Pockels cell by having a single device select both the diagonal and circular bases. A single device that rotates around the {1,1,1} axis, equidistant from $|H\rangle$, $|D\rangle$, and $|R\rangle$, can select the three analysis bases by applying a voltage of $-V$, 0 or $+V$, where an applied voltage V causes a 1/3 wave shift, shown in Fig. 5.3. Pockels cells can only rotate around a axis corresponding to linear polarization (they behave exactly as variable waveplates), but this effect can be simulated with the addition of fixed waveplates.

The entangled state created by Alice is subject to several stages of polarization rotation on both photons.

[2]Photon number-splitting is in effect a non-linear absorption. Decoy state protocols use pulses of different mean brightness to detect this form of attack.

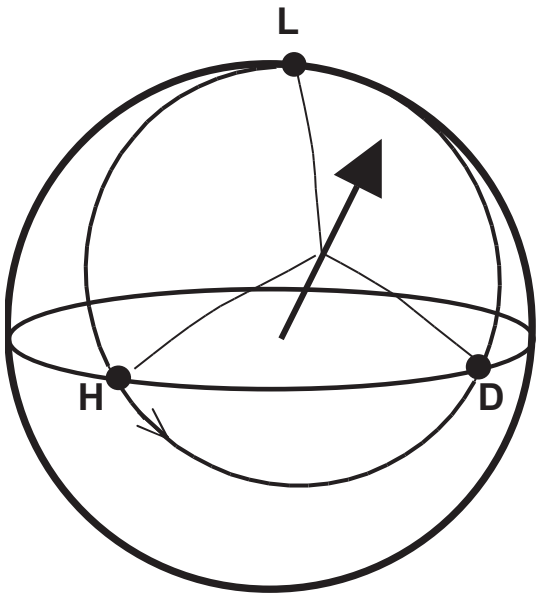

Figure 5.3: Rotation around the $\{1, 1, 1\}$ axis on the Poincaré sphere allows access to all three communications bases in the six-state protocol with a single Pockels cell. With no voltage applied, the projection is in the H/V basis, while applying a positive or negative voltage allows measurement in the D/A or R/L basis.

First, the state created acquires some phase shift traveling through the birefringent crystals. The produced state is thus $(|HH\rangle + e^{i\phi}|VV\rangle)/\sqrt{2}$. The specific states created by remote state preparation depend on the value of this phase. Furthermore, both photons acquire a geometric phase traveling through the fiber-optic mode filters, and Bob's photon experiences a rotation in the storage cavity. All of these polarization shifts act to effectively randomize the final entangled state that Alice and Bob share. However, as long as the rotations are unitary, they preserve the entanglement, and any maximally entangled state is equivalent to any other by unitary rotations on one of the photons. Therefore, instead of measuring and correcting for each step of polarization rotation, we simply make a complete state tomography (as described in section 2.7) between Alice and Bob with all the rotations included, then set Alice's analysis waveplates to match the three bases that Bob can project into. This allows us to eliminate all fixed waveplates except for a single HWP after the Pockels cell. We use this HWP to tune the effective angle between the Pockels cell rotation axis and the PBS, which needs to be the same angle (on the Poincaré sphere) as the angle between $|H\rangle$ and the {1,1,1} direction, i.e., 27.4°. This drastically reduces the number of waveplates needed, with the disadvantage that we no longer use the convenient axial, diagonal, and circular bases, but instead three essentially random bases, though they retain their relative orientations, begin at right angles to each other on the Poincaré sphere.

5.5.4 Eavesdropper

In order to show the dependence of the various protocols on error rate, we simulate an eavesdropper, and show the resulting increase in BER and decrease in key rate. The basic function of an eavesdropper is to intercept each qubit, and entangle its state with ancilla qubits, then measure those ancilla to guess the state of the original qubit. The simplest eavesdropper is known as an intercept-resend attack, in which Eve measures photons randomly in one of the communication bases, then resends a duplicate photon with the detected polarization. It is easy to show that if Eve does this on every photon, she introduces a 25% BER for the BB84 protocol: 50% of the time, Eve will randomly select the wrong basis, and 50% of those cases will result in an error detected by Alice and Bob. For the six-state protocol, Eve will choose incorrectly 67% of the time, and induce a 33% error rate. However, this is not optimal for Eve. She can do better by measuring in a direction equidistant from each basis. For the standard BB84 protocol, this is the 22.5° basis, halfway between $|H\rangle$ and $|D\rangle$. For the six-state protocol, this is the so-called Britebard basis, in the {1,1,1} direction on the Poincaré sphere [54]. In fact, this is one way of seeing the advantage of the six-state protocol: the optimal eavesdropping basis is 27.4° from the communication bases, the larger angle causing more disturbance and gaining the eavesdropper less information than she could get in the BB84 protocol. We simulate an optimal eavesdropping strategy using quartz decoherers. These are thick, birefringent optical elements that cause polarization decoherence by entangling polarization to the frequency degree of freedom. The net effect is the same as if entanglement were to a qubit Eve kept, although of course in reality she sends her ancilla to Bob along with the disturbed polarization state. Nevertheless, the induced disturbance on the QKD polarization qubit is identical to that expected in an actual eavesdropping attack.

5.6 Results

We show an improvement in total secret key rate for RQC with both the classical laser source and the entangled photon source. Relative data-rates are shown in Fig. 5.4. With the classical source, we only demonstrated BB84, and the gains are relatively small (12%), mostly due to loss in the storage cavity. In this system, we used a 454-ns delay line with a total efficiency of only 66%. This, combined with a slightly higher error rate when the delay is included, lead to a marginal improvement over the non-relativistic protocol. We can expect that a similar design using the six-state protocol would have a greater relative gain: the cavity loss and noise would be approximately the same, but the sifting loss we avoid would be much higher. In addition, as discussed above, we assumed that the polarization state was not available to Eve through some other degree of freedom. We expect the spatial modes of the different state are identical,

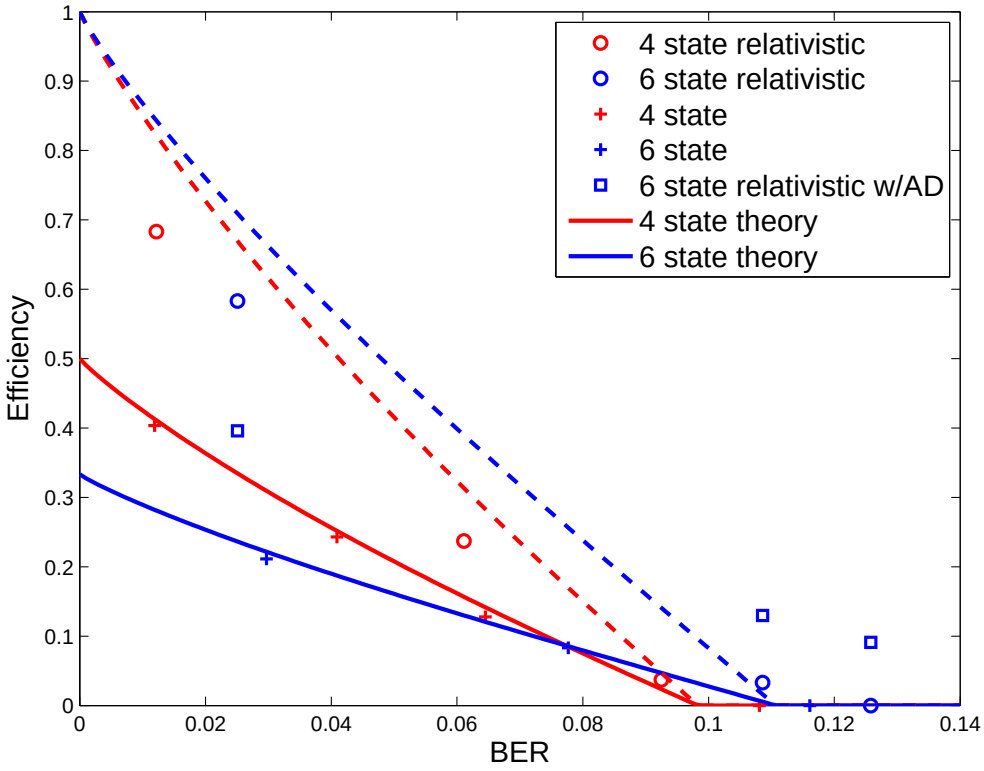

Figure 5.4: Relative efficiency for several variations of RQC. The efficiency is the number of final secret key bits over the number of photons detected by Bob, with a correction factor for loss during storage. The higher error rates for each protocol have a simulated eavesdropper.

due to the use of a fiber mode-filter; however, we can only claim that the states are indistinguishable by frequency or timing to within our ability to measure them. The entire system operates at a 10 kHz repetition rate, with a final key rate of 840 bit/s.

With the entangled-photon source we show a greater advantage for the relativistic protocols. This is primarily due to reducing the required delay time, using only a 350-ns delay with approximately 85% transmission efficiency. We first characterized the entangled state using quantum state tomography [21], to verify we have nearly a maximally entangled state. While this is not strictly necessary, as Alice and Bob automatically place a bound on the degree of entanglement from their measure BERs, it is experimentally useful. The tomography shows the entangled state to have a tangle $T = 93.1\%$ and a fidelity with the nearest maximally entangled state of $F \geq 97.8\%$, resulting in a minimum average BER of 2.5%, and an efficiency of 0.73 bits/photon detected by Bob. Due to loss in the cavity, we consider this to be equivalent to only 0.62 bits/photon for comparison against the non-relativistic protocol. Nevertheless, this rate is considerably greater than what we achieved in the non-relativistic variant (which we implemented by by-passing the

optical delay line), and greater than is possible even in a noise-free system.

Although our system operated at a relatively low total rate (425 raw bits/second) due to the entanglement brightness and maximum switching speed of the Pockels cell, it shows a distinct advantage over what a similar system could accomplish without incorporating the relativistic element. Furthermore, the storage aspect does not limit the speed of operation at any realistic level. Since the storage is a simple delay, it can automatically store many photons in different time bins. A high-speed RQC system, using many other quantum memories would require one qubit for each "in-flight" photon. We consider it unlikely an implementation could use less than 100 ns of delay, which for a 1.25 GHz QKD system would require 13 separately addressable storage qubits.

Chapter 6

Quantum Orienteering

The contents of this chapter have previously been published by Jeffrey, Altepeter, Colci, and Kwiat as "Optical Implementation of Quantum Orienteering", *Phys. Rev. Lett.*, 96:150503 (2006).

6.1 Introduction

Quantum orienteering is a quantum communication protocol addressing the problem "how can Alice most efficiently communicate a direction in space to Bob?" While cryptography involves communication resistant to a malicious adversary, there is no such concept here. Rather, Alice and Bob are cooperating parties, attempting to communicate a direction in the most efficient manner possible, i.e., the most precision with the fewest resources. This problem can be solved using classical or quantum communication. If Alice and Bob share a reference frame (such as that provided by well-known distant stars, or the earth's magnetic field), they can transmit a classical signal, e.g., a binary encoding of the polar and azimuthal angles, allowing straightforward communication of a direction to a precision limited only by the number of bits they use. This protocol is simple and efficient; however, if they do not share a frame, Alice and Bob must transmit a *physical* object, such as a spinning gyroscope whose angular momentum points in the desired direction. No amount of "pure" information can convey a physical direction due to the symmetry of space. These physical objects can be either quantum or classical. Classical objects such as gyroscopes may, in principle, be measured to arbitrary precision, yielding a perfect transmission fidelity. However, for small systems, such as a single spin-1/2 particle, quantum mechanical uncertainty limits the accuracy with which Bob can estimate the direction. Here we examine how Alice and Bob can cooperate to maximize the information Bob can extract from such a system.

For quantum orienteering using *two* spin-1/2 particles, it has been proven that a straightforward measurement of the two particles independently does not yield the optimum fidelity of communication, but can be beaten by a joint measurement on the two particles [55, 56]. Furthermore, if Alice changes her encoding

by flipping the direction of the second spin, Bob can make even more precise measurements [57]. Others have investigated optimal strategies for more than two spin-1/2 particles [58, 59] but here we consider only one- and two-spin cases. Quantum orienteering is similar to optimal detection [60] and unambiguous state discrimination [61, 62], protocols for distinguishing between a fixed set of non-orthogonal states. Recently, Pryde *et al.* [63] showed the advantage of using collective measurement to distinguish whether a pair of spins were parallel or anti-parallel. In contrast to these, orienteering involves providing the best estimate for an arbitrary state or direction. We have designed and implemented a sample protocol using photon polarization as our effective spin-1/2 system. Below we discuss the advantages and limitations of such an implementation.

6.2 Optimal Measurements

In order to compare different protocols for orienteering, we must quantify the accuracy of transmission. We define the average transmission fidelity as

$$F \equiv \int d\Omega \frac{1+\cos(\theta)}{2}, \tag{6.1}$$

where θ is the angle between Bob's guess and the direction Alice tried to send, and the integral is an average over all possible directions Alice might choose. For a random guess, then, the average fidelity will be 1/2, while perfect transmission gives a fidelity of 1. When Alice is only allowed a single spin-1/2 particle to transmit a direction, Bob's measurement is simple: he picks any direction $\hat{r}$ and measures the spin in that direction, e.g., with a Stern-Gerlach magnet. Bob's guess then is $\pm\hat{r}$, depending on whether the outcome of his measurement is $\pm 1/2$. This achieves a transmission fidelity of 2/3, the optimal value using only a single spin [56].

When Alice sends Bob *two* spins polarized in the same direction, Bob has several options. The simplest extension uses two independent measurements, one on each particle. The best way to do this is to measure in bases at 90° to each other, e.g., along $\hat{X}$ and $\hat{Y}$. This gives a fidelity of 73.3%. However, if Bob instead makes a single collective measurement on the two spins, he can make a slight improvement to $F = 75\%$ [56] by projecting into the basis

$$|\psi_k\rangle = \frac{\sqrt{3}}{2}|\hat{n}_k, \hat{n}_k\rangle + \frac{1}{2}e^{i\phi_k}|\psi^-\rangle. \tag{6.2}$$

The $|\hat{n}_k\rangle$ are the tetrahedral states shown in Fig. 6.1, while $|\psi^-\rangle$ is the anti-symmetric Bell state $\frac{|HV\rangle - |VH\rangle}{\sqrt{2}}$. Adjusting the phase term $e^{i\phi_k}$ allows the set of states to be made orthonormal (see Table 6.1).

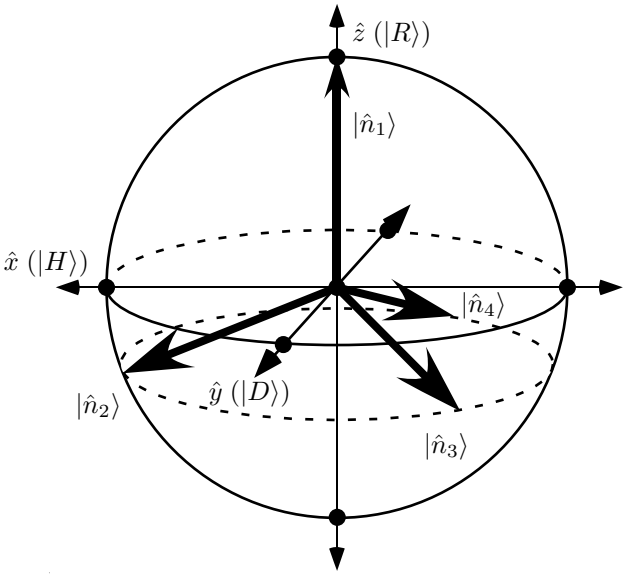

Figure 6.1: The four $\hat{n}_j$ equally spaced directions in space corresponding to the corners of a regular tetrahedron. In our implementation we associate directions in real space with directions on the Poincaré sphere of polarization states: right-circular polarization ($|R\rangle$) is mapped to the $\hat{z}$ direction, horizontal polarization ($|H\rangle$) is mapped to $\hat{x}$, and 45° polarization ($|D\rangle$) to $\hat{y}$.

Remarkably, even this is not the optimal protocol, though it is the best Bob can do if Alice polarizes both spins in the same direction. If instead Alice polarizes the second spin in the opposite direction (and informs Bob she is doing so!) Bob can make a different set of joint measurements [64]:

$$|\psi'_k\rangle = \frac{\sqrt{3}}{2}\frac{|\hat{n}_k, -\hat{n}_k\rangle + |-\hat{n}_k, \hat{n}_k\rangle}{\sqrt{2}} + \frac{1}{2}e^{i\phi'_k}|\psi^-\rangle. \tag{6.3}$$

Again, the ϕ'_k's are chosen to make this an orthonormal basis. Using these measurements, Bob can increase the fidelity of his guess to 78.9%, the maximum achievable using 2 spin-1/2 particles [65]. The advantage of 2 anti-parallel spins over 2 parallel spins has frequently been qualitatively explained as follows: the former states span the full Hilbert space, while the latter do not (they have no overlap with the singlet entangled state). Thus, the parallel spins have a smaller effective "alphabet" for encoding the direction. However, this intuitive argument has not been made rigorous, and recent work has in fact shown that the argument is incomplete and actually fails in some cases [66].

6.3 Photon Implementation

While it is common to regard all two-level systems as equivalent, for this application they are not. Here we are not only interested in the geometry of state space, but its *embedding into a physical object*. It

	(X,Y,Z)	ϕ_k	ϕ'_k
$\hat{n_1}$	$(0, 0, 1)$	0	0
$\hat{n_2}$	$\frac{1}{3}(\sqrt{8}, 0, -1)$	π	0
$\hat{n_3}$	$\frac{1}{3}(-\sqrt{2}, \sqrt{6}, -1)$	1.897	0
$\hat{n_4}$	$\frac{1}{3}(-\sqrt{2}, -\sqrt{6}, -1)$	-1.897	0

Table 6.1: The coordinates of the four tetrahedral directions, along with a set of phases sufficient to make the measurement sets $|\psi_k\rangle$ and $|\psi'_k\rangle$ orthonormal.

is thus important to note that while the spin of electrons and other massive fermions can be imagined to "point" in a specific direction (within the limits of the uncertainty principle), that is not the case for photons [67]. Consequently, photons cannot directly implement orienteering as described above. Neither our implementation, nor any other using photon polarization as the communication medium, can transmit a direction in space without some previously agreed upon reference [68]. Nevertheless, we are still able to demonstrate the operating principles of orienteering using photons, and in particular show explicitly the advantages of collective measurements and anti-parallel encodings of the direction.

A number of factors make a direct photon implementation difficult. The most important of these is that, as mentioned above, the polarization of a photon does not necessarily point in a particular direction. However, Alice and Bob can use the representation of polarization states on the Poincaré sphere to construct a mapping, such as that shown in Fig. 6.1, to convert photon polarization estimates into directions. Obviously, there are an infinite number of equivalent mappings. Bob and Alice's agreement on a particular one is tantamount to possessing a shared reference frame.

The second difficulty is that making the required joint measurements is challenging, equivalent to full Bell-state analysis. While one or two Bell-states may be measured at a time, it is not, in general, possible to perform an efficient, arbitrary, four-outcome measurement of the two-photon polarization state with current technology [1]. However, if Alice and Bob *do* share a reference frame (required for the polarization-to-direction mapping already), they can combine part of the measurement into state creation. Specifically, since an arbitrary joint measurement without ancilla can be thought of as a two-qubit unitary operator $\hat{U}$ followed by a separable measurement $\hat{M}$ [64], we can directly apply the unitary to Alice's source state $|\psi\rangle$, allowing Bob to make a simple separable measurement: $(\hat{M} \cdot \hat{U})|\psi\rangle = \hat{M}(\hat{U}|\psi\rangle)$. This requires Alice to

[1] It is possible to do full Bell state analysis probabilistically using optical C-NOT gates [69, 70, 71], or to do so deterministically if the photons are also entangled in another degree of freedom [72].

precisely create specific partially entangled states such as

$$\begin{aligned}|\psi_{111}\rangle \approx 0.683\,|HH\rangle + (0.463 + 0.358i)|HV\rangle \\ +(-0.208 + 0.383i)|VH\rangle \\ +(-0.011 - 0.025i)|VV\rangle,\end{aligned} \tag{6.4}$$

equivalent to a parallel encoding of the direction $\frac{1}{\sqrt{3}}(1,1,1)$.

Fig. 6.2 shows the system used to create these entangled states and perform the separable measurements on them. Alice uses a tunable source of polarization entanglement based on spontaneous parametric down-conversion (SPDC) in a pair of non-linear crystals (BBO) [73], as described in Chapter 2. These crystals have their optic axes oriented such that when pumped by a UV laser (in our case, an Ar^+ laser at 351.1 nm), the first crystal generates horizontally polarized down-conversion pairs ($|HH\rangle$) while the second crystal generates vertically polarized pairs ($|VV\rangle$). If the coherence length of the pump is long, those processes are indistinguishable, resulting in the entangled state $\cos\theta|HH\rangle + \sin\theta|VV\rangle$. The weight of the two terms—and thus the degree of entanglement—is controlled by the pump polarization; the other 5 parameters in a general two-qubit pure state can be set by wave plates after the crystals [74]. In order to achieve the high state quality needed to be able to resolve the optimal fidelities of the different encodings, we also implemented a phase-compensation technique, described in detail elsewhere [1][2]. In each case, the required states for our experiment were created with a fidelity greater than 98%. (In fact, although transferring the unitary operator onto the state preparation step leads to particular relative phases, as in Eqn. (6.4), Bob's separable measurement is actually insensitive to these, depending only on the *magnitudes*.)

6.4 Results

We measured the transmission fidelity for three different two-photon protocols (Table 6.2; for comparison we also list the measured fidelities if Alice sends only a single photon). First, we implemented the naïve approach where Bob makes separable measurements. The second protocol used the optimal joint measurements for a parallel encoding, while the final case used the optimal measurements for anti-parallel encoding. In each case, we determined the fidelity for a large number of possible directions (see Fig. 6.3). The average orienteering fidelities (Table 6.2) are near the theoretical values for all trials, thereby verifying the principle of using quantum encoding for orienteering. In particular, the experimental results for joint measurements exceed

[2]SPDC generates very pure entangled states, but the specific state created varies over the collection solid angle. By inserting a birefringent element that cancels this direction-dependent state rotation, we generate a more pure source of entangled photons than previously available.

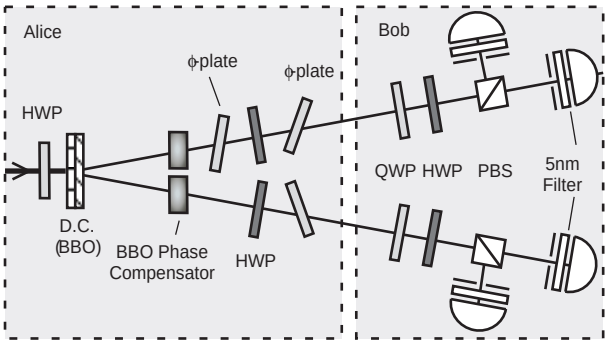

Figure 6.2: Alice creates the necessary entangled states using parametric down-conversion in two BBO crystals. By changing the first half wave plate (HWP) she can create a non-maximally entangled state of the form $\cos\theta|HH\rangle + e^{i\phi}\sin\theta|VV\rangle$. The BBO phase compensators correct for a spatial dependence of ϕ in the initial entangled state, allowing greater state purity [1]. Following these are several wave plates which allow Alice to create an arbitrary pure two-qubit state. The plates marked "ϕ-plate" are wave plates (with their optic axes at 0°) which can be tipped to provide an arbitrary phase (ϕ) between $|H\rangle$ and $|V\rangle$, while the HWPs perform rotation by π about any linear axis on the Poicaré sphere. Bob uses a QWP, HWP, and polarizing beam splitter (PBS) in each arm, enabling an arbitrary projection on each qubit. This allows him to make the separable measurement for orienteering, and also to perform full state tomography [2].

class	single spin	separable	joint parallel	anti-parallel
sphere	66.5 [66.7]	73.2 [73.6]	74.0 [75.0]	78.2 [78.9]
equator	74.8 [75.0]	84.1 [85.3]	74.0 [75.0]	78.4 [78.9]
tetrahedron	69.3 [69.4]	73.6 [73.6]	82.5 [83.3]	94.9 [95.5]

Table 6.2: The average fidelities for each of the four protocols — the single-spin case and three variations on the two-spin protocol. We also show the average when Alice is confined to transmitting a direction on the equatorial plane, or picking one of the four tetrahedral directions. The (statistical) error on each value is ±0.1%; the theoretical limits are shown in [ ].

the upper bound for separable measurements, and the the experimental results for the anti-parallel protocol exceed the theoretical bound for parallel spin protocols.

In addition, we examined cases where Alice is restricted to only send a subset of possible directions. For example, if Alice is constrained to sending states on the equatorial plane (equivalent to transmitting a *phase*), the optimal fidelity of 85.3% [58] is realized using orthogonal separable measurements (also in the plane). The greatest advantage of the anti-parallel encoding is observed when Alice chooses to send one of the four tetrahedral directions, yielding a fidelity of 95.5% [65], much higher than the 83.3% using joint measurements on a parallel encoding.

6.5 Discussion

Our results confirm the possibility of superior orienteering using quantum resources, i.e., that joint measurements on anti-parallel spins outperform separable measurements. It is important to realize, however,

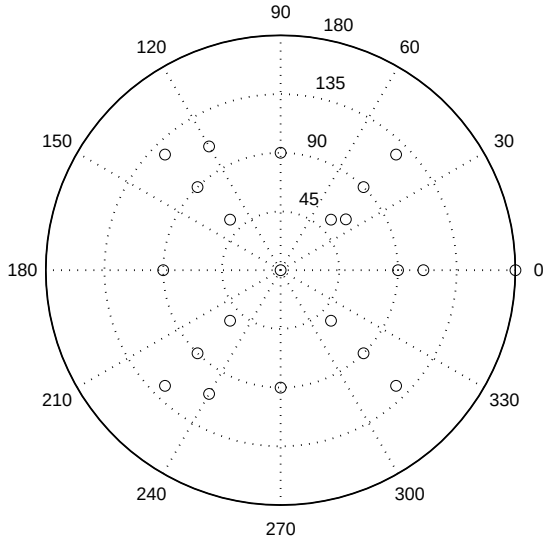

Figure 6.3: For all four protocols, each of the directions indicated on this polar plot of the surface of the Poincaré sphere was encoded. These include the four tetrahedral directions, the cardinal directions: $\pm\{\hat{X}, \hat{Y}, \hat{Z}\}$, four additional states on the equator at $\pm 45°$ and $\pm 135°$, and all eight points at the corner of an inscribed cube: $(\pm 1, \pm 1 \pm 1)/\sqrt{3}$. The center of the plot corresponds to the state $|R\rangle$ while the outer rim represents $|L\rangle$, encoding the states $\pm\hat{z}$, respectively. Linear polarization states lie on the middle circle, with a polar angle ϕ of 90°.

that none of these algorithms is actually optimal when Alice and Bob share a reference frame. In that case, they actually do the best if Alice simply sends a 2-bit *classical* encoding of the direction, which amounts to telling Bob which of the tetrahedral directions from Fig. 6.1 is closest to her desired direction. This gives an average transmission fidelity of 85.3%, significantly higher than the optimal measurement on anti-parallel spins (78.9%).

Quantum orienteering in the absence of a shared reference frame is interesting theoretically since it is a recent addition to a relatively small set of quantum communication protocols. Nevertheless, realizing a practical implementation is likely to be problematic. While a system using electron spins, nuclear spins, or Rydberg atoms (which can also be used for a related protocol, transmitting a reference frame [75]) could avoid the photon-implementation issues discussed above, it introduces others, most notably sensitivity to stray magnetic fields, which would cause precession of the spins. For example, it seems unlikely that two parties could be in a position where it was technologically difficult or impossible to share a reference frame, yet they could control the magnetic fields well enough to communicate using quantum orienteering.

However, this opens up a new avenue for exploration: instead of agreeing on a global idea of a direction, Alice and Bob may wish to *measure* the rotation applied by the background fields. This, for instance, would allow partners with quantum computers to align their computational basis in the same direction, defined so

that when Alice sends a '1', that is what Bob receives, and similarly for '0'. Now, Alice is no longer trying to transmit a physical direction, and so this application may be performed not only with massive spins, but, e.g., with photon polarization. In this instance, two parties connected by a fiber optic cable, which in general performs an arbitrary rotation of the polarization, may identify a basis in which they can communicate most efficiently.

Chapter 7

Conclusions

Currently, only a small but growing number of quantum communication protocols are known, and even fewer have been experimentally implemented. Nevertheless, their existence illuminates some of the fundamental differences between classical physics and quantum mechanics. Some of these protocols, such as quantum key distribution, allow us to accomplish tasks impossible with only classical physics, while others, such as orienteering, show improved behavior when implemented with quantum channels.

One particularly interesting avenue, which has received less attention, is the intersection between quantum physics and special relativity. There are a few known communication protocols relying on special relativity, such as secure coin tossing [76]. Coin tossing allows Alice and Bob to work together to generate a random bit (heads or tails), even when they do not trust each other. It has been proven that no quantum information protocol (or any standard classical protocol) can implement this securely [77, 78] – either Alice or Bob will be able to cheat, swaying the outcome in their favor. However, a very simple relativistic protocol solves this problem neatly: Alice and Bob each publicly declare a random bit, revealing it before they can see what the other has chosen. The coin toss is simply the exclusive-or of Alice and Bob's bits; therefore, either one of them can defeat any cheating strategy by simply picking their bit randomly. The security of this simple protocol rests on neither Alice nor Bob being able to alter their action based on the other.

In our relativistic QKD protocol we have shown that by combining constraints from both relativity and quantum mechanics we can gain an advantage that neither holds alone. The only other protocol we know of to combine the two is secure bit-commitment [79, 80], which has not yet been implemented. The essence of all three of these protocols with a relativistic component is that they force the parties involved to commit to a decision before some in-transit information reaches them, a capability intrinsic to relativity and its notion of space-like separated parties.

Quantum orienteering is a purely quantum protocol, but one that shows the practical inequivalence of different quantum systems. The intent is to transmit a direction in space without referring to an established reference frame. We implemented the protocol using photon polarization, which is formally identical to the originally proposed spin-1/2 particles. However, while the results are the same, the interpretation must be

different: electron spin has a natural mapping to physical space, while photon polarization does not – in practice, communicating a direction with polarization states does require a shared reference frame. Without the shared frame, a linearly polarized photon only naturally identifies an axis, not a unique direction. For instance, if a photon is vertically polarized (in some reference frame) it cannot distinguish between pointing "up" and "down".

The essence of the practical difference between photon polarization and electron spin is that the photon massless. For a massive particle such as a neutron or electron, the spin, and therefore the angular momentum, always points in a definite direction. It is always possible to transform into the rest frame of the particle, and measure its angular momentum. Photons are massless, and therefore have no rest frame. In particular, they are spin-1 particles whose $m_z = 0$ states are forbidden. This prevents them from pointing unambiguously in any direction. Other quantum systems also do not communicate direction, simply because the physical quantity is not directional. For instance, encoding into the state of an atom may or may not encode directional information, depending on the symmetry of the states used. The ground state of an atom is spherically symmetric, so conveys no direction, while some excited states have an electric dipole. It would even be possible to select an atomic qubit where one basis state was rotationally symmetric and the other was not. All of these factors are normally ignored, assuming that "all two level systems are equivalent". Certainly they are mathematically equivalent; however, for some types of communication the specific type of system is important.

These protocols extend the field of quantum communication, still in its early stages. While QKD already shows promise as a practical communication tool, other protocols are only in the development stages. It is likely the true impact of quantum communication will not be realized until it can be coupled with quantum computing. Quantum channels are not merely a convenience, but a requirement for connecting quantum computers. In this, we may find applications for protocols similar to quantum orienteering, used to align the computation basis between two separated computers or registers within a single computer.

Appendix A

Cavity mirror alignment

The key to building a useful optical delay line is selecting the optimal configuration. A cavity constructed using cylindrical mirrors can be described by three parameters, N, m_x, and m_y, as described in Chapter 3. N is the total number of round trips, while m_x and m_y describe the trajectory of the beam. N, m_x, and m_y must be relatively prime, and $m_x < m_y < N$ (the solution is symmetric under exchange of m_x and m_y so we can take either to be smaller). The first step is to select an appropriate set of parameters. From Eqn. (3.6) we find that the eigenvalues of the ray matrix $A(d, \theta)$ are:

$$\phi_1 = \frac{\pi m_x}{N} \quad \text{and} \quad \phi_2 = \frac{\pi m_y}{N}. \tag{A.1}$$

The ray matrix $A(d, \theta)$ is a function of two settings, d, the spacing between the mirrors, and θ, the twist angle. The best way to pick a configuration is to numerically solve the equation

$$eig(A(d, \theta)) = \{\pm\phi_1, \pm\phi_2\} \tag{A.2}$$

for all values of m_x and m_y with the desired N.

To select a specific solution from the many, the most important factor is that the nearest bounce to the central coupling hole be as far from it as possible. To do this, we calculate the entire trajectory of the input ray as it propagates through the cavity. Specifically, we select an appropriate input ray x_0 and calculate the position of each bounce according to the formula

$$x_i = A^i x_0. \tag{A.3}$$

A good first choice for x_0 is just $(0, 1, 0, 1)$. That is,the initial position in x and y is zero, while the slope is 1. For each cavity configuration, find the ratio of the closest approach to the origin to the furthest deviation from the origin. If the scale is set such that the furthest deviation is within a finite bound (the size of the mirror), that allows an exact calculation of the nearest approach. The most stable, lowest-loss configuration

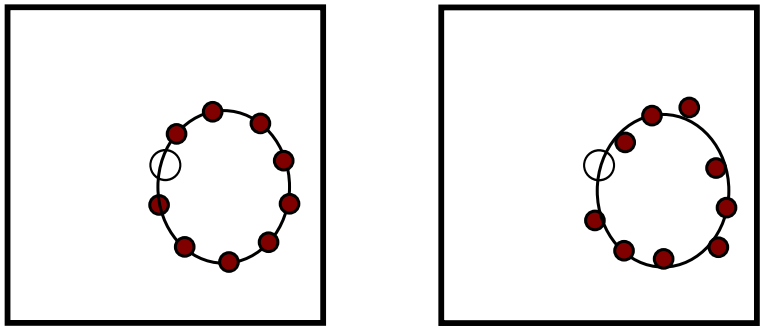

Figure A.1: Example spot patterns for orientations very close to $\theta = 90°$. Left: When the twist angle is exactly 90°, the reflections lie on an ellipse. Right: If the the twist angle is near, but not exactly 90°, the pattern roughly follows the ellipse, but is distorted, as shown.

will be one where the nearest approach is as large as possible.

Once the target configuration is chosen, we must attempt to align the mirrors to match the predicted d and θ. This is a process in several steps. First, we align the two mirrors in a cavity configuration with separation very close to the calculated value and the axes of curvature orthogonal. We then shine an alignment laser through the coupling hole in the first mirror onto the center of the second mirror. We adjust the tip and tilt of the second mirror so that the beam returns directly through the coupling hole. This aligns the second mirror to normal to the cavity axis. We now slightly misalign the second mirror so that the returning beam reflects near the coupling hole instead of passing through it, and adjust the tip and tilt of the first mirror such that it too is normal to the cavity axis.

Once the mirrors are aligned carefully on the cavity axis, we attempt to reach a known setting of twist angle and separation. To do so, we adjust the alignment laser off of the central axis. The beam will be shown to roughly trace out an ellipse. When the axes of the two mirrors are actually orthogonal, the spot pattern on each mirror will indeed be an ellipse. If the axes are slightly off, the elliptical pattern will distort, as shown in Fig. A.1. Adjust the twist angle until this pattern is as closely elliptical as possible. Slight variation in the radii of curvature of the mirrors may prevent perfect alignment.

At this point, we could set the separation exactly by counting the number of reflections on the elliptical spot pattern, and then comparing that to the simulations. However, the twist angle is more difficult to set than the separation. As long as the initial placement is accurate within 1 cm or so, it is usually possible to go directly from the elliptical alignment to the target configuration. We twist the second mirror by the correct number of degrees. This will bring about something very close to the target spot pattern. We then push the translation state back and forth by hand in the longitudinal direction, watching for the desired pattern. Next we adjust the tip and tilt of the cavity end mirrors to adjust the overall spot pattern size, to make sure that the reflections stay safely away from the outside edges of the mirrors.

The final step is mode matching. In order to have high efficiency, we want the beam to be uniformly small when it reflects off of the end mirrors. While we presently do not have detailed calculations for this, it appears that a reasonably good approximation is to gently focus the input beam such that there is a beam waist near the center of the cavity. The periodic focusing behavior of the cavity makes it relatively insensitive to the initial focusing, as long as the input beam is not too large or diverging. The output beam is usually only slightly separated from the input. At a distance of $\sim$1 m from the cavity entrance, the output beam is far enough away to pick off using a small mirror or prism.

These techniques have been sufficient to align delay cavities of up to a few microseconds of delay. With very high-reflectivity mirrors, it should be possible to have efficient storage for 10s of microseconds. Configurations with that long of a delay will be more sensitive to alignment. It is possible these techniques will not be effective in that case. Probably the best way to achieve these more sensitive alignments is to first identify a nearby solution of lower N, then calculate the small adjustment needed to move to the larger N configuration. In this way it should be possible to generate delay lines with thousands of cycles, for $> 10\mu s$ delay time.

Appendix B

Mirror Surface Quality and Cleaning

In order to achieve low loss storage cells, we require very high reflectivity mirrors. The highest reflectivity mirrors available are multi-layer dielectric mirrors coated using ion beam deposition. This process has been used to create resonant cavities with a Q greater than 10^6 and a corresponding mirror reflectivity $R > 99.999\%$ [25]. Mirrors of this quality would potentially allow multi-pass storage cells to have extremely low loss, comparable to that of single-mode fibers. However, the reflectivity of the mirrors used in the storage cavity describe in Chapter 3 are considerably lower, at 99.8%.

The first consideration is the manufactured quality of the mirrors. Ion beam deposition provides very uniform layers; however, the mirror quality is limited by the smoothness of the bare substrate. Highly polished flat or spherical mirrors are widely available, but we require cylindrical surfaces for our delay cavity. These are harder to manufacture and not widely available with such high surface quality. Fig. B.1 shows the surface of a cylindrical mirror manufactured by HellmaOptik and a flat mirror manufactured by MLD Technologies. Both mirrors have a multi-layer coating composed of alternating layers of SiO_2 and Nb_2O_5 applied by MLD technologies. The Hellma cylindrical mirror has a number of features which will cause scattering. Primarily, the surface is covered with 7 nm "peaks" which are likely either nano-particles that stuck to the surface and were coated over, or peaks that are part of the substrate. These features dominate the 1.2 nm RMS roughness. In the area between these peaks, there is also considerable variation in height, indicating that even if those bumps are removed, the surface would be rougher than the flat MLD substrate.

Surface roughness causes scatter, which appears to be the primary loss mechanism for our storage cavity mirrors (we measure the transmission through the mirror to be less than 0.01%, and it is unlikely that there is significant absorption in the dielectric coating). Scattering from a dielectric surface is a complex process that depends on the exact structure of the surface defects involved. However, if the surface is characterized by Gaussian noise with an RMS roughness much smaller than the wavelength, the integrated scatter over

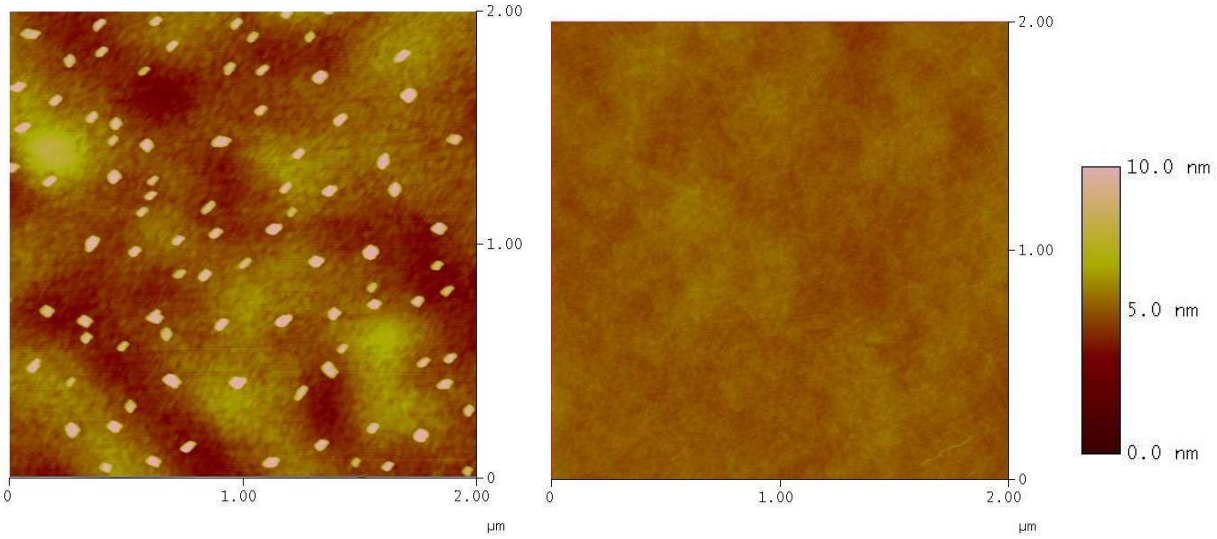

Figure B.1: A: AFM scan of cylindrical cavity mirror coated with ion beam deposition B: Flat mirror with the same coating. The z-range on each scan is 10 nm. The flat mirror has an RMS roughness of 0.21 nm, while the cylindrical mirror roughness is 1.2 nm. In addition, the cylindrical mirror surface is dominated by 7 nm nanoparticles.

all angles (total integrated scatter, or TIS) is approximated by [81]

$$\mathrm{TIS} = 1 - e^{-(4\pi\delta\cos(\theta_0)/\lambda)^2}, \tag{B.1}$$

$$\delta = \frac{\lambda}{4\pi\cos(\theta_0)}\sqrt{TIS}. \tag{B.2}$$

Here, λ is the wavelength of light (670 nm), θ_0 is the angle of incidence (0°), and δ is the RMS roughness. To explain the 0.2% loss we see requires a surface roughness of 2.4 nm. This is greater than the roughness measured from the AFM scans of 1.2 nm, however, much less that the 7 nm size of the nano particles on the surface. Since the surface structure shown in Fig. B.1 is definitely not Gaussian, the model given here may not directly apply. More investigation would be required to determine whether this scattering model is insufficient or there is some other form of loss. Regardless, this roughness is of the correct order for the observed loss.

Additionally, the reflectivity degraded somewhat with time. This indicates that the surfaces became contaminated. We investigated a number of cleaning techniques to attempt to restore the surface to its original quality. All cleaning processes require manipulating the surface of the mirror, and therefore have some possibility to cause damage. Therefore, we attempted to proceed from perceived less-invasive techniques to more vigorous cleaning. In addition to potentially avoiding unneeded steps, the early cleaning processes remove large dust particles which could scratch the surface if dragged along it later. Below is the order

we arrived at. Ultimately, while none of them caused permanent surface damage, most of them provided no discernible benefit, consistent with the belief that the most prominent problem was the inherent surface quality.

The first line of cleaning is non-contact dusting, using compressed nitrogen. The nitrogen is clean and dry, so avoids leaving residue on the surface, or causing scratches from spraying dust onto the mirror. This technique removed a few visible dust particles, but did not noticeably affect the reflectivity. Presumably, the dust did not cover a significant fraction of the surface area. The second technique is a solvent rinse. We mount the mirrors at an angle, and spray organic solvent along the top edge, allowing it to flow over the surface of the mirror, then use compressed nitrogen to blow the remaining liquid off the surface to minimize the solvent residue. We used acetone first, which we expect to remove most any hydrocarbons on the surface, but will itself leave a residue, then followed with either 2-propanol or methanol. These solvents are intended to remove the acetone residue. We again saw no noticeable improvement after this cleaning. The next technique is drop-and-drag with lens tissue and solvents. It is important that the surface be free from large dust particles which could scratch the surface during this procedure, so should only be used after at least dusting the surface with compressed air or nitrogen. Here, a piece of lens tissue is placed on the surface of the mirror, then wet with a solvent and dragged off the surface. Friction from the lens tissue should remove contaminants more tightly adhered to the surface. Again, to avoid solvent residue, this should be done either with methanol only, or acetone followed by methanol (or 2-propanol). The next technique is OptiClean. This is an acetone-based gel which is applied to the surface. It then dries into a plastic film, which is peeled off, removing contaminants with it. It is supposed to both dissolve organic contaminants and trap particulates in the film, while reducing the risk of dragging particles across the surface. Of all the procedures we tried, this is the only one which noticeably improved the mirror reflectivity, increasing it from 99.7% to 99.8%, at or near the initial level. At the suggestion of MLD, we tried two more procedures. First, we soaked the mirrors in warm DI water with Alconox soap for several hours, followed by DI water and solvent cleaning to remove the soap residue. Next, we baked the substrate in a low-pressure argon environment, intended to volatilize any organics on the surface, while preventing oxidation. Again, neither of these procedures improved the reflectivity, but neither did they appear to damage the surface.

We believe this cleaning regimen should remove contaminants with as little chance of damaging the surface as possible. While only the OptiClean procedure yielded a noticeable improvement in our case, this is probably due to the mirror surfaces starting relatively clean. They are mounted vertically, which reduced the amount of dust, and haven't been exposed to many likely sources of contamination. The storage cavity is mounted inside a Plexiglas enclosure which reduces wind and dust, though may increase the organic

contaminants, especially if the glue used to assemble the enclosure was not cured completely.

It should be possible to construct a delay line with low loss and long (several microseconds) delay. The primary barrier seems to be the availability of sufficiently high quality cylindrical substrates. It is difficult to tell if this is an intrinsic problem, or simply a lack of demand for high quality cylindrical optics.

Appendix C

Low Latency Communication

Low-latency communication is a key component of relativistic quantum key distribution, as described in Chapter 5. Communication latency directly translates into increased storage requirements with the attending loss. Most common communications protocols (such as Ethernet) trade latency for bandwidth. By bundling data into large blocks, they decrease the overhead associate with each transfer at the expense of increased latency. For our application, this is the opposite of what we want. A single basis selection only requires 2 or 3 distinct messages, but the latency is critical.

We built a suitable low-latency communication system using a Xilinx 9500 series CPLD (complex programmable logic device). This device is an array of logic gates connected to flash memory that allows programmable interconnects between the gates, enabling a sort of virtual bread-board for logic systems. More sophisticated devices known as FPGAs (field programmable gate arrays) typically provide more features (such as memory and special purpose connections) and operate faster. However, the CPLD is simple to use, and has high-current outputs which allow it to easily drive other devices. In particular, we use the output directly to drive a laser diode.

We programmed the CPLDs using VHDL (VHISC Description Language), a programming language widely used for this purpose. This language allowed us to define a state-machine. A finite state machine (FSM) is a computing device that can be in one of a finite number of internal states, stored in a small amount of memory. The FSM receives a clock input, and on each clock cycle moves to a new state in a way that depends on the state of the inputs, as well as the current state of the machine. Each output is then defined by whether it is "high" (i.e., logical 1, +3.3 V), or "low" (i.e., logical 0, 0 V) for each state. A simplified version of the modulator state machine is shown in Fig. C.1.

We used a 5-bit encoding for the encoded messages. While the three bases could be communicated with only 2 bits, it is convenient to use 5 for a number of reasons. The first consideration is to give each code a start bit of 1. Since a photon detection can come at any time, it is needed to instruct the demodulator to begin decoding. We also disallow 10000 as a code, since a single glitch will be recognized as a valid signal. The second consideration is that we use a discriminator to recover the signal from a photo-diode.

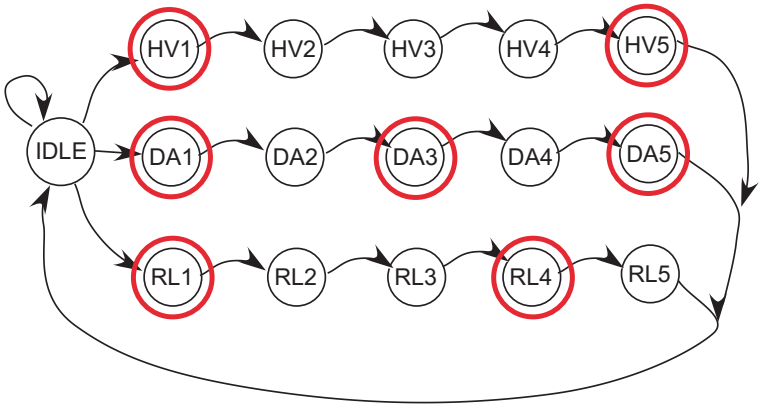

Figure C.1: Simplified finite state machine (FSM) for the modulator. The default state is idle. If one of the quantum basis inputs is triggered, the state machine follows the transition onto one of the three paths. Each clock cycle then moves the FSM down the path until it reaches the fifth state. The FSM then returns to the idle condition. The states marked with a red circle are those in which the laser output is on. This sets the output code for each basis.

The discriminator responds to a rising edge with a pulse of a single clock cycle. Thus, we do not allow two sequential 1s.

The core of the modulator has 16 states. The initial state is the quiescent state, in which no signal has been seen recently. If exactly 1 of the 6 input signals is triggered (each corresponding to one of the 6 possible polarizations in the six-state QKD protocol), the device transfers into one of 3 states (each corresponding to one of the bases used in the six-state protocol). From there, each clock cycle causes the device to move along a linear chain of states. Each machine state of each chain corresponds to a single output bit for the selected basis. Upon reaching the end of a chain, the device returns to the idle state, and awaits another signal.

This design is functional, but we have made a number of enhancements to this simple 16-state machine. The first is to allow a switchable output delay. We added a bank of 3 switches that allow the output to be delayed between 0 and 7 clock cycles. This is useful because it allows us to ensure that the classical signal is not transmitted until after the quantum signal (the precise security requirements are discussed in Chapter 5). We also add a programmed dead-time. The modulator will not respond to a photon click for a few cycles after it finishes transmitting a signal. This ensures that the demodulator is completely reset by the time the next signal begins. Finally, the modulator sends a signal to Alice's computer indicating which of the six states triggered the modulator. Because Alice's IO card only samples at 10 MHz, the demodulator holds this signal for several clock cycles to make sure she sees the signal.

The demodulator operates similarly to the modulator. The internal state of the demodulator is captured in a shift register that holds the last 5 bits received. Each clock cycle, the 4 most recent bits are shifted down

one, while the input sampled is loaded into the first bit of the register. When the entire register corresponds to one of the magic key sequences, several outputs are triggered. One set of outputs triggers the Pockel's cell to set the analysis basis for the DA and RL bases. A second set of outputs goes to Bob's computer, and in the same fashion as with the modulator, must hold those signals for many clock cycles. After all of the signals are activated, the demodulator resets itself and waits until the shift register again recognizes one of the basis selection codes.

The demodulator must be able to operate with a clock that is not phase-locked to the modulator. The clock rising edge can come at any time during a bit, and the relative phase between the clock and signal may drift. One way to solve this is to use clock-recovery, a technique that extracts the clock directly from the data channel. This ensures that the clock is always synchronized to the data stream, but is not easy to implement in our system. Instead, we implement two independent demodulators on the same chip. The shift register for the second demodulator is clocked on the *falling* edges of the clock signal. This makes it sample the input 180° out of phase with the normal rising edge demodulator. As long as the clocks are close in frequency, one of the two demodulators will correctly recognize every signal. If the two demodulators generate the same output, or only one recognizes a valid message, the message is output. If the two demodulators disagree, there is no way to determine which is correct, and the message is ignored. This is important because basis selection errors can compromise the security of the quantum key distribution protocol, but dropped messages do not.

References

[1] Joseph B. Altepeter, Evan R. Jeffrey, and Paul G. Kwiat. Phase compensated ultra-bright source of entangled photons. *Optics Express*, 13(22), 2005.

[2] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White. Measurement of qubits. *Phys. Rev. A*, 64(5):052312, November 2001.

[3] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. doi: 10.1137/S0097539796300921. URL `http://link.aip.org/link/?SMJ/26/1411/1`.

[4] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *ArXiv Quantum Physics e-prints*, August 1995.

[5] A. K. Lenstra and H. W. Lenstra Jr. *The development of the number field sieve*. Springer, 1993.

[6] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC '96: Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, New York, NY, USA, 1996. ACM Press. ISBN 0-89791-785-5. doi: http://doi.acm.org/10.1145/237814.237866.

[7] Adrian Kent. Quantum bit string commitment. *Phys. Rev. Lett.*, 90(23):237901, Jun 2003. doi: 10.1103/PhysRevLett.90.237901.

[8] Matthias Fitzi, Nicolas Gisin, and Ueli Maurer. Quantum solution to the byzantine agreement problem. *Phys. Rev. Lett.*, 87(21):217901, Nov 2001. doi: 10.1103/PhysRevLett.87.217901.

[9] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 87(16):167902, Sep 2001. doi: 10.1103/PhysRevLett.87.167902.

[10] Andrew Chi-Chih Yao. On the power of quantum fingerprinting. In *STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 77–81, New York, NY, USA, 2003. ACM Press. ISBN 1-58113-674-9. doi: http://doi.acm.org/10.1145/780542.780554.

[11] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Royal Society of London Philosophical Transactions Series A*, 356:1733, August 1998.

[12] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller. Quantum repeaters: The role of imperfect local operations in quantum communication. *Phys. Rev. Lett.*, 81(26):5932–5935, Dec 1998. doi: 10.1103/PhysRevLett.81.5932.

[13] Zhi Zhao, Tao Yang, Yu-Ao Chen, An-Ning Zhang, and Jian-Wei Pan. Experimental realization of entanglement concentration and a quantum repeater. *Phys. Rev. Lett.*, 90(20):207901, May 2003. doi: 10.1103/PhysRevLett.90.207901.

[14] Julio T. Barreiro, Nathan K. Langford, Nicholas A. Peters, and Paul G. Kwiat. Generation of hyperentangled photon pairs. Phys. Rev. Lett., 95(26):260501, 2005. doi: 10.1103/PhysRevLett.95.260501. URL `http://link.aps.org/abstract/PRL/v95/e260501`.

[15] Carl A. Kocher and Eugene D. Commins. Polarization correlation of photons emitted in an atomic cascade. *Phys. Rev. Lett.*, 18(15):575–577, Apr 1967. doi: 10.1103/PhysRevLett.18.575.

[16] Alain Aspect, Philippe Grangier, and Gérard Roger. Experimental tests of realistic local theories via bell's theorem. *Phys. Rev. Lett.*, 47(7):460–463, Aug 1981. doi: 10.1103/PhysRevLett.47.460.

[17] Ryan S. Bennink, Yun Liu, D. Duncan Earl, and Warren P. Grice. Spatial distinguishability of photons produced by spontaneous parametric down-conversion with a focused pump. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 74(2):023802, 2006. doi: 10.1103/PhysRevA.74.023802. URL `http://link.aps.org/abstract/PRA/v74/e023802`.

[18] Akira Tomita and Raymond Y. Chiao. Observation of berry's topological phase by use of an optical fiber. *Phys. Rev. Lett.*, 57(8):937–940, Aug 1986. doi: 10.1103/PhysRevLett.57.937.

[19] John S. Bell. On the problem of hidden variables in quantum mechanics. *Rev. Mod. Phys.*, 38(3): 447–452, Jul 1966. doi: 10.1103/RevModPhys.38.447.

[20] Stuart J. Freedman and John F. Clauser. Experimental test of local hidden-variable theories. *Phys. Rev. Lett.*, 28(14):938–941, Apr 1972. doi: 10.1103/PhysRevLett.28.938.

[21] Joseph B. Altepeter, Evan R. Jeffrey, and Paul G. Kwiat. *Advances in AMO Physics, 2006, Chapter 3: Photonic State Tomography*. Elsevier, Munich, Germany, 2006.

[22] A. E. Siegman. *Lasers*. University Science Books, Mill Valley, CA, 1986.

[23] T. B. Pittman, B. C. Jacobs, and J. D. Franson. Demonstration of quantum error correction using linear optics. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 71(5):052332, 2005. doi: 10.1103/PhysRevA.71.052332. URL `http://link.aps.org/abstract/PRA/v71/e052332`.

[24] Peter Fritschel. Second generation instruments for the laser interferometer gravitational wave observatory (ligo). *ArXiv Quantum Physics e-prints*, August 2003.

[25] Christina J. Hood, H. J. Kimble, and Jun Ye. Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity. *Phys. Rev. A*, 64(3):033804, Aug 2001. doi: 10.1103/PhysRevA.64.033804.

[26] J. M. Bennett and D. Rönnow. Test of Opticlean Strip Coating Material For Removing Surface Contamination. Appl. Opt., 39:2737–2739, June 2000.

[27] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, UK, 1992.

[28] B. Jun and P. Kocher. The intel random number generator. *white paper prepared for Intel Corp*, 1999.

[29] Evaluation of via c3 nehemiah random number generator. `http://www.via.com.tw/en/downloads/whitepapers/initiatives/padlock/evaluation_padlock_rng.pdf`, 2003.

[30] Zvi Gutterman, Benny Pinkas, and Tzachy Reinman. Analysis of the Linux random number generator. In *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy (S&P'06)*, pages 371–385, Washington, DC, USA, 2006. IEEE Computer Society. ISBN 0-7695-2574-1. doi: http://dx.doi.org/10.1109/SP.2006.5.

[31] M. Isida and H. Ikeda. Ann. inst. stat. math. Tokyo 8, 119 (1956).

[32] A. Stefanov, N. Gisin, O. Guinnard, L. Guinnard, and H. Zbinden. Letter Optical quantum random number generator. *Journal of Modern Optics*, 47:595–598, April 2000.

[33] T. Jennewein, U. Achleitner, G. Weihs, H. Weinfurter, and A. Zeilinger. A fast and compact quantum random number generator. *Review of Scientific Instruments*, 71:1675–1680, April 2000. doi: 10.1063/1.1150518.

[34] A. Spinelli, M.A. Ghioni, S.D. Cova, and L.M. Davis. Avalanche detector with ultraclean response for time-resolved photon counting. *IEEE Journal of Quantum Electronics*, 34:817–821, 1998.

[35] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[36] A F Webster and S E Tavares. On the design of s-boxes. In *Lecture notes in computer sciences; 218 on Advances in cryptology—CRYPTO 85*, pages 523–534, New York, NY, USA, 1986. Springer-Verlag New York, Inc. ISBN 0-387-16463-4.

[37] Fips 180-2 secure hash standard. `http://csrc.nist.gov/publications/fips/fips180-2/fips180-2withchangenotice.pdf`.

[38] Xiaoyun Wang, Hongbo Yu, and Yiqun Lisa Yin. Efficient collision search attacks on sha-0, 2005.

[39] The marsaglia random number cdrom including the diehard battery of tests of randomness. `http://www.csis.hku.hk/ diehard/cdrom/`, 1995.

[40] Fips 140-2: Security requirements for cryptographic modules. `http://csrc.nist.gov/cryptval/140-2.htm`, 2001.

[41] Ueli M. Maurer. A universal statistical test for random bit generators. *J. Cryptol.*, 5(2):89–105, 1992. ISSN 0933-2790.

[42] W. K. Wooters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802, 1982.

[43] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proc. of the IEEE Int. Conf. on Computers, Systems, and Signal Processing, Bangalore, India*, Proc. IEEE, page 175. IEEE, New York, 1984.

[44] J Daemen and V Rijmen. *The Design of Rijndael: AES–the Advanced Encryption Standard*. Springer, 2002.

[45] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978. ISSN 0001-0782. doi: http://doi.acm.org/10.1145/359340.359342.

[46] D. G. Enzer, Phillip G. Hadley, Richard J Hughes, Charles G Peterson, and Paul G Kwiat. Entangled-photon six-state quantum cryptography. *New Journal of Physics*, 4:45.1, 2002.

[47] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Phys. Rev. Lett*, 77:2818, 1996.

[48] G Brassard and L Salvail. Secret-key reconciliation by public discussion. *Lecture Notes in Computer Science*, (765):410, 1994.

[49] Christopher A. Fuchs, Nicolas Gisin, Robert B. Griffiths, Chi-Sheng Niu, and Asher Peres. Optimal eavesdropping in quantum cryptography. i. information bound and optimal strategy. *Phys. Rev. A*, 56 (2):1163–1172, Aug 1997. doi: 10.1103/PhysRevA.56.1163.

[50] D. Bruss. Optimal eavesdropping in quantum cryptography with six states. *Phys. Rev. Lett.*, 81: 3018–3021, 1998.

[51] Hoi-Kwong Lo, H.F. Chau, and M. Ardehali. Efficient quantum key distribution scheme and a proof of its unconditional security. *Journal of Cryptology*, 18:133–165, 2005.

[52] N. A. Peters, J. T. Barreiro, M. E. Goggin, T.-C. Wei, and P. G. Kwiat. Remote State Preparation: Arbitrary Remote Control of Photon Polarization. *Physical Review Letters*, 94(15):150502, April 2005. doi: 10.1103/PhysRevLett.94.150502.

[53] H.-K. Lo, X. Ma, and K. Chen. Decoy State Quantum Key Distribution. *Physical Review Letters*, 94 (23):230504, June 2005. doi: 10.1103/PhysRevLett.94.230504.

[54] Christopher A. Fuchs, Nicolas Gisin, Robert B. Griffiths, Chi-Sheng Niu, and Asher Peres. Optimal eavesdropping in quantum cryptography. i. information bound and optimal strategy. *Phys. Rev. A*, 56 (2):1163–1172, Aug 1997. doi: 10.1103/PhysRevA.56.1163.

[55] A. Peres and W. K. Wootters. Optimal detection of quantum information. *Phys. Rev. Lett.*, 66:1119–1122, March 1991. doi: 10.1103/PhysRevLett.66.1119.

[56] S. Massar and S. Popescu. Optimal Extraction of Information from Finite Quantum Ensembles. *Phys. Rev. Lett.*, 74:1259–1263, February 1995. doi: 10.1103/PhysRevLett.74.1259.

[57] N. Gisin and S. Popescu. Spin Flips and Quantum Information for Antiparallel Spins. *Physical Review Letters*, 83:432–435, July 1999. doi: 10.1103/PhysRevLett.83.432.

[58] R. Derka, V. Buzek, and A. K. Ekert. Universal Algorithm for Optimal Estimation of Quantum States from Finite Ensembles via Realizable Generalized Measurement. *Physical Review Letters*, 80:1571–1575, February 1998. doi: 10.1103/PhysRevLett.80.1571.

[59] E. Bagan, M. Baig, A. Brey, R. Muñoz-Tapia, and R. Tarrach. Optimal Strategies for Sending Information through A Quantum Channel. *Phys. Rev. Lett.*, 85:5230–5233, December 2000. doi: 10.1103/PhysRevLett.85.5230.

[60] R. B. M Clarke, V. M. Kendon, A. Chefles, S. M. Barnett, E. Riis, and M. Sasaki. Experimental realization of optimal detection strategies for overcomplete states. *Phys. Rev. A*, 64(1):012303, July 2001.

[61] B. Huttner, A. Muller, J. D. Gautier, H. Zbinden, and N. Gisin. Unambiguous quantum measurement of nonorthogonal states. *Phys. Rev. A*, 54:3783–3789, November 1996.

[62] M. Mohseni, A. M. Steinberg, and J. A. Bergou. Optical Realization of Optimal Unambiguous Discrimination for Pure and Mixed Quantum States. *Physical Review Letters*, 93(20):200403, November 2004. doi: 10.1103/PhysRevLett.93.200403.

[63] G. J. Pryde, J. L. O'Brien, A. G. White, and S. D. Bartlett. Demonstrating Superior Discrimination of Locally Prepared States Using Nonlocal Measurements. *Phys. Rev. Lett.*, 94:220406, 2005.

[64] Terry Rudolph. private communication.

[65] H. Bechmann-Pasquinucci and N. Gisin. Incoherent and coherent eavesdropping in the 6-state protocol of quantum cryptography. *Phys. Rev. A*, 59:4238, 1999.

[66] S. L. Braunstein, S. Ghosh, and S. Severini. Estimation of pure qubits on circles. December 2004.

[67] Terry Rudolph. Quantum Information is physical too... April 1999.

[68] A. Peres and P. F. Scudo. Unspeakable quantum information. January 2002.

[69] T. B. Pittman, B. C. Jacobs, and J. D. Franson. Demonstration of nondeterministic quantum logic operations using linear optical elements. *Phys. Rev. Lett.*, 88:257902, 2002.

[70] Jeremy L. O'Brien, Geoffrey J. Pryde, Andrew G. White, Timothy C. Ralph, and David Branning. Demonstration of an all-optical quantum controlled-not gate. *Nature*, 426:264, 2003.

[71] Sara Gasparoni, Jian-Wei Pan, Philip Walther, Terry Rudolph, and Anton Zeilinger. Realization of a photonic controlled-not gate sufficient for quantum computation. *Phys. Rev. Lett.*, 93:020504, 2004.

[72] P. G. Kwiat and H. Weinfurter. Embedded Bell-state analysis. *Phys. Rev. A*, 58:R2623, October 1998.

[73] P. G. Kwiat, E. Waks, A. G. White, I. Appelbaum, and P. H. Eberhard. Ultrabright source of polarization-entangled photons. *Phys. Rev. A*, 60:R773, August 1999.

[74] Tzu-Chei Wei, Joseph B. Altepeter, David Branning, Paul M. Goldbart, D. F. V. James, Evan Jeffrey, Paul G. Kwiat, Swagatam Mukhopadhyay, and Nicholas A. Peters. Synthesizing arbitrary two-photon polarization mixed states. *Phys. Rev. A*, 71:032329, 2005.

[75] N. H. Lindner, A. Peres, and D. R. Terno. Elliptic Rydberg states as direction indicators. May 2003.

[76] Adrian Kent. Coin tossing is strictly weaker than bit commitment. *Phys. Rev. Lett.*, 83(25):5382–5384, Dec 1999. doi: 10.1103/PhysRevLett.83.5382.

[77] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Phys. Rev. Lett.*, 78 (17):3414–3417, Apr 1997. doi: 10.1103/PhysRevLett.78.3414.

[78] H.-K. Lo and H. F. Chau. Why quantum bit commitment and ideal quantum coin tossing are impossible. *ArXiv Quantum Physics e-prints*, May 1996.

[79] Adrian Kent. Unconditionally secure bit commitment. *Phys. Rev. Lett.*, 83:1447–1450, 1999.

[80] Adrian Kent. Secure classical bit commitment over finite channels.

[81] Karl H. Guenther, Peter G. Wierer, and Jean M. Bennett. Surface roughness measurements of low-scatter mirrors and roughness standards. *Applied Optics*, 23(21):3820, 1984.

Curriculum Vitae

Evan Robert Jeffrey

University of Illinois at Urbana-Champaign		Phone:	+1-217-356-5803
1110 West Green Street		Email:	ejeffrey@uiuc.edu
Urbana, IL 61801 USA		Citizenship:	United States

Education

Doctor of Philosophy in Physics *August 2000 - October 2007 (expected)*

University of Illinois at Urbana-Champaign

Bachelor of Arts in Physics *August 1996 - May 1999*

Washington University

Research Experience

University of Illinois *Urbana, IL* *May 2001 - Present*

Optical Quantum Information Group

Advisor: Dr. Paul Kwiat, *Position:* Graduate Student

Research Foci: Quantum communication, including quantum key distribution and quantum orienteering. Creation and manipulation of entanglement, including entangled photon sources, quantum state tomography, and quantum process tomography.

Washington University *St. Louis, MO* *May 1997 - August 1997*

Cardiovascular Biophysics Group

Advisor: Dr. Sandor Kovacs, *Position:* Howard Hughes Undergraduate Fellow

Research Foci: Mechanical models of the heart used to understand blood flow during ventricular filling. Application as diagnostic indicators of heart malfunction.

Work Experience

Engineering Animation Incorporated *Ames, IA* *June 1999 - August 2000*

Position: Senior Software Engineer

Responsibilities: Developed Internet-based visual collaboration environment for CAD data and engineering documents. Our tools securely and reliably facilitated greater manufacturer / supply chain communication within the automotive and aerospace industries.

Iowa Thin Film Technologies *Ames, IA* *June 1998 - August 1998*

Position: Student Employee

Responsibilities: Built equipment and software for roll-to-roll manufacturing of thin film solar cells, as well as test and measurement instruments for process control.

Iowa State University *Ames, IA* *June 1996 - August 1996*

Position: Student Employee, Microelectronics Research Center

Responsibilities: Developed climate control automation software for for zero-G hydroponics experiments.

Teaching Experience

General Physics, TA (Thermal Physics & Quantum Physics) *UIUC, Spring 2001*

General Physics, TA (Mechanics) *UIUC, Fall 2000*

Object Oriented Software Design, TA *Wash. U., Fall 1998, Spring 1999*

Awards and Honors

3rd Place Team: Charlie Townes *Amazing Light* Young Researcher's Competition *2005*

UIUC Excellence in Teaching Award *2001*

UIUC Excellence in Teaching Award *2000*

Howard Hughes Undergraduate Fellowship *1997*

Publications

E. R. Jeffrey, J. B. Altepeter, and P. G. Kwiat, "Optical Implementation of Quantum Orienteering", *Physical Review Letters* **9**6, 150503 (2006).

J. B. Altepeter, E. R. Jeffrey, and P. G. Kwiat, "Phase-compensated ultra-bright source of entangled photons", *Optics Express* **1**3, 22 (2005).

J. B. Altepeter, E. R. Jeffrey, P. G. Kwiat, S. Tanzilli, N. Gisin, and A. Acin, "Experimental Methods for Detecting Entanglement", *Physical Review Letters* **9**5, 033601 (2005).

J. B. Altepeter, E. R. Jeffrey, and P. G. Kwiat, "Photonic Qubit Tomography", Chapter, *Advances in Atomic, Molecular, and Optical Physics, Vol. 52* , P. Berman, ed. (Elsevier 2005).

E. R. Jeffrey and J. B. Altepeter, "Quantum Tomography" (Matlab code) `http://research.physics.uiuc.edu/QI/Photonics/Tomography/`.

T. C. Wei, J. B. Altepeter, D. Branning, P. M. Goldbart, D. F. V. James, E. Jeffrey, P. G. Kwiat, S. Mukhopadhyay, and N. A. Peters, "Synthesizing arbitrary two-photon polarization mixed states", *Physical Review A* **7**1, 032329 (2005).

N. A. Peters, J. B. Altepeter, D. Branning, E. R. Jeffrey, T. C. Wei, and P. G. Kwiat, "Maximally Entangled Mixed States: Creation and Concentration", *Physical Review Letters* **9**2, 133601 (2004).

E. R. Jeffrey, M. W. Brenner, and P. G. Kwiat, "Delayed-choice quantum cryptography" *Proc. SPIE* 5161 (2004).

P. G. Kwiat, J. B. Altepeter, J. Barreiro, D. Branning, E. Jeffrey, N. A. Peters, and A. VanDevender, "Optical Technologies for Quantum Information Science", *Proc. SPIE* 5161 (2004).

E. Jeffrey, N. Peters, and P. G. Kwiat, "Towards a periodic deterministic source of arbitrary single-photon states", *New Journal of Physics* **6**, 100 (2004).

J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, Paul Kwiat, R. T. Thew, J. L. O'Brien, M. A. Nielsen, and A. G. White, "Ancilla-Assisted Quantum Process Tomography", *Physical Review Letters* **90**, 193601 (2003).

N. A. Peters, J. B. Altepeter, E. Jeffrey, D. Branning, and P. G. Kwiat, "Precise Creation, Characterization, and Manipulation of Single Optical Qubits", *Quantum Information and Computation* **3**, 503 (2003).

P. G. Kwiat, J. B. Altepeter, D. Branning, E. Jeffrey, N. Peters, and T. C. Wei, "Taming Entanglement", *Proceedings of the 6th International Conference on Quantum Communication, Measurement, and Computing* (QCMC '02), J. H. Shapiro and O. Hirota, eds. Rinton Press, 117 (2003).