Architectures for photon-mediated quantum information processing

by

Mihir Pant

B.Eng., Nanyang Technological University (2011)
S.M., Massachusetts Institute of Technology (2015)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2018

Author ..... Signature redacted .....
Department of Electrical Engineering and Computer Science
January 31, 2018

Certified by ..... Signature redacted .....
Dirk R. Englund
Associate Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by ..... Signature redacted .....
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Architectures for photon-mediated quantum information processing

by

Mihir Pant

Submitted to the Department of Electrical Engineering and Computer Science
on January 31, 2018, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this thesis, I present architectures for quantum information processing where photons are used as the quantum bit (qubit) or for mediating entanglement between other qubits. The emphasis of this research is to simplify the basic building blocks required in such processors. The all-photonic repeater and computing architectures do not require material nonlinearities, and their resource requirements are reduced by several orders of magnitude. The photon-mediated atomic memory architecture is designed to work with faulty memories and experimentally demonstrated values of coherence time and photonic coupling efficiency. In the quantum network architecture, the only operation at every node is probabilistic Bell measurement.

Thesis Supervisor: Dirk R. Englund
Title: Associate Professor of Electrical Engineering and Computer Science

Acknowledgments

I would like to thank Dirk Englund for being an amazing supervisor. I am extremely grateful for the academic freedom that I had while working in his group. My interests during the last four years have changed significantly, but Dirk has always been supportive. At the same time, he always had the time to guide me, and kept my theoretical work relevant to experimentalists. One conversation in particular will always stay with me. At the end of a conference, he came to me and said "Let's figure out what the requirements for making a useful quantum system are. Let's just get the numbers, regardless of how challenging they may be". As I was to find out (see Chapter 2), the requirements are indeed very daunting, but going through that calculation actually helped us find a path forward (covered in later chapters).

Saikat Guha had a major impact on my PhD. Saikat's guidance allowed me to grow as a theorist while being in an experimental group, and almost all my papers have involved collaborations with him. I have seen myself mature as a scientist and become more careful about stating my ideas precisely during my time with him. My numerous conversations with him exposed me to ideas ranging from percolation theory, information theory to classical error correcting codes. His breadth of work which spans across classical and quantum domains has been a major inspiration for me, and something I hope to emulate in the future.

I would also like to thank my two other thesis committee members, Jeff Shapiro and Terry Rudolph. Jeff Shapiro's course on quantum optical communication was very important to me, and I will continue to refer to his notes in my future research. I also loved his questions at conferences which were extremely effective at getting to the heart of the matter. My initial acquaintance with Terry was from his papers which have been central to linear optical quantum computing. During the last two years, I was able to talk to him through the MIT-Imperial College London Seed Fund program, which inspired a lot of the work in this thesis and shaped my view of quantum computing in general.

I was a visiting scientist at at Raytheon BBN Technologies for two years during my

PhD, and also did two summer internships there under Saikat Guha and Mohammad Soltani. In addition, I had the opportunity to interact with a number of scientists including Hari Krovi, Prithwish Basu, Zachary Dutton, Boulat Bash, Jonathan Habif, Kin Chung Fong and the rest of the quantum information processing group.

There are a number of people at MIT that I was able to collaborate with including Hyeongrak Choi, Jacques Carolan, Mikkel Heuck, Nicholas Harris, Darius Bunandar, Christopher Panuski, Greg Steinbrecher, Jacob Mower, Catherine Lee, Ryan Hamerly and Mihika Prabhu. Furthermore, my labmates at the Quantum Photonics lab have been important to me during my PhD.

I had the privilege of working with Don Towsley from UMass Amherst and his student Arman Mohseni Kabir on problems in percolation theory. It was also great to work with Liang Jiang and Sreraman Muralidharan from Yale on repeater problems. From Yale, I also enjoyed Leandros Tassiulas' insights on network theory.

My work has also been influenced by conversations with Mercedes Gimeno-Segovia, Pete Shadbolt, Naomi Nickerson, Mihai Vidrighin, Mark Thompson, Jeremy O'Brien, Josh Silverstone, Karl Berggren, Faraz Najafi, Ananthram Swami, Paul Yu, Stefano Pirandola, Koji Azuma, Hoi-Kwong Lo, Zheshen Zhang, Murphy Niu, Franco Wong, Ben Dixon, Scott Aaronson, Alex Arkhipov, Gian Guerreschi, Aram Harrow, Saeed Mehraban, Joshua Combes, Barry Sanders, and Ish Dhand.

I immensely enjoyed the courses that I took at MIT. I have already mentioned Jeff Shapiro's course on quantum optical communication but there are two more professors I would like to thank in particular. I loved Peter Hagelstein's course on applied quantum and statistical physics because of the window it provided into his thought process, and his fresh perspective on a subject that I had studied before. I am extremely glad that I took Michael Sipser's course on the theory of computation. In addition to providing me an excellent introduction to new areas, he was patient enough to answer the many questions that I had for him, even after the course was over.

I really enjoyed my semester as a teaching assistant for 6.011: Signals, Systems and Inference. I would like to thank George Verghese, Alan Oppenheim, and Peter

Hagelstein for this extremely enriching experience. My role as a teacher improved my ability to explain research ideas. I would also like to thank the two other TAs from the class: Nalini Singh and Allan Sadun.

I would also like to thank my friends here (there are so many that if I start naming them, I fear to miss out on someone!) who have been my family away from home in my journey. I'm grateful to have a family that has always lent me unconditional support and has been a patient sounding board for all my professional and personal aspirations.

I would like to thank the funding agencies involved in my PhD. I was funded by the Air Force Research Laboratory (AFRL/RITA), Grant No. FA8750-14-2-0120, DARPA project Scalable Engineering of Quantum Optical Information Processing Architectures (SEQUOIA), US Army contract number W31P4Q-15-C-0045, the Air Force Office of Scientific Research MURI (FA9550-14-1-0052), Army Research Laboratory (ARL) Center for Distributed Quantum Information (CDQI), and Office of Naval Research program Communications and Networking with Quantum Operationally-Secure Technology for Maritime Deployment (CONQUEST), awarded under prime contract number N00014-16-C-2069.

It's hard to recall every person that played an important role in the last four years of your life, particularly at a time of such major changes. I apologize to anyone I may have forgotten to mention here.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Quantum mechanical effects are prevalent in computers today. Quantum effects like tunneling and discretization of energy levels form the basis of solid state physics used in transistors. However, even though the internal dynamics of these computers is dictated by quantum mechanical effects, their operation can be described classically i.e., a mechanical system can simulate a transistor based computer with a number of components proportional to the number of transistors. Hence, conventional computers are still *classical* computers.

However, recent experiments have shown that it is possible to isolate quantum mechanical systems from the environment sufficiently well so that superposition and entanglement play a major role. Unlike classical computers, these systems cannot be efficiently described with a system of any components following the laws of classical mechanics. This forms the basis of *quantum* information processors which can be used to accomplish computing and communication tasks beyond the capabilities of classical systems.

1.1 Quantum computing

Quantum computers use superposition and entanglement to perform several computational tasks, including prime factorization [7], optimization [8, 9], database search [10], machine learning [11], and quantum chemistry simulations [12] more efficiently than

classical computers. In some of these examples, quantum computers with resources (both space and time) scaling polynomially in the size of the problem can solve problems that have no polynomially scaling algorithm on a classical computer. Hence, in addition to practical motivations to building quantum computers, a quantum computer would provide strong evidence against the extended Church-Turing Thesis.

A quantum computer is composed of quantum bits or qubits. Since a quantum bit represents a quantum state, it is often written in the Dirac notation as $|0\rangle_L$ for a qubit in the zero state and as $|1\rangle_L$ for a qubit in the one state. However, since a quantum bit can be in a superposition of zero and one, a single qubit in general can be in the state $c_0|0\rangle_L + c_1|1\rangle_L$ with $|c_0|^2 + |c_1|^2 = 1$. Furthermore, a system of n entangled qubits can be in the state $\Sigma_{i\in B_n} c_i|i\rangle_L$ where B_n is the set of all n bit binary numbers and $\Sigma_{i\in B_n}|c_i|^2 = 1$. Hence, tracking the evolution of a general system of n qubits requires tracking the evolution of 2^n coefficients (c_i) which is inefficient in general. However, this does not automatically mean that an n qubit quantum computer allows us to obtain the results of 2^n computations; measurement of an n qubit system collapses the state and only gives n bits of information. Quantum algorithms maneuver the state in such a way that the measurement result contains information about a large number of coefficients, and calculating the measurement outcome would require a large amount of classical computation.

There have been proposals of analog computing in systems ranging from mechanical gears [13] to soap bubbles for solving NP-complete problems (see [14] for a review). However, the maximum problem size in these systems is limited by errors; in all the analog computing proposals to date, the output becomes noisier as the size of the system increases. As a result, these systems do not present a path to obtaining better *scaling* as compared to conventional computers. Quantum computers on the other hand are *not* analog computers. Quantum error correction codes and fault tolerant architectures have been developed which can suppress errors for any system size with an overhead that is poly-logarithmic in the system size. There has been a recent theoretical proposal for fault-tolerant quantum computing architecture with constant overhead [15], although there are some problems which need to be solved including

efficient decoders, raising the threshold and minimizing long range gates.

While these theoretical results show that it is possible to build a quantum computer that far exceeds the capabilities of any classical computer in principle, practical issues have prevented the experimental demonstration of such a system. Firstly, fault-tolerant quantum computation requires the error in every component to be below a certain finite threshold. Secondly, even if the error correction overhead is poly-logarithmic in the size of the problem, the actual number of physical qubits required per logical qubit is large when operating near the error correction threshold (the overhead diverges at the threshold). Finally, even though a quantum computer scales better than a classical computer, reaching a problem size where this scaling advantage shows is challenging because of the incredible maturity of CMOS fabrication used in classical computers. A recent paper investigating the resource requirements for building an error-corrected quantum computer with ion-traps powerful enough to solve problems that are beyond current capabilities, using currently proposed architectures, found that the quantum computer would have the size of a football field [16]. Building so many components while maintaining an error rate below threshold is a daunting engineering challenge and it is quite likely that architectural improvements will be required to bring this down.

1.2 Photonic quantum computing

Photonic quantum computing computing uses photons to encode qubits. In *dual-rail* encoding, one photon in two spatial modes encodes a single qubit. i.e. $|10\rangle \equiv |0\rangle_L$ and $|01\rangle \equiv |1\rangle_L$. The subscript L signifies the qubit state. The modes could be temporal, spatial, polarization or any other degree of freedom. In this thesis, the modes of the dual rail qubit will be assumed to be spatial modes in different waveguides unless specified otherwise.

Single qubit rotations in this encoding can be implemented with a single tunable beamsplitter. Large tunable arrays of such beamsplitters have been demonstrated with high visibility [17, 18, 19]. Furthermore, optical coupling between two spatially

separated waveguides is negligible, and as a result, the qubit decoherence rate is smaller than in other qubit systems [20]. Photonic qubits are prone to errors due to photon loss. However, since photon loss takes a dual-rail qubit outside the qubit encoding, photon loss, even on a single qubit without error correction, is a detectable error. As a result, error correction thresholds for photon loss are as high as 50 % which is much higher than thresholds for errors that keep the qubit in the code space.

One major issue with photonic qubits is the difficulty in implementing two qubit gates. Direct implementation of a two qubit gate requires an extremely large material nonlinearity [21, 22, 23]. Even with such large nonlinearities, the causal, noninstantaneous response of the nonlinearity [24, 25] and unwanted spectral entanglement caused by the nonlinearity [26] lowers the gate fidelity. There has been work on overcoming such effects [27], but the fidelity only approaches one in a large system of coupled nonlinearities.

An alternative is to use effective nonlinearities induced by photon measurement [28, 29]. *Cluster state quantum computing* [1] is a popular quantum computing paradigm which uses such measurement induced nonlinearities. A cluster state is an entangled state that can be depicted as a graph in which each node represents a qubit in the state $(|0\rangle_L+|1\rangle_L)/\sqrt{2}$ and edges represent controlled-Z gates (see Fig. 1-1(a)). A seminal paper by Raussendorf et al. [1] showed how any gate in circuit model quantum computing can be mapped to adaptive measurements on a square grid cluster state as shown in Fig. 1-2. In other words, a series of adaptive measurements on a square grid cluster state can implement any quantum algorithm. Note that other lattice topologies which can be *renormalized* to a square lattice can also be used to implement any quantum algorithm (see chapter 3 for more details). Cluster state quantum computing involves classical computing between measurements, and this makes cluster state quantum computing more efficient than gate-based quantum computing in some ways. As an example, Shor's Algorithm can be implemented using a cluster state of constant depth in the size of the problem [30]. Furthermore, there are well-developed schemes for fault-tolerant quantum computing with cluster states [31, 32, 33].

Cluster state quantum computing reduces the problem of implementing two-qubits

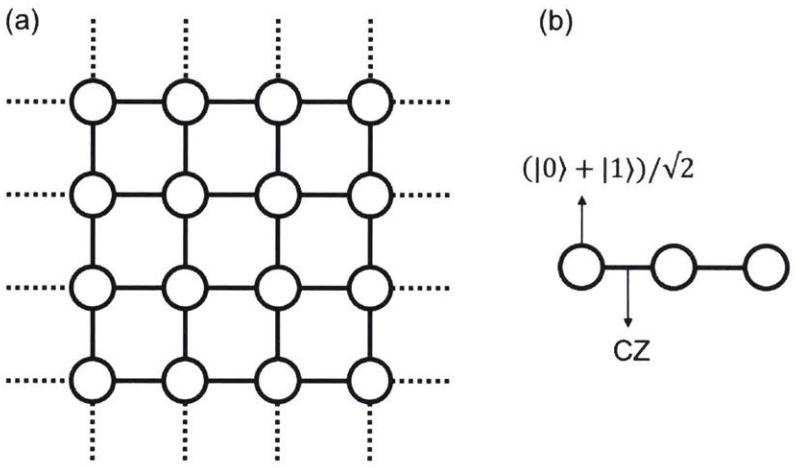

Figure 1-1: (a) A cluster state. The circles represent qubits in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ and edges represent controlled Z gates. Any quantum algorithm can be implemented with adaptive measurements on such a square lattice i.e. such a square lattice is a resource for universal quantum computation. (b) A three-qubit GHZ state.

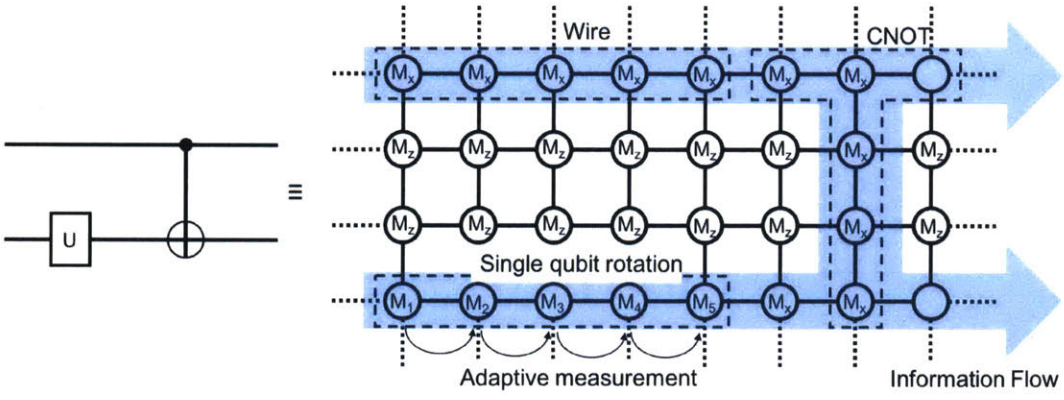

Figure 1-2: Any circuit model quantum computation (left) can be mapped to adaptive measurements on a square grid cluster state (right). M_x, M_z represent measurement in the X and Z bases respectively. M_1 - M_5 are measurements in bases that depend on the intended single qubit rotation and previous measurement results. The thick blue arrows represents information flow. For more details, see Ref. [1].

gates to creating the entangled cluster state and adaptive measurement on this state. Starting from single photons, the cluster creation process can start with the creation of 3-qubit Greenberger-Horne-Zeilinger (GHZ) states [34] which are equivalent to the three qubit cluster state in Fig. 1-1(b), up to single qubit rotations. Since single qubit operations with linear optics are relatively easy, we often compare states in terms of equivalence up to single qubit rotations. 3-qubit GHZ states can be created from six single photons using a probabilistic but heralded linear optic circuit with a success probability of 1/32 [2]. Subsequently, small cluster state can be merged together to form a larger cluster state using Type-I and Type-II fusion gates proposed in Ref. [3], both of which succeed with probability 50%. Subsequent work has shown how the success probability of Type-II fusion can be increased by injecting additional ancilla states into the circuit [35, 36, 37]. Because of this enhancement in the success probability, and the natural loss tolerance of Type-II fusion, we use only Type-II fusion to merge cluster states in this thesis. We defer further discussion of 3-qubit GHZ creation and fusion gates to Chapter 2.

The entangling operations used to create cluster states with linear optics are probabilistic. We will use two methods to create cluster states with probabilistic gates: multiplexed generation of cluster states, which will be explored in Chapter 2, and percolation based creation which will be explored in Chapters 3 and 4.

1.3 Quantum communication

Given the widespread use of light as an information carrier in modern communication systems, photons have an advantage over other systems as carriers of quantum information.

Quantum communication can accomplish a number of tasks including the generation of shared secrets whose security relies only on the laws of physics (better known as quantum key distribution or QKD) [38, 39], distributed quantum computing [40], improved sensing [41, 42], blind quantum computing (quantum computing on encrypted data) [43], and secure private-bid auctions [44].

Since a Bell pair shared between two parties can provide one bit of quantum communication using teleportation, we quantify the quantum communication rate with the entanglement generation rate.

For any direct-transmission protocol over the pure-loss optical channel of transmissivity η, assuming unlimited authenticated two-way public classical communication, it was recently shown that the entanglement generation rate cannot exceed $\log_2(1/(1-\eta))$ bits per mode [45], which is $\approx 1.44\eta$ for $\eta \ll 1$. We will refer to this as the Pirandola-Laurenza-Ottaviani-Banchi (PLOB) upper bound. For a pure-loss channel, the PLOB bound improves over the Takeoka-Guha-Wilde (TGW) bound [46] by a factor of 2 in the $\eta \ll 1$ regime. The TGW bound is an upper bound on the secret-key agreement capacity with unlimited two-way classical communication $P_2(\mathcal{N})$, applicable to a general quantum channel $\mathcal{N}$. For the pure-loss channel $\mathcal{N}_\eta$, the PLOB bound coincides with the best-known achievable rate [47], thus establishing $P_2(\mathcal{N}_\eta) = \log_2(1/(1-\eta))$ bits per mode. The ebits/s rate is obtained by multiplying the bits/mode rate by the spatio-temporal-polarization bandwidth (modes/s), which is governed by the channel geometry, and the transmitter and detector bandwidth. Since loss increases exponentially with distance L in optical fiber (i.e., $\eta = e^{-\alpha L}$), for $\eta \ll 1$, the entanglement generation rate in single-hop quantum communication must decay exponentially with the range L.

One of the big applications of quantum communication is quantum key distribution (QKD) which enables two distant authenticated parties Alice and Bob, connected via a quantum (e.g., optical) channel, to generate information-theoretically secure shared secret bits. No knowledge of the channel conditions (noise model, or any channel estimate) is required a priori to ensure security. However, the shared secret is generated at a rate commensurate with the worst-case adversary physically consistent with the channel conditions actually presented to Alice and Bob. The reason is that all the perceived channel imperfections (anything that causes the channel map to deviate from a noiseless identity transformation) is attributed to the actions of the most powerful adversary allowed by physics—even though some (or all) of that deviation of the channel from an identity map may actually stem from non-adversarial

sources, such as losses due to free-space diffraction, fiber loss, detection inefficiency, thermal noise from blackbody at the operating temperature and wavelength, and detector noise. An important consequence of this assumption is that all the signal power transmitted by Alice that is not collected by Bob is made available coherently to the eavesdropper, Eve. This model for Eve is the intuition behind why the secret key rate for a direct-transmission based QKD protocol must decrease linearly with η, the Alice-Bob power transmissivity, in the $\eta \ll 1$ regime [46, 45].

Quantum repeaters are devices that, when inserted along the length of the optical channel, can help generate shared secret at a rate that surpasses the PLOB bound [45], even when the repeaters themselves are not trusted.

1.4 Thesis outline

Chapter 2 presents schemes for an all-optical repeaters, several improvements to all-optical repeater schemes and calculates the resource requirements for building useful all-optical quantum repeaters using multiplexed creation of cluster states, accounting for losses in detectors, sources and waveguides in the system. Our calculations show that relatively small cluster states can act as useful quantum repeaters. However, the multiplexed creation of cluster states is very inefficient. In chapter 3, we study a more efficient path towards cluster state creation with non-deterministic gates using percolation theory. We present a mapping from the success probability of linear optic fusion operations to the bond percolation threshold of a logical graph, which allows us to find lattices with properties desirable for linear optic quantum computing. Chapter 4 uses the percolation based creation of cluster states in an architecture for quantum computing with atomic memories, that is tailored to nitrogen vacancy centers in diamond. In Chapter 5, we present the design of a quantum network architecture that only requires entanglement swaps at the repeater nodes, but provides several advantages over a linear repeater chain. Finally, Chapter 6 presents both experimental and theoretical work on devices for photonic quantum information processing that I was involved in.

Chapter 2

Resource costs for all-optical quantum repeaters

2.1 Introduction

Quantum repeaters, proposed in Ref. [48], are devices which when inserted along the length of the optical channel, can help generate shared secret at a rate that surpasses the PLOB bound [45]. Quantum repeaters need not be trusted or physically secured to ensure the security of the keys generated. If n quantum repeaters are inserted along the length of a channel of transmissivity η connecting the communicating parties Alice and Bob, and if there are absolutely no physical constraints placed on the repeater nodes (i.e., the repeaters are assumed to be lossless, error-corrected, general purpose quantum computers), then the maximum key rate achievable by Alice and Bob is given by $-\log_2(1-\eta_{\text{min}})$ bits/mode, where $\eta_{\text{min}} \equiv \min(\eta_1, \eta_2, \ldots, \eta_{n+1})$ (with $\eta = \eta_1 \ldots \eta_{n+1}$) is the transmissivity of the lossiest link between successive repeater nodes [49] (see [50] for a different upper bound based on squashed entanglement [46]). Given n repeater nodes, their optimal placement is to lay them equally spaced in transmissivity, in which case, the maximum rate is $-\log_2(1-\eta^{1/(n+1)})$ bits/mode. As $n \to \infty$, the rate is unbounded. However, assuming repeaters to be lossless error-corrected quantum computers is not practical. A more practically relevant question to ask is if the repeater nodes have finite resources with lossy and imperfect compo-

nents (where 'resources' may be different physical entities depending upon the type of quantum repeater and the protocol employed), then what rate can Alice and Bob achieve, and what would it take to build repeater nodes to significantly outperform $R_{\text{direct}}(\eta) = -\log_2(1-\eta)$ bits/mode. This is the topic addressed in this chapter, for repeaters that are built solely using photonic components—single-photon sources, detectors, electro-optic feedforward, but no matter-based quantum memories. As we will see later, given physical constraints on a repeater node, placing more repeaters (higher n) between Alice and Bob may not always improve the rate, i.e., depending upon the total distance L (or equivalently, the transmissivity η) between Alice and Bob, and given the physical device constraints in a repeater node, there may be an optimal number $n^*(\eta)$ of nodes, which achieves the highest end-to-end rate.

In this chapter, we study two different all-optical repeater protocols: a protocol based on mode multiplexing and entanglement swapping, and a one-way protocol based on logical Bell measurements. We found several improvements for both protocols to reduce resource requirements. We consider losses in sources, waveguides and detectors in every step of the protocol, including the creation of the resource states for our protocol. However, we do not consider 'multi-photon' errors that may stem from multi-photon emissions from the source, detector dark clicks, or errors due to mode mismatch in the sources. We should note however that the error correction schemes analyzed here also provide some protection against depolarizing noise [51, 52], a variant of which arises when one assumes multi-photon errors, and errors stemming from imperfect mode matching within the passive linear optical circuits at the repeater nodes.

2.2 Preliminaries

We work with dual-rail photonic qubits, where the logical $|0\rangle_L$ and $|1\rangle_L$ are encoded by a single photon in one of two orthogonal (spatial) modes. We use the term *qubit mode* to refer to a set of 2 channels used to transmit a dual-rail qubit between repeater stations. Hence, m qubit modes correspond to $2m$ physical channels.

A photonic *cluster* state (or, graph state), on a graph $G(V, E)$ with vertices in set V and edges in set E, can be constructed by preparing each of the $|V|$ qubits (one stationed at each vertex) in the state $(|0\rangle_L + |1\rangle_L)/\sqrt{2}$, and applying $|E|$ controlled-phase operations (a two-qubit unitary gate that applies a Pauli Z gate to the second qubit if the first qubit is in the $|1\rangle_L$ state and applies an identity otherwise) on each pair of vertices that share an edge [1]. An ideal loss-less photonic cluster state on graph G is a pure state, $|\psi\rangle_G$. A lossy cluster state on G is obtained when all the photonic qubits of $|\psi\rangle_G$ are transmitted through independent pure-loss beamsplitters each of transmissivity η. We call $1 - \eta$ the *loss rate* of such a lossy cluster state. Clearly, the loss rate of $|\psi\rangle_G$ itself is 0.

Arbitrary photonic cluster states can be prepared—with non-unity probability—using ideal single photons, passive linear optics (i.e., beamsplitters and phase shifters) and single photon detectors [3]. As examples, in the absence of losses, a two-photon maximally entangled (Bell) state can be prepared with success probability 3/16 [53], whereas a three-photon maximally entangled (GHZ) state can be prepared with success probability 1/32 [2]. Browne and Rudolph introduced linear-optic Type I and Type II two-qubit fusion gates, which if successful (with probability 1/2), can fuse two cluster fragments into one, according to specific rules [3]. These fusion gates, in conjunction with Bell states and GHZ states, can be used to construct arbitrary cluster states. The success probability of the fusion gates can be improved to 3/4 if additional (ancilla) single photons are available to be injected on-demand into an otherwise-passive linear optical circuit, and if the detectors have up to two-photon number resolution [36]. We will assume such *boosted* fusion gates in our all-optical repeater construction described in this chapter.

We will model a lossy single photon source of efficiency η_s as one that emits, on demand, the mixed state $\eta_s |1\rangle \langle 1| + (1 - \eta_s) |0\rangle \langle 0|$. We will use η_d for the efficiency of all detectors in the system. We will consider that the cluster is created on a photonic chip to allow for easier scalability after which the photons are coupled, with efficiency P_c, to fiber with loss coefficient α and speed of light c_f. $P_{\text{ff}} = e^{-\beta \tau_s c_{\text{ch}}}$ will denote the survival probability of a photon on-chip during one feed-forward step, where β is the

loss coefficient, c_{ch} is the speed of light and τ_s is the feed-forward time, all on-chip. We will also use $P_0 \equiv P_{\text{GHZ}} P_c$ where $P_{\text{GHZ}} = \eta_s \eta_d / (2 - \eta_s \eta_d)$ [2] is the survival rate of the photons that are input into a linear-optic circuit to produce 3-photon maximally-entangled GHZ states. The final measurement step requires feed-forward in fiber. The survival probability, $P_{\text{ff(fib)}}$, during feed-forward time in fiber, τ_f, is $P_{\text{ff(fib)}} = e^{-\alpha \tau_f c_{\text{f}}}$. The values for device performance assumed for the plots that appear later in the chapter, are summarized in Table 2.1.

2.3 Quantum repeaters based on mode multiplexing and entanglement swapping

In this section, we will focus on a class of quantum repeaters that rely solely on probabilistic BSMs, quantum memories, and multiplexing, i.e., the ability to 'switch' qubits across (spatial, spectral, or temporal) modes. The essence of such a repeater protocol was developed by Sinclair *et al.* [54], which employed spectral multiplexing in multimode quantum memories across m parallel spectral channels, and entanglement swapping using linear optics and single photon detectors (the success probability of which can at most be 50%). Guha *et al.* analyzed the secret key rates achievable by the above protocol, with a fixed m (memory size) and found that even when photon loss is the only source of noise, the achievable key rate is of the form $D\eta^s$, where D, and $s < 1$ are constants that are functions of various losses in the system (viz., detection efficiencies, coupling losses, memory loading and readout efficiencies, and BSM failure probability) [55]. Since the exponent of η, i.e., $s < 1$, the key rate beats the PLOB bound (which scales as: $\propto \eta$ for $\eta \ll 1$) beyond a certain minimum distance determined by the actual values of the system's loss parameters, which is around a couple of hundred kilometers for reasonable estimates of the losses [55]. However, since $\eta = e^{-\alpha L}$ in fiber, the rate achieved by this repeater protocol for a fixed memory size still scales exponentially with the range L, albeit with a smaller exponent compared to the best possible rate without any repeater, which could turn

into a huge absolute improvement in the end-to-end secret key rate.

Azuma *et al.* recently proposed an all-photonic variant of this protocol in which they substituted matter based quantum memories with optical cluster states [51], based on a proposal by Varnava *et al.* to mimic a quantum memory (i.e., protect against photon losses) by appending each physical photonic qubit by an entangled 'tree cluster' state [56]. As long as the losses incurred by each photon (i.e., photons being protected as well as the additional photons in the trees added for loss protection) is less than 3 dB, the effective loss of the logical qubit can approach zero, by increasing the size of the tree cluster, i.e., the number of photons in the logical qubit [57]. Thus, Azuma *et al.*'s proposal showed the theoretical feasibility of a quantum repeater architecture (i.e., one that can beat the scaling of direct-transmission QKD) using only flying qubits, with the repeater *nodes* being equipped only with single photon sources, passive linear-optical circuits (beamsplitters and phase shifters), single photon detectors, and classical feedforward.

Azuma *et al.*'s result marked a promising conceptual leap towards all-optical quantum repeaters. However, important unanswered questions remained, including the achievable secure key generation rate and how it scales with distance (or loss), as well as the physical resource requirements: e.g., the number of photon sources and detectors at the repeater nodes. As an example, a calculation in their paper shows that at a range of $L = 5000$ km, an entanglement-generation rate of 69 kHz is achievable in a fiber based linear optic system with 100 kHz repetition rate, 150 ns feed forward time and a source-detector efficiency product of 95% whereas sharing a single entangled pair via a direct transmission scheme with the same parameters would require 10^{81} years. The level of error protection required to achieve the aforesaid repeater performance at $L = 5000$ km would require one to build entangled clusters of $\sim 10^4$ photons per clock cycle at each repeater node. Building such a cluster using linear optics and feed-forward [2, 58] would require around 10^{24} photon sources at each repeater node. Furthermore, since every photon used for error correction is sent between repeater nodes in [51], their scheme would require around $20,000$ parallel channels connecting the neighboring nodes. Thus, while Ref. [51] crucially showed

the theoretical feasibility of all-optical repeaters, further work is needed to address their practical feasibility. These results open up a compelling line of research to investigate improved all-photonic repeater architectures of various genres which could be built with practically feasible resources and also a thorough comparative study of rates achievable with each such all-optical repeater scheme.

Our contributions in this section are twofold. The first is a rigorous analysis of: (a) the secret key rates achievable with the aforesaid all-photonic repeater architecture given the size of the cluster at each repeater station, and (b) the resources required (e.g., number of single photon sources and detectors required at each repeater node) to build that cluster, while taking into account each step in building the required clusters using a network of passive linear optics (i.e., beamsplitters and phase shifters), imperfect on-demand sources with loss (see section 2.2 for a description of the source), single photon detectors (with some number resolving capability), and feed-forward. We find that the achievable secret key rate scales as $D\eta^s$ bits/mode, where D and $s < 1$ are functions of the number of photon sources at each repeater node (the resource constraint—which is parametrically related to the size of the cluster), all the 'inline' losses (e.g., losses in the optical fiber or waveguide used while creating the cluster, independent of the fiber loss between repeater stations), and the source and detector efficiencies. With $\eta \sim e^{-\alpha L}$ in fiber, the key rate still scales exponentially with L, but with a smaller exponent compared to the best direct-transmission protocol. This is no surprise given the analysis of [55], since the tree-cluster construction of [51] essentially mimics an imperfect quantum memory, but one whose efficiency cannot simply be modeled by a constant per mode as in Ref. [55]. Using the cluster building scheme proposed by Li et al. [58], we find that to a good approximation, the resource requirements are determined by the number of fusion steps k required to build the cluster and hence, we calculate the performance of the best cluster that can be built in k fusion steps. We use the scheme of Li *et al.* because it has been shown to be more efficient than the scheme of Varnava *et al.* [2] at building clusters [58]. Given all the inline and device losses, we evaluate the number of photon sources (and detectors) needed at each repeater node to beat the PLOB bound at a given total range L

between Alice and Bob. We also prove that given the device losses, there is an optimal spacing between the repeater nodes (which evaluates to roughly 1.5 km for a set of system parameters), regardless of the overall range L.

Our second major contribution is a significant improvement to the all-photonic repeater architecture in [51]—both in terms of the resources required at each node and the number of parallel optical channels connecting the neighboring nodes. We find that barely beating the PLOB bound using the all-optical scheme of [51] requires more than 10^{11} photon sources at each repeater node for realizing the required optical cluster states and measurements. It also requires 208 parallel channels connecting neighboring nodes, even when assuming very optimistic device-loss parameters. Assuming the same device losses, our improved repeater architecture reduces the number of photon sources (to beat the PLOB bound) by 5 orders of magnitude, while reducing the number of channels to 8. In both these calculations, each source is used only once per clock cycle, i.e., they are not temporally multiplexed. We prove a tight analytical lower bound for the performance of our improved scheme. These performance advances are enabled primarily by the following: (1) using boosted fusion logic that improves the success probability of the BSM to 75% by using four ancilla single photons [36], (2) employing a more resource-efficient scheme for creating tree clusters, building on the work of Refs. [2, 58], (3) retaining all the ancilla photons used for loss protection (i.e., to mimic a quantum memory) locally at the repeater nodes in a lossy waveguide, and (4) optimizing the timing of several single qubit measurements.

2.3.1 Counterfactual error correction

The cluster state represented by a graph $G(V, E)$ is an eigenstate of the $|V|$ stabilizer operators $X_i \ \Pi_{j\in\mathcal{N}(i)}Z_j$, where the index i runs over all the vertices, X_i and Z_j are Pauli X and Z operators on qubit i and qubit j respectively, and $\mathcal{N}(i)$ is the set of all nearest neighbor vertices of vertex i. One simple observation, given that the cluster state is an eigenstate of the aforesaid stabilizer operators, is that an X measurement on qubit i, and Z measurements on all but one of the qubits in $\mathcal{N}(i)$, would deterministically reveal what the outcome of a Z measurement on that

unmeasured qubit in $\mathcal{N}(i)$ would have been, *even if* that unmeasured qubit had been lost. This realization is at the heart of the tree-based counterfactual error correction for protection against photon losses, developed by Varnava et al. [57]. The idea is to attach a tree cluster to each physical photonic qubit in the graph state that needs to be protected against qubit loss. One can then deduce the result of any measurement on that qubit via an appropriate sequence of measurements on the qubits of the attached tree. The physical qubit and the qubits of the tree together form a protected (logical) qubit. We consider regular trees described by the *branching vector* $\vec{b} \equiv \{b_0, b_1, \ldots, b_m\}$, which signifies that the root of the tree has b_0 children nodes, and each of those nodes have b_1 children nodes, and so on until $b_0 b_1 \ldots b_m$ nodes at depth m. For such regular trees used for loss-error protection, one can write an explicit, yet recursive, expression for the success probability P of performing an arbitrary single-qubit measurement on the protected qubit [57]. It was shown that one can push P arbitrarily close to 1 as long as the probability of losing each photon is less than 1/2. Fig. 2-1 illustrates how to attach a $\{3, 2, 2\}$ tree, shown by the dark (purple) shaded nodes, to a physical qubit of a cluster, shown by light (green) shaded nodes. Note that after the tree cluster is attached to the physical qubit, X basis measurements must be performed on the physical qubit itself and the root node of the tree. These X basis measurements, if successful, create additional edges (shown in dashed blue in Fig. 2-1) between each neighboring qubit of the root node and each neighboring qubit of the physical qubit, after which the tree-protected logical qubit is ready to use.

2.3.2 Repeater architecture

Before discussing the all-photonic repeater architecture, it is instructive to review a generic quantum repeater architecture based on multimode quantum memories, probabilistic BSMs, and multiplexing over m orthogonal qubit modes depicted in Figs. 2-2(a) and (b), which was proposed by [54], and analyzed in Ref. [55]. Alice and Bob are separated by optical fiber of length L (i.e., end-to-end transmissivity, $\eta = e^{-\alpha L}$), interspersed with n repeater stations spaced $L_0 = L/n$ apart, with Alice

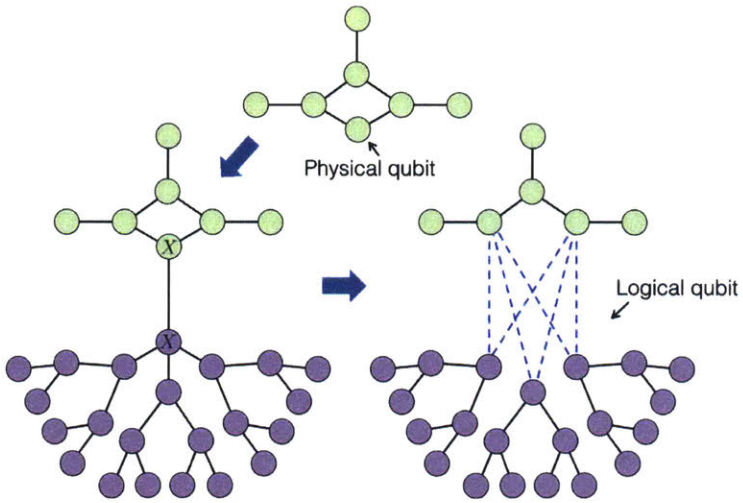

Figure 2-1: Attaching a $\{3, 2, 2\}$ tree to a node of a photonic cluster.

and Bob $L_0/2$ away from the terminal repeaters in the chain.

Each of the n repeater nodes (or, 'major nodes'), shown by a gray box, consists of a multimode quantum memory straddled between sources of m Bell pairs on its left and another m on its right. Each major node loads one half of an entangled Bell state onto the memory, while transmitting the other half towards the middle of the adjoining *elementary link*. Each major node does the above synchronously on every clock cycle. At the center of each elementary link is a 'minor node', shown as dark-blue-shaded boxes at the bottom in Fig. 2-2(b). After the qubits from the major nodes reach the minor nodes (i.e., after propagation through a distance $L_0/2$), each minor node, simultaneously, performs BSMs on each of the m pairs of qubits received from the repeater nodes on its either side. The successful BSMs within each elementary link are shown by thick (green) line segments. Immediately after the minor node BSMs, each minor node sends back the information—about which (of the m) qubit modes were successfully measured—to its two neighboring major nodes, on an authenticated classical channel. Upon receipt of that information, each major node performs a BSM on two qubits held in its memory that had been entangled halves of qubits that participated in successful BSMs at the minor node to the left of that major node, and the minor node to its right, respectively. Simultaneous with the minor-node BSMs, Alice and Bob measure, in one of the two randomly-chosen mutually-unbiased bases,

the m qubits they receive at their respective ends of the terminal half-elementary-link segments (see Fig. 2-2(b)), and send the information about which channels generated a 'click' on their detectors, back to their respective neighboring major nodes. Finally, each major node sends the information on whether its BSM succeeded, to Alice and Bob. Hence, at every clock cycle, with some probability (i.e., if all the minor nodes heralded at least one success each, all major node BSMs were successful, Alice and Bob both detect a photon on at least one of the m qubit modes each while using the same measurement bases), Alice and Bob obtain a shared (raw, sifted) bit. A long sequence of sifted bits is thereafter used to distill a quantum-secure shared secret via error correction and privacy amplification.

The all-optical repeater architecture we will now discuss builds upon a recent proposal by Azuma *et al.* [51], although there are some important differences, which we will point out later in Section 2.3.7. The key idea is to mimic a quantum memory (whose goal is essentially to protect photonic qubits against loss for a certain time duration) by using the tree cluster approach described in Section 2.2. The authors of [51] went one step further and subsumed the functionalities of all the subcomponents of the major node (the quantum memory as well as the $2m$ Bell pair sources) into one single giant optical cluster state, which we will describe next. Fig. 2-2(c) illustrates the construction of this cluster. We start with a depth-2 star cluster with a degree-$2m$ root node, and $4m + 1$ total qubits. The 'outer' qubits, shown as white circles, play an analogous role to the white qubits in Fig. 2-2(a) that are transmitted to the minor nodes on fiber channels. The $2m$ 'inner' qubits, shown as gray circles, are each attached with a tree cluster of an appropriately-chosen branching vector $\vec{b}$, thereby creating a giant tree cluster. The loss-protected (logical) inner qubits play a dual role, that of the black qubits in Fig. 2-2(a) that are held in the quantum memories locally at the major nodes, and that of the memories themselves. We make the two X measurements corresponding to each tree appended to the star, as described in the previous section (i.e., a total of $4m$ X measurements). Finally, we make a Y measurement on the root node of the star, which has an effect of creating a clique among all the (logical) inner qubits, shown by black circles in Fig. 2-2(c). The clique

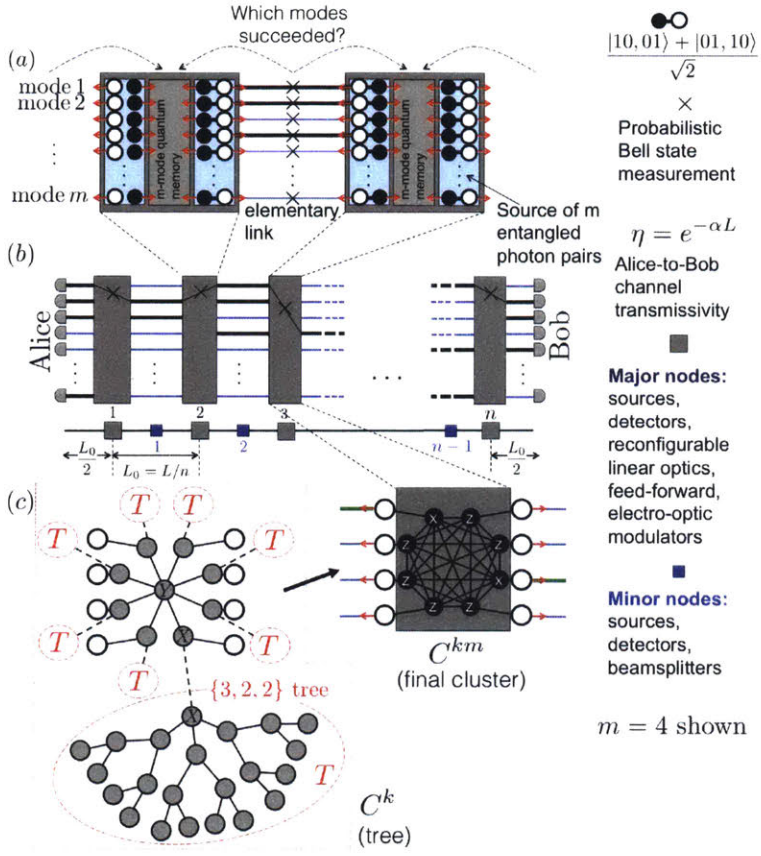

Figure 2-2: (a) and (b) show schematics of one elementary link, and a chain of them connecting Alice and Bob, respectively, for a repeater architecture that employs quantum memories, Bell pair sources, probabilistic BSMs, and multiplexing over m orthogonal qubit modes (parallel channels). (c) depicts the construction of a photonic cluster state that can subsume the roles of the quantum memory and the Bell pair sources, thereby resulting in a quantum repeater architecture based solely on 'flying' qubits. The outer (white) photonic qubits are transmitted on the fiber channels, and the inner (black) qubits are held locally in a (lossy) waveguide at the repeater node. See text for a detailed description.

of the $2m$ logical inner qubits, connected to the $2m$ outer qubits, forms the full photonic cluster state that each major node creates every clock cycle, and sends out the $2m$ outer qubits (the white circles) towards the neighboring minor nodes (m to the left and m to the right) on fiber channels. Note that the final cluster state (after the X and Y measurements) is not a tree.

Each major node is equipped with single photon sources, reconfigurable passive linear optics, and single photon detectors. The clusters are created using linear op-

tics and feed-forward [2, 58]. Since the cluster creation process is probabilistic, the resources (number of photon sources, detectors, size of linear optic circuit) must be chosen to ensure a near-unity success probability of creating the cluster in every clock cycle (see Fig. 2-5).

The minor nodes are identical to what was described earlier. The remainder of the protocol proceeds exactly as described at the beginning of this Section in the context of the memory-based architecture, except for the following difference of the action at the major nodes. When the information about which modes were successful comes back at a major node (from the two neighboring minor nodes), instead of doing a BSM between a pair of qubits held in a memory, the major node applies X measurements on the two logical inner qubits corresponding to the successful modes on either side of the clique, and makes Z measurements on the remaining $2m-2$ logical inner qubits (see Fig. 2-2(c)). The X measurements have the effect of fusing the successful outer qubits into an entangled chain, and the Z measurements have the effect of removing the extraneous qubits from the cluster.

So, in any given clock cycle, if the photonic clusters at each major node are successfully created (which includes success in performing the $4m$ X measurements and one Y measurement), if all the minor nodes herald at least one BSM success, if the logical (inner) qubits survive the local storage at the major nodes while the outer qubits fly to the minor nodes and the classical information (about which modes were successful) arrives back, if the two X measurements and $2m-2$ Z measurements done to prune the clusters at the major nodes using that classical information are successful, and if Alice and Bob get at least one click each while using same measurement bases, then Alice and Bob obtain a raw sifted shared bit. In Section 2.3.6, we will explicitly calculate this overall success probability, and the resulting secret-key generation rate. As we will see, larger error-protection trees will afford better rate performance (up to a limit governed by the device loss rates), but creating larger clusters at the major nodes will require more resources (sources and detectors).

In Section 2.3.3, we will describe in detail the construction of the clusters at the major nodes using linear optics, and calculate the success probability. In Section 2.3.4,

we will describe how the measurements on the major-node clusters are done, after the BSMs at the minor nodes, to stitch together an end-to-end entangled state between Alice and Bob.

2.3.3 Constructing the clusters at the major nodes

The cluster as described above, prepared at each major node in every clock cycle, is pieced together by fusing single photons into progressively larger cluster fragments, probabilistically, using linear-optic circuits and photon detectors. The optimal algorithm for create photonic cluster states using linear optics—in terms of minimizing the total number of photons consumed and maximizing the eventual probability of success—is not known even for a general N-node line cluster. With losses from sources detectors and waveguides during cluster construction, finding the optimal recipe becomes even harder. One design knob is the number of redundant cluster fragments attempted at each step. A higher number of attempts improves the probability of successfully creating the final cluster, but with a higher number of required photon sources and detectors. We will refer to this trick of attempting the creation of multiple identical cluster fragments at each step of the process as *multiplexing*.

We will describe now the resource counts and success-probability calculations for two methods to create the cluster at the major node. The first one is a method implied by previous rough estimates of the resource requirements [58, 2]. We will then discuss an improved scheme that decreases the resource requirements during the creation process. Fig. 2-4 provides a schematic for these two schemes, which we will refer to in the discussion below.

Let us label the final cluster C^{km} (see Fig. 2-2) where the letter m signifies that the Y measurement required to turn the inner qubits of the star into a clique (a fully interconnected graph) and the X measurements required to connect the error protection trees to the inner qubits have already been applied. Before these measurements, the (tree) cluster is labelled as C^k. The daughter clusters that are fused together to create C^k will be labelled as C_1^{k-1} and C_2^{k-1}. The daughter clusters that are fused together to create C_1^{k-1} are: $C_{1,1}^{k-2}$ and $C_{1,2}^{k-2}$. The clusters that are fused together

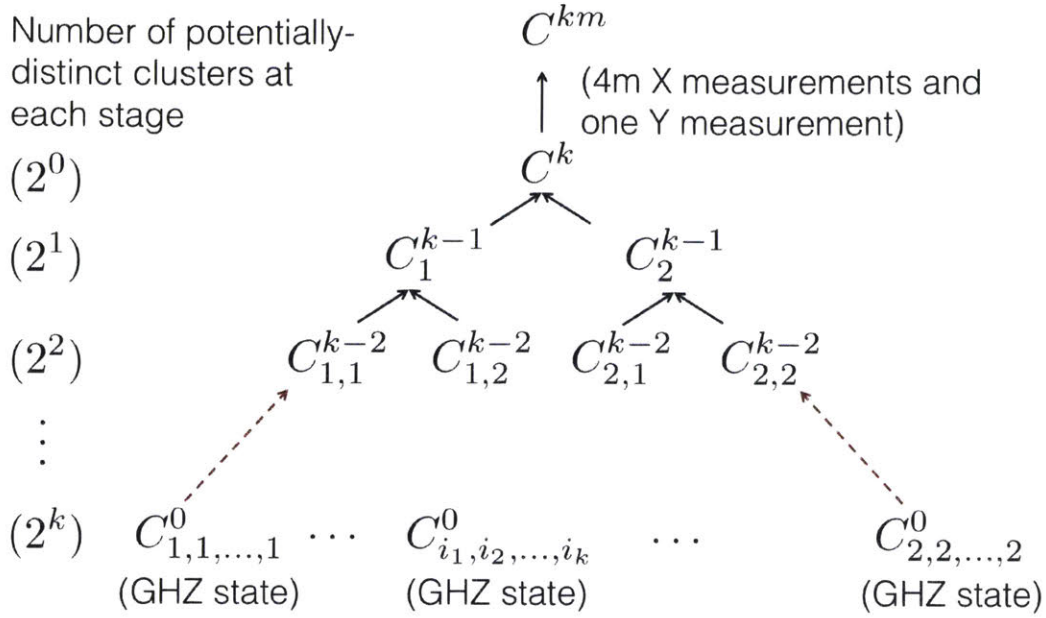

Figure 2-3: The tree cluster C^k (and the final cluster C^{km} after the X and Y measurements), shown in Fig. 2-2, are created by a sequence of probabilistic linear-optic fusion-II operations, starting from 3-photon maximally-entangled (GHZ) states.

to create $C_{1,2}^{k-2}$ are: $C_{1,2,1}^{k-3}$ and $C_{1,2,2}^{k-3}$, and so on (See Fig. 2-3). At the bottom of the stack are 3-photon GHZ states, $C_{\boldsymbol{i}}^0$ with $\boldsymbol{i} \equiv i_1, i_2 \ldots, i_k$, which are in turn created by groups of 6 photons fed into linear-optic circuits that generate the 3-photon GHZ states with probability $S_{\mathrm{GHZ}} = \left[\eta_s\eta_d(2-\eta_s\eta_d)\right]^3/32$ [2]. The loss rate of the heralded GHZ states is, $1 - P_{\mathrm{GHZ}}$ where $P_{GHZ} = \eta_s\eta_d/(2-\eta_s\eta_d)$ [2].

We will assume that the cluster C^k can be prepared in a series of k fusion steps, where at each step, clusters of roughly equal sizes are fused together, thus roughly doubling the cluster size in each step [58]. This assumption becomes accurate in the limit of large clusters. This method ties the final size of the intended cluster ($Q = 2^k + 2$ photons) to the number of fusion steps (k), and this relationship becomes increasingly exact as k becomes large. In other words, we will assume that $C_{\boldsymbol{i},1}^{l-1}$ and $C_{\boldsymbol{i},2}^{l-1}$ are two clusters each of p photons, which when fused successfully using a fusion-II gate (applied to one photon each of the above two clusters) creates the $2p-2$ photon cluster $C_{\boldsymbol{i}}^l$, $\boldsymbol{i} \equiv i_1, i_2 \ldots, i_{k-l}$. Starting with the 3-photon GHZ states $C^0_{i_1,i_2\ldots,i_k}$, the size of C^k is 2^k+2 photons. Hence, the minimum number of fusion steps

required to build a Q photon cluster is $k = \lceil \log_2(Q-2) \rceil$. The label k, the number of fusion-II steps used to arrive at C^k, also translates to the resource requirements, and the loss rate of each photon in the final cluster, as we show below. Note that k is a function of the branching vector $\vec{b}$ of the error-correction trees used. The larger the error-correction trees, the larger will be the final cluster C^k, and the larger will be the number of steps k required to prepare that cluster.

The naive multiplexing scheme

Let us now examine the cluster creation process (depicted for $k = 2$ in Fig. 2-4(a)). At every point we need the cluster fragment C_i^l, we attempt to create n_B copies of that identical cluster ($n_B = 3$ shown in Fig. 2-4(a)), of which hopefully one is successfully created and heralded for further use. Therefore, creating one usable copy of C^k requires $(2n_B)^k$ GHZ states $C^0_{i_1,i_2\ldots,i_k}$ at the bottom of the stack. Each GHZ state is picked from $n_{\rm GHZ}$ parallel-attempted GHZ states ($n_{\rm GHZ} = 4$ shown in Fig. 2-4(a)), and creating each GHZ state requires 6 single photons. Therefore, creating one usable copy of C^k requires $(2n_B)^k \times 6n_{\rm GHZ}$ single photons. Finally, at the top of the chain, we create $n_{\rm meas}$ copies of C^k in parallel ($n_{\rm meas} = 4$ shown), on each of which the $4m$ X measurements and one Y measurement are performed, to prepare copies of the final required cluster C^{km}. We choose $n_{\rm meas}$ such that we obtain with high probability one successfully-created copy of C^{km}. Therefore, the total number of single photon sources (shown by black dots at the bottom of Fig. 2-4(a)) that need to simultaneously fire on every clock cycle, $N_s = 6n_{\rm GHZ}\, n_{\rm meas}(2n_B)^k$.

The probability of successfully creating a GHZ state $C^0_{i_1,i_2\ldots,i_k}$ is $S_0 = 1-(1-S_{\rm GHZ})^{n_{\rm GHZ}}$. The success probability of fusion at the l-th step—i.e., that of combining $C^{l-1}_{i,1}$ and $C^{l-1}_{i,2}$ into C^l_i—is given by $S_{B(l)} = (P_{\rm GHZ}P_{\rm ff}^l)^2/2$. The success probability of heralding one cluster C_i^l (from the n_B parallel copies attempted) is given by the recursive formula, $S_l = 1-(1-S_{l-1}^2 S_{B(l)})^{n_B}$, with S_0 given as above. The $4m+1$ (X and Y) measurements required to convert C^k to the final cluster C^{km} succeed with probability $S_{\rm mXY} = \left(P_{\rm GHZ}P_{\rm ff}^{k+1}\right)^{4m+1}$. Since this step is multiplexed over $n_{\rm meas}$ parallel attempts, the success probability of heralding one copy of the final cluster

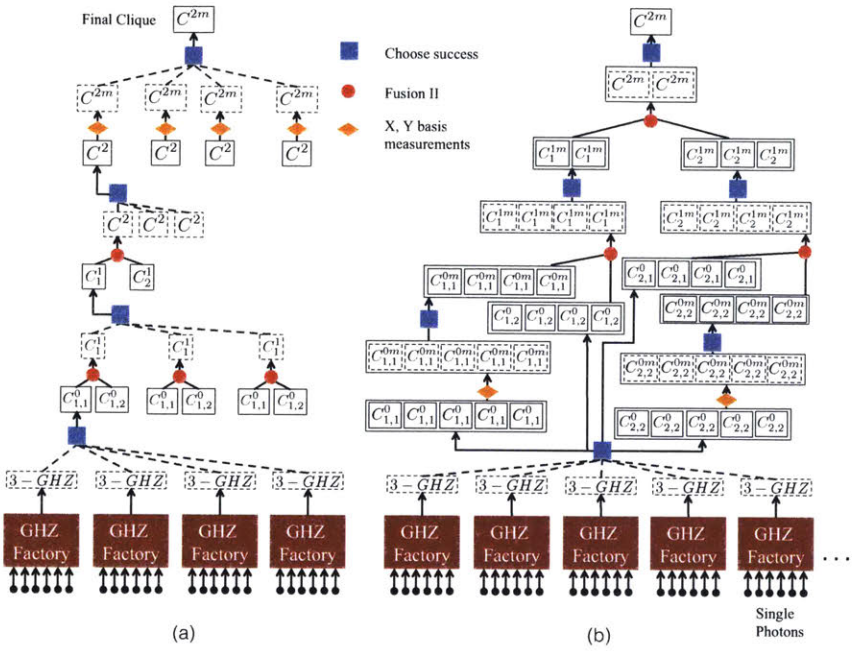

Figure 2-4: (a) naive multiplexing scheme. A dashed rectangle represents a cluster that has some probability of having been been created after a probabilistic fusion step (red circle) or at the output of creating GHZ states using linear optics starting from six single photons (labeled 'GHZ Factory'). A solid rectangle represents a cluster state that is successfully created with high probability by choosing a successful outcome (blue square) out of several identical copies attempted (dashed boxes). (b) the improved multiplexing scheme. A box surrounding clusters of the same type represents a bank of clusters and any operation applied to the bank is applied to all the clusters in it.

at a major node is given by, $P_{c1} = 1 - (S_{B(k)} S_{\mathrm{mXY}})^{n_{\mathrm{meas}}}$. The success probability of all n repeater nodes creating the clusters C^{km} locally during any given clock period, is $P_{cn} = P_{c1}^{n}$. The blue (dashed) plot in Fig. 2-5 shows P_{cn} as a function of N_s for $n = 250$ repeater stations (major nodes), $k = 7$, and for device parameters as given in Table 2.1.

The improved multiplexing scheme

The improved multiplexing scheme we now describe addresses the following deficiencies of the scheme described above.

- The protocol presented above does not make the most optimal use of the multi-

ple copies of identical clusters that are successfully created at a given step. To illustrate this point, let us consider the $n_B = 3$ copies of (attempted) C^2 clusters that are shown in Fig. 2-4(a), of which one successfully created C^2 is picked. The first of those three attempted C^2 clusters is shown to be created by fusing a C_1^1 cluster and a C_2^1 cluster. The C_1^1 is chosen out of $n_B = 3$ copies of (attempted) C_1^1 clusters, as shown. If two of those three copies of C_1^1 are actually successfully created, the second success goes to waste. Note however that the second and the third (of the three attempted) C^2 clusters will also each need to be created by fusing a C_1^1 and a C_2^1. Those two C_1^1 clusters will also be picked from $n_B = 3$ copies each of (attempted) C_1^1 clusters (not shown in the figure). It is thus simple to see that at each time step, a total of $(n_B)^k = 9$ copies of C_1^1 are attempted, but the selection of successes only happen within groups of three, which is clearly inefficient. A far more efficient approach is to maintain one single "bank" of copies of C_1^1 and similarly one single bank for copies of C_2^1, and attempt fusions on clusters from these two banks pairwise (and throw away the excess clusters in the bank that has more copies), to produce a single bank of C^2 clusters. This way, one does not have to choose the multiplexing numbers n_B, n_{GHZ} and n_{meas}, and the total number of single photons N_s directly translates to an overall probability of success P_{c1} of creating the final cluster C^{km}. In general, we maintain single banks of each distinct cluster fragment consumed in the entire stack shown in Fig. 2-3, and for each fusion step shown in Fig. 2-3, we apply pairwise fusion to *all* cluster copies from the two banks corresponding to the two daughter clusters (and throw away the excess clusters from the bank that has more).

- The X and Y measurements that were performed at the very end (on $4m + 1$ nodes of the tree cluster C^k, to convert it to the required final cluster C^{km}) can be performed at the very beginning—on the appropriate photons (which would eventually become those $4m+1$ photons in C^k)—while they are still part of the 3-photon GHZ states, i.e., before any of the fusion-II operations begin.

Making these measurements at the bottom of the stack makes failures much less costly, which in turn significantly reduces the resource requirements (i.e., the N_s required to achieve a given final success probability P_{cn}). Section 2.3.5 rigorously explains why these measurements can be done on the photons while they are still parts of the GHZ states.

- The success probability of each of the fusion-II operations (at all k steps in the cluster creation process) can be improved from 1/2 to 3/4 by injecting ancilla single photons [36]. These success probability numbers diminish with source and detection inefficiencies. But, the cost of using additional photons needed (as ancillas) to realize these *boosted* fusion gates is far outweighed by the effect of the success-probability improvement, thereby improving the effective tradeoff between N_s and P_{cn}.

We start with N_s photons and send them all through GHZ factories, hence attempting the creation of $\lfloor N_s/6 \rfloor$ 3-photon GHZ states. The number of GHZ states x successfully created follows a binomial distribution $B(x, \lfloor N_s/6 \rfloor, S_{\text{GHZ}})$ where $B(x,n,p) = \binom{n}{x} p^x (1-p)^{n-x}$. Hereafter, let us follow an illustrative set of numbers for a $k = 2$ cluster, which is depicted schematically in Fig. 2-4(b). Suppose we get $x = 18$ successfully-created GHZ states. These GHZ states are now split into 4 banks corresponding to $C^0_{1,1}$, $C^0_{1,2}$, $C^0_{2,1}$ and $C^0_{2,2}$. Out of these, let us say $C^0_{1,1}$ and $C^0_{2,2}$ consist of photons that would be eventually measured in C^k. As discussed in section 2.3.5, these qubits can be measured now. Since the measurement of photons has a success probability $P_{\text{ff}} P_{\text{GHZ}}$, the number of $C^{0m}_{1,1}$ cluster states (x) created as a result of making measurements on y $C^0_{1,1}$ states follows a binomial distribution $B(x, y, P_{\text{ff}} \eta_{\text{GHZ}})$. The banks corresponding to $C^0_{1,1}$ and $C^0_{2,2}$ are given a fraction $1/(P_{\text{ff}} P_{\text{GHZ}})$ more GHZ states. Hence, these banks have 5 GHZ states each whereas the other two have 4 each. Suppose that measuring the 5 copies of $C^0_{1,1}$ results in 4 copies of $C^{0m}_{1,1}$, and measuring the 5 copies of $C^0_{2,2}$ results in 4 copies of $C^{0m}_{2,2}$. The first fusion step is now attempted (i.e., fusing $C^{0m}_{1,1}$ with $C^0_{1,2}$, and fusing $C^0_{2,1}$ with $C^{0m}_{2,2}$) resulting in 2 successfully created copies of C^{1m}_1 and 3 copies of C^{1m}_2 (the maximum possible number of successes in both cases

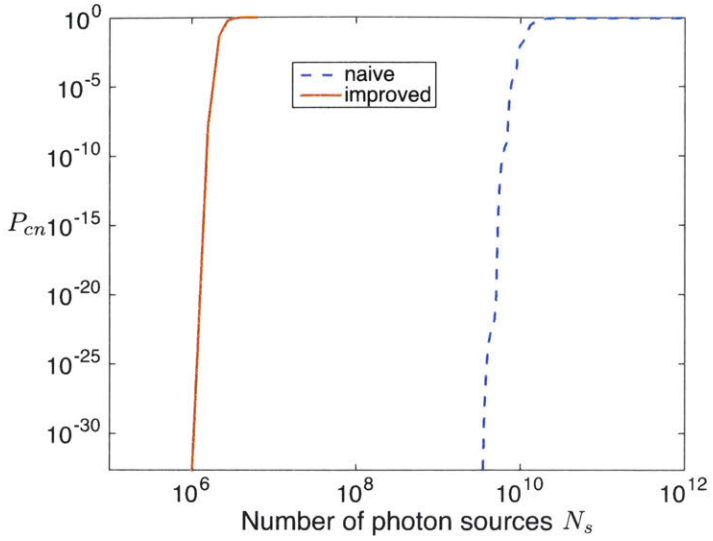

Figure 2-5: The probability that all $n = 250$ major nodes are simultaneously successful in creating clusters of size $k = 7$ fusion steps (i.e., $2^k + 2 = 130$ photon clusters), using the naive and the improved multiplexing schemes.

was 4). In the final step, there are 2 fusion attempts from which we get one copy of the final cluster state C^{2m}.

In general, in a level-l fusion step in Fig. 2-3, and with y_1 and y_2 copies in the respective banks of the two daughter clusters, the distribution of the number x of fused states C_i^l is, $B\left(x, \min\{y_1, y_2\}, S_{B(l)}^{im}\right)$, where $S_{B(l)}^{im} = \eta_{B(l)}^2 \left(\frac{1}{2}(\eta_s\eta_d)^2 + \frac{1}{4}(\eta_s\eta_d)^4\right)$ [36] and $\eta_{B(l)} = P_{\text{GHZ}} P_{\text{ff}}^{l+1}$ is the survival rate of photons up to before the l^{th} fusion step. The success probabilities of this scheme, P_{c1} (and P_{cn}) are calculated using Monte Carlo simulations.

In Fig. 2-5, we plot the probability P_{cn} of successfully building clusters C^k (with $k = 7$), simultaneously at $n = 250$ major nodes, for both schemes. n_B, n_{GHZ} and n_{meas} are optimized for the naive scheme to maximize P_{cn} for any given N_s. The plot clearly shows that the improved scheme leads to resource savings by a factor of $\sim 10^4$. We further observe that, for both schemes, P_{cn} undergoes a rapid percolation-like transition from zero to one as N_s is increased beyond a certain threshold value. P_{cn} is only a function of k, n, and N_s. We fix $P_{cn} = 0.9$ and calculate the corresponding minimum N_s required, for every value of k and n. This sharp-transition behavior of

P_{cn} allows us to conveniently split the problem of designing the repeater architecture into two parts:

(1) choosing an error-protection level by choosing m (number of parallel qubit channels) and $\vec{b}$ (the branching vector of the error protection trees), which gives us k (indicative of the total cluster size), and using this to calculate the key rate vs. distance achieved—both with n repeater stations, and also the resulting envelope over all n; and

(2) given the design choices (m and $\vec{b}$), calculating the number of photon sources N_s so as to achieve a close-to-unity P_{cn} (probability that all n nodes create the required clusters on every clock cycle), for a given value of k (cluster size at each repeater node), and n (the number of repeater nodes).

2.3.4 Measuring the clusters and connecting the chain

Once the clusters are created, the outer qubits are sent to minor nodes at the middle of the elementary links, as shown by the arrows in Fig. 2-2(c). The outer qubits are measured in the Bell basis at the minor nodes using ancilla-assisted boosted fusion gates [36]. The loss rate seen by the outer qubits is $\epsilon_{\text{trav}} \equiv 1 - \eta^{\frac{1}{2n}} P_{\text{ff}}^{k+2} P_0$ where $\eta^{1/2n}$ is the transmissivity of half of an elementary link (of range $L/2n$). All the physical qubits corresponding to the inner (logical) qubits are stored locally in a fiber bundle with the same attenuation as the communication fiber between the repeater stations. Due to the classical-communication delay, the core qubits see more loss than the outer qubits do, which we define as $\epsilon_{\text{stat}} = 1 - \eta^{\frac{1}{n}} P_{\text{ff}}^{k+2} P_{\text{ff(fib)}} P_0$. However, it is important to note that, just like in the architecture of [54, 55], this delay only leads to a latency in the scheme and does not affect the clock rate of the system.

When the result of the BSMs on the m qubit channels at the two neighboring minor nodes arrive back at a major node, the major node picks one successful qubit channel on either side (if none of the m BSMs were a success on any one of the sides, then that time period is an overall failure). The logical inner qubits corresponding to all the outer qubits that are not deemed part of the successful BSMs are removed from the cluster by measuring them in the Z basis [57] (note that this Z measurement is a

logical one, which benefits from the loss-protection trees). On the two logical qubits (one on either side) corresponding to the successful channels, X basis measurements are performed, which has an effect of extending the entanglement. Alice and Bob, simultaneous with the minor node BSMs, detect the m outer photons sent to them by the first and the last major node in the repeater chain, over links of length $L_0/2$, using one of two randomly-chosen mutually-unbiased bases. Assuming the clusters at all n repeater nodes were successfully created (which happens with probability P_{cn}), the conditional probability of generating an end-to-end entangled pair between Alice and Bob, in one clock cycle, is given by the probability that all $n-1$ minor nodes herald at least one successful BSM, and all the pruning logical X and Z measurements on the clusters at all n major nodes are successful, and Alice and Bob both obtain successful detects on at least one of the m qubit channels:

$$P_{\text{meas}} = P_Z^{2(m-1)n} P_X^{2n} \left[1-(1-P_B)^m\right]^{n-1} P_{\text{end}}^2, \tag{2.1}$$

where P_X and P_Z are the probabilities of successful X and Z basis measurements on the logical inner qubits, respectively. P_{end} is the probability that Alice (resp., Bob) obtains at least one successful detection in one of the m qubit channels.

We quantify the performance of the repeater architecture in terms of the number of shared secret bits generated per mode (i.e., per clock cycle per spatial channel, where m is the number of spatial channels employed). Since, the channel noise comprises only photon loss, the success probability divided by the number of spatial channels per attempt is the bits per mode generated by this scheme i.e. $S = P_{cn}P_{\text{meas}}/2m$. Note that the bits per mode is obtained by dividing by the number of spatial channels that is twice the number of qubit channels ($2m$).

2.3.5 Re-ordering measurements in the cluster-creation process

In this section, we explain why the X measurements required to attach trees for counterfactual error correction and the Y measurement required to create the "clique"

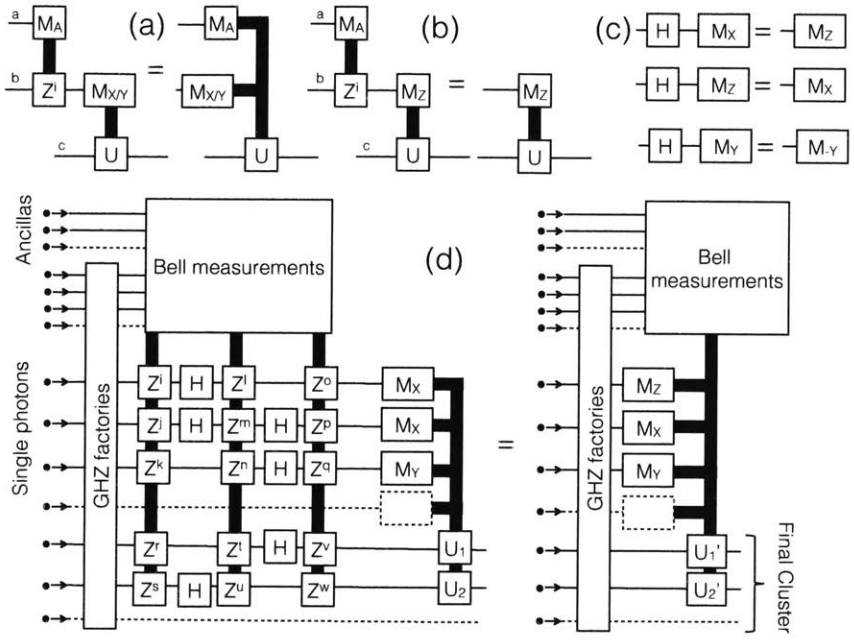

Figure 2-6: Explanation of why single qubit measurements can be applied before fusion operations. (a) X and Y basis measurements can be moved before conditional Z operators. (b) Z operators before Z basis measurements can be removed. (c) Hadamard gates followed by measurement in the X, Y or Z basis are equivalent to direct measurements in a different Pauli basis. (d) Single qubit measurements on the final cluster can be moved before fusion operations.

from the "star" cluster can be applied before the fusion operations. This makes the cluster creation process more efficient. The reordering of the operations is depicted in Fig. 2-6. Thin lines here represent photonic qubits, thick lines represent feed-forward operations, boxes labelled X, Y, Z, and H represent single qubit X, Y, Z rotations, and Hadamard gates respectively, and boxes labelled M_X, M_Y, and M_Z represent measurement in the X, Y, and Z bases, respectively.

First, we show some results regarding re-ordering of single qubit measurements and rotations. In the left side of Fig. 2-6(a), the unitary operation U on qubit c is conditioned on the result of an X or Y basis measurement on qubit b (that is determined beforehand). In addition, there is a conditional operation Z^i on the qubit b which depends on a feed-forward signal from a different part of the circuit, which in this case is the result of measurement M_A on qubit a. The application of a Z

gate before X or Y measurement simply has the effect of flipping the result of the measurement. Hence, the measurement $M_{X/Y}$ can be performed before M_A and the feed-forward result of M_A can simply be used to flip the result of $M_{X/Y}$ as shown on the right side of Fig. 2-6(a). The system in Fig. 2-6(b) is identical to the system in Fig. 2-6(a) except for the fact that measurement in the X/Y basis is replaced by measurement in the Z basis. Since application of a Z rotation does not influence the outcome of the Z measurement, the Z gate and the associated feed-forward can be removed entirely. In Fig. 2-6(c), we depict that a Hadamard gate followed by an X basis measurement is equivalent to a Z basis measurement, a Hadamard gate followed by a Z basis measurement is equivalent to an X basis measurement, and a Hadamard gate followed by a Y basis measurement is equivalent to a Y basis measurement with the result flipped.

We now use these results to show how measurements can be pushed earlier in the cluster creation process at the major nodes. The left side of Fig. 2-6(d) shows the system with measurements applied after the fusion operations. Single photons that are sent through GHZ factories to create 3-photon GHZ states, which are then fused using Bell measurements using ancilla photons. The surviving photons require some Hadamard and conditional Z rotations as part of the controlled-phase and parity-projection operations [58]. Finally, some of the surviving photons require X and Y basis measurements, the results of which are fed forward to photons in the final "clique" cluster. As shown in Fig. 2-6(a), (b) and (c), measurements in the Pauli basis can be pushed in front of Hadamard and conditional Z rotations by simply moving to a different Pauli basis or flipping the result of the measurement result. Hence, the system is equivalent to the right side of Fig. 2-6(d) in which single qubit Pauli measurements are applied before the fusion operation.

2.3.6 Rate calculations

In this section, we will evaluate the secret key bits generated per mode using the all-optical repeater architecture described above. We will first derive an exact analytical expression for the rate, given the design choices (m and $\vec{b}$), and compare it with

the best rate achievable without the use of quantum repeaters. A given choice of m and $\vec{b}$ determines k, which quantifies the size ($2^k + 2$ photons) of the cluster C^k generated at each major node on every clock cycle. Next, we will choose a value of k, and numerically optimize the choice of m and $\vec{b}$ (consistent with the chosen k) that maximizes the rate.

The probabilities of success for fault-tolerant X and Z measurements on one of the (logical) inner qubits at a major node cluster, P_X and P_Z, can be expressed in terms of the probabilities R_i of a successful 'indirect' Z measurement (as described in Section 2.2) on a qubit at the i-th level of the error-protection tree [51, 57]:

$$P_X = R_0 \tag{2.2}$$

$$P_Z = (1 - \epsilon_{\text{stat}} + \epsilon_{\text{stat}} R_1)^{b_0}, \tag{2.3}$$

where

$$R_i = 1 - \left[1 - (1 - \epsilon_{\text{stat}})(1 - \epsilon_{\text{stat}} + \epsilon_{\text{stat}} R_{i+2})^{b_{i+1}}\right]^{b_i}, \tag{2.4}$$

and $i \leq l$, $R_{l+1} = 0$, $b_{l+1} = 0$.

Let us assume a tree depth of $d = 2$, i.e., $\vec{b} = [b_0\ b_1]$, which is consistent with our numerical findings on the optimal branching vector as described later in this section (see Table 2.2). For a depth-2 branching vector, using Eq. 2.4, we find that $R_0 = 1 - \left[1 - \left(1 - \epsilon_{\text{stat}}\right)^{b_1+1}\right)\right]^{b_0}$ and $R_1 = 1 - \epsilon_{\text{stat}}^{b_1}$. Thus,

$$P_X = 1 - \left[1 - \left(\eta^{\frac{1}{n}}\right)^{b_1+1} B^{b_1+1}\right]^{b_0}, \text{ and} \tag{2.5}$$

$$P_Z = \left[1 - \left(1 - \eta^{\frac{1}{n}} B\right)^{b_1+1}\right]^{b_0}, \tag{2.6}$$

and the Bell measurement success probability becomes

$$P_B = \frac{AB^2}{m}\eta^{\frac{1}{n}}, \tag{2.7}$$

where $A = m\left(\frac{1}{2}(\eta_s\eta_d)^2 + \frac{1}{4}(\eta_s\eta_d)^4\right)/P_{\text{ff(fib)}}^2$, $B = P_{\text{ff}}^{k+2}P_{\text{ff(fib)}}P_0$.

The probability of at least one successful detection at Alice's (or Bob's) end is given by

$$P_{\text{end}} = 1 - \left(1 - \eta^{\frac{1}{2n}}C\right)^m, \tag{2.8}$$

where $C = P_{\text{ff}}^{k+2}P_0$.

We now have the bits-per-mode rate achievable with an n-repeater-node chain,

$$S_n(L) = \frac{P_{cn}}{2m}P_{\text{end}}^2 P_Z^{2(m-1)n}P_X^{2n}\left[1-(1-P_B)^m\right]^{n-1}, \tag{2.9}$$

with P_X, P_Z, P_B and P_{end} as given in Eqs. 2.5, 2.6, 2.7 and 2.8, with $\eta = e^{-\alpha L}$ the transmissivity of the end-of-end channel (of range L). See the magenta (dotted) lines in Fig. 2-7 for the plots of $S_n(L)$ as a function of L for a few chosen values of n. To obtain an achievable rate-distance envelope, we pick the transmissivity (resp., range) value given by

$$\eta_n = e^{-nz\ln(AB^2)} \tag{2.10}$$

on the rate-distance function $S_n(L)$, for each n, and evaluate the rate-loss function $S_{\text{LB}}(L)$ by computing a locus of the pair $(\eta_n, S_n(-\ln(\eta_n)/\alpha))$ over $n \in \{0,1,2,\ldots\}$, which by construction is a lower bound to the true envelope of the functions $S_n(L)$ over all n. The parameter z in Eq. 2.10 is numerically optimized to maximize the value of the lower-bound envelope $S_{\text{LB}}(L)$.

Let us evaluate P_X, P_Z, P_B and P_{end} at $\eta = \eta_n$ (i.e., substitute $\eta^{1/n} = (AB^2)^{-z}$ in the respective expressions), and define the following quantities:

$$p_X = 1 - \left[1 - \left(AB^2\right)^{-z(b_1+1)}B^{b_1+1}\right]^{b_0}, \text{ and} \tag{2.11}$$

$$p_Z = \left[1 - \left(1 - \left(AB^2\right)^{-z}B\right)^{b_1+1}\right]^{b_0}, \tag{2.12}$$

$$p_B = \frac{1}{m}\left(AB^2\right)^{1-z}, \text{ and} \tag{2.13}$$

$$p_{\text{end}} = 1 - \left(1 - \left(AB^2\right)^{-z/2} C\right)^m, \tag{2.14}$$

using which let us define the following: $C_1 = p_Z^{2(m-1)} p_X^2$, $C_2 = 1 - (1 - p_B)^m$, and $C_3 = p_{\text{end}}^2$. We can now express $S_n(L)$ evaluated at $L = L_n \equiv -\ln(\eta_n)/\alpha$ as $S_n(L_n) = (C_1 C_2)^n \frac{C_3 P_{cn}}{2mC_2}$. To obtain the locus of $(\eta_n, S_n(-\ln(\eta_n)/\alpha))$ over all n, we eliminate n from $S_n(L_n)$, and the expression of η_n in Eq. 2.10, and obtain the following rate-loss envelope lower bound (expressed as a function of $\eta = e^{-\alpha L}$):

$$S_{\text{LB}}(\eta) = D\eta^s, \tag{2.15}$$

where $D = \frac{C_3 P_{cn}}{2mC_2}$ and the exponent $s = -\frac{\ln(C_1 C_2)}{z \ln(AB^2)}$. We now numerically optimize the choice of the parameter z such that the value of the exponent s is minimized (note that C_1, C_2 and C_3 are functions of z). The repeater scheme, even when built with lossy components, can achieve $s < 1$. This establishes that the rate achieved by this repeater architecture scales better with the end-to-end channel loss than the best possible scaling achievable without the use of quantum repeaters. The rate corresponding to the PLOB bound, an upper bound to the best possible key rate achievable without repeaters, scales as $\propto \eta$, i.e., achieves an exponent $s = 1$. The value of the exponent s achievable by the repeater scheme can be improved (lowered) by enhancing the level of error correction (i.e., choosing a larger $\vec{b}$). Doing so increases the size of the photonic clusters needed at the repeater nodes, and hence increases the number of photon sources N_s required locally at each node.

In Fig. 2-7, we plot $S_n(L)$ (bits per mode) as a function of L (in km) for $n = 1, 10, 24, 56, 133$, and 314 (magenta dotted plots), with $\vec{b} = \{7, 3\}$ trees and $m = 4$ parallel qubit channels (which translates to $k = 8$ clusters at each major node), and other device parameters as summarized in Table 2.1. We also show the analytical rate-envelope lower bound $S_{\text{LB}}(L)$ (black solid line), obtained by choosing the optimal z numerically. For the numbers chosen, we get $D = 0.11$ and $s = 0.37$. Our lower bound $S_{\text{LB}}(L)$ is seen to be remarkably close to the true rate-distance envelope, and

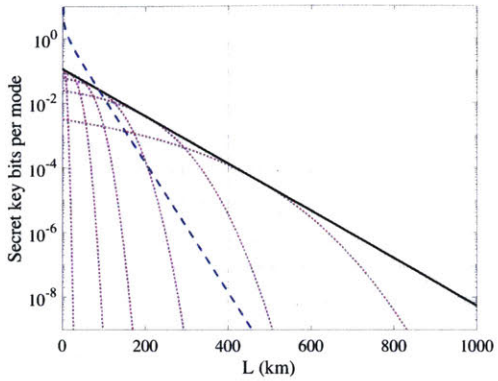

Figure 2-7: The key rate (in bits per mode) $S_n(L)$ achieved by an n-node repeater chain shown as a function of range L, for $n = 1, 10, 24, 56, 133$, and 314 (magenta dotted plots). Our analytical lower bound to the rate-distance envelope $S_{\mathrm{LB}}(L)$ (black solid plot) is indistinguishable for the numerically-obtained rate-distance envelope, which are very close to one another. For all the rate-distance plots, we choose $m = 4$ parallel channels and $\vec{b} = \{7, 3\}$ error-protection trees (which translates to $k = 8$ clusters). The PLOB bound is shown for comparison (blue dashed plot).

the bound remains an excellent approximation to the rate-distance envelope for all parameters we tried. For the plots in Fig. 2-7, the maximum deviation of $S_{\mathrm{LB}}(L)$ from the envelope is 0.09% (at $n = 1$), and the deviation is less than 0.005% for $L > 32$ km. This close agreement of the analytical lower bound to the true envelope suggests that the ansatz in Eq. 2.10 is a very good approximation.

One implication of $\eta_n = (AB^2)^{-nz}$ lying on the rate-distance envelope is that the distance between each repeater (major) node,

$$L_0 \equiv \frac{L}{n} = \frac{\ln(AB^2)}{\alpha} \tag{2.16}$$

is a constant and independent of the total range L. In other words, given the device parameters and the choice of the major-node cluster size (i.e., m and $\vec{b}$), there is an optimal gap with which repeaters should be placed—no more, and no less. For the numbers used for the plots in Fig. 2-7, $L_0 = 1.49$ km. Fig. 2-7 also shows the PLOB bound for comparison (blue dashed plot), which the repeater scheme is seen to outperform beyond a range of 87 km.

Device parameter	symbol	value
fiber loss coefficient	α	$0.046\ \mathrm{km}^{-1}$ (0.2 dB/km)
on-chip loss coefficient	β	$0.62\ \mathrm{m}^{-1}$ (2.7 dB/m)
feed-forward time in fiber	τ_f	102.85 ns
feed-forward time on-chip	τ_s	20 ps
chip to fiber coupling efficiency	P_c	0.99
source detector efficiency product	$\eta_s\eta_d$	0.99
speed of light in fiber	c_f	$2 \times 10^8 m/s$
speed of light on chip	c_{ch}	$7.6 \times 10^7 m/s$

Table 2.1: Assumed values for device performance parameters. The source detector efficiency product $\eta_s\eta_d$ is sufficient for the purposes of the calculations in this chapter, and need not be specified separately. Recall that $P_{\text{ff}} = e^{-\beta\tau_s c_{\text{ch}}}$, $P_{\text{ff(fib)}} = e^{-\alpha\tau_f c_{\text{f}}}$, and $P_0 = \eta_s\eta_d P_c/(2-\eta_s\eta_d)$. For computational ease, we choose τ_f such that $P_{\text{ff}} = P_{\text{ff(fib)}}$. The device parameters are aggressive but may be achievable in the near future given rapid advances in integrated photonics.

2.3.7 Discussion

In this section, we will go back to the all-photonic repeater architecture proposed by Azuma *et al.* [51], and discuss the main modifications (improvements) we considered in the architecture we described and analyzed above. We will also show a comparative study of the resource requirements and rate performance of the naive scheme and our modified scheme. Following are the salient differences between the architecture we analyzed above, and the one proposed in [51].

Retaining vs. transmitting the clusters—In the proposal of [51], all the logical inner qubits, along with the outer qubits (i.e., all the N photons of the cluster at a major node) are sent to the minor node, whereas we store the inner qubit photons in a fiber spool locally at the major nodes. The former has an advantage that no classical communication needs to happen from minor nodes back to major nodes before the logical X and logical Z measurements are done to the logical inner qubits, since all those qubits are present locally at the minor nodes when the BSMs are performed there on outer-qubit pairs from neighboring major node clusters. The advantage of

our (latter) scheme is that the number of parallel physical channels needed ($2m$) is much smaller as compared to the number needed (N) for the scheme in [51]. For the numbers in Fig. 2-7, that is 8 as opposed to 208 parallel fiber channels connecting successive repeater nodes.

Difference in the bits-per-mode rate—Further, the bits per mode achieved by the architecture in [51] would be given by $P_{cn}P_{\text{meas}}/N$, whereas the bits per mode achieved by our modified architecture would be $P_{cn}P_{\text{meas}}/2m$. The P_{meas} of the former is higher (due to lower loss incurred by the photons of the logical inner qubits of the clusters as they do not need to wait in a lossy fiber spool while waiting for the classical information to fly back from the minor nodes). However, the other improvements described below more than compensate for the better P_{meas}, and the latter scheme achieves a far better bits-per-mode performance (see Fig. 2-8).

Linear optic vs. boosted linear optic fusion gates—We propose the use of the improved Bell-state measurement scheme of Ewert *et al.* [36] that inject four single photons to boost the success probability of the fusion-II gate. Our calculations show that the cost of using these additional ancilla photons is far outweighed by the effect of the improved success probability, in the performance of the repeater architecture, despite assuming lossy sources and detectors.

Improved multiplexing scheme for cluster generation—We use an improved multiplexing scheme to create the clusters at the major nodes, as described in Section 2.3.3 and depicted in Fig. 2-4(b). Previous studies have estimated the resource requirements for cluster generation based on the average number of attempts required for each probabilistic steps [2, 58]. However, in order to generate the required cluster at every repeater station on every clock cycle with high probability, the resources required at each repeater station need to be greater than the number that would allow for cluster creation "on average". To our knowledge, this is the first study that explicitly looks at how probabilistic operations need to be multiplexed in a real system.

Pushing the measurements ahead during cluster creation—The single qubit measurements that do not depend on the outcomes of Bell measurements at the minor

nodes, are performed before the fusion operations, directly on the photons of the GHZ states, very early during the cluster creation process.

Let us now see what the above modifications to the architecture do to the rate performance. The bits-per-mode rates for the naive and the improved schemes are plotted in Fig. 2-8(a) and (b), respectively. We assume device loss parameters as listed in Table 2.1 for both sets of plots. In each plot, we compute the rate-distance performance (envelopes taken over n, the number of repeater nodes) for four different error-protection levels (i.e., $k = 7, 8, 9$, and 10). For every point on each rate-distance envelope, m and $\vec{b}$ are optimally chosen (consistent with the given k). Each rate-distance plot exhibits the $D\eta^s = De^{-s\alpha L}$ behavior, and the exponent s diminishes as a higher k is chosen. For the naive scheme, the minimum k for which the repeater can beat the PLOB bound (pink-dashed line) is $k = 8$ and the optimized clusters at the major nodes have 192 photons each. Hence, the scheme would require 208 parallel fiber links connecting successive nodes. In comparison, in the improved scheme, $k = 7$ is sufficient to beat the PLOB bound, and requires $2m = 8$ parallel fiber links. The optimal tree depth, for this $k = 7$ rate plot is found to be $d = 2$, which is consistent with the analytical development in Section 2.3.6.

Table 2.2 lists, at a range of $L = 300$ km, and for each of the cases ($k = 7, 8, 9, 10$), the optimal values of m for the naive (m_{naive}) and new schemes (m_{new}), the optimal branching vector for the naive ($\vec{b}_{\text{naive}}$) and new schemes ($\vec{b}_{\text{new}}$), and the number of parallel fiber links needed in the naive scheme (N_{naive}). In the case of the new scheme, the number of parallel fiber links needed is simply $2m_{\text{new}}$.

k	m_{naive}	N_{naive}	$\vec{b}_{\text{naive}}$	m_{new}	$\vec{b}_{\text{new}}$
7	5	100	$\{3,2\}$	4	$\{4,2\}$
8	8	208	$\{4,2\}$	5	$\{5,3\}$
9	11	462	$\{5,3\}$	6	$\{7,4\}$
10	12	864	$\{7,4\}$	8	$\{10,5\}$

Table 2.2: For $k = 7, 8, 9$, and 10, at $L = 300$ km range, m_{naive} and m_{new} are the optimal values of m for the naive and new schemes respectively. $\vec{b}_{\text{naive}}$ and $\vec{b}_{\text{new}}$ are the optimal values of $\vec{b}$ for the naive and new schemes respectively. N_{naive} is the corresponding number of parallel fiber links needed between successive repeater nodes in the naive scheme. For the new scheme, the number of parallel links is $2m_{\text{new}}$.

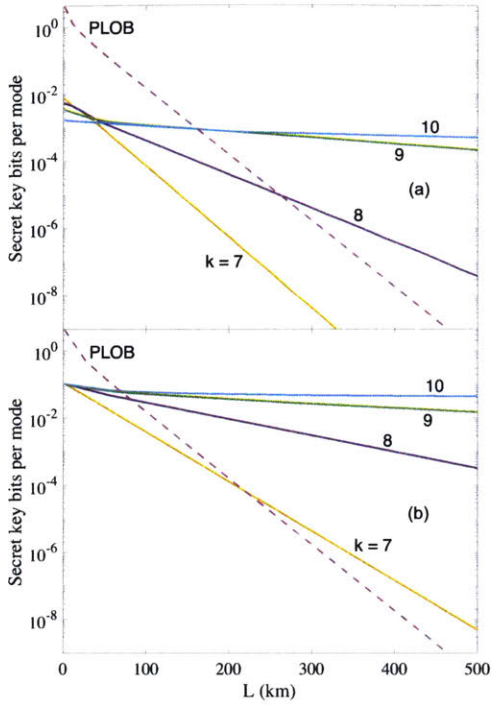

Figure 2-8: Scaling of the bits per mode as a function of distance L for different numbers of fusion steps k for the (a) naive scheme and (b) with the improvements of this chapter. The PLOB bound is the pink dashed line. $Q = 2^k + 2$ is the total number of photons in the cluster at each repeater station.

Let us now compare the resources (number of photons, N_s) required to build the major node clusters, for the respective cases that can (barely) beat the PLOB bound. The naive scheme requires 1.9×10^{11} photon sources at each major node, while the new scheme requires 3.3×10^6 sources, an improvement of 5 orders of magnitude (see Fig. 2-5). It is also interesting to note that if the primitive resources were 3-photon GHZ sources rather than single photon sources, 15 thousand GHZ sources would be required, a relatively smaller number.

Given the size of the earth, for terrestrial long distance communications, it is useful to quantify the performance of our (improved) all-optical repeater scheme at say 5000 km. Without quantum repeaters, the best QKD protocol realized with ideal devices cannot exceed a key rate of 2.9×10^{-99} bits per mode at this distance. Our all-optical repeater scheme, with 954-photon clusters ($k = 10$) at each repeater node

can attain a key rate of 8×10^{-3} bits per mode using $2m = 18$ parallel channels and $n = 12411$ repeater nodes, which translates to a 144 kHz key generation rate assuming a 1 MHz repetition rate. If we employed 518-photon clusters ($k = 9$) instead, the rate achieved would only be 4×10^{-8} bits per mode using $2m = 14$ parallel channels and $n = 12255$ repeater nodes. The number of photon sources required at a repeater node to create the required clusters (using linear optics) for the above two example cluster-size constraints are 1.2×10^8 and 3.6×10^7, respectively.

In the presence of losses in the waveguide, there is a maximum sustainable size of the clusters at the major nodes, at least for the error protection methods described in this chapter. A larger cluster requires a greater creation time and hence, each photon in the cluster sees a larger effective loss rate (stemming from the P_{ff}^k term in ϵ_{trav} and ϵ_{stat}). Since the error correction scheme has a maximum loss tolerance of 50%, there is a maximum size of the clusters that can be created and thus a maximum level of error protection that a qubit can have. So, given a set of device losses, increasing the error protection level (viz., k) cannot indefinitely improve the rate performance.

The aforesaid detrimental effect of loss with an increasing cluster size has more serious implications for linear optical cluster state quantum computing (CSQC) in general, using only tree-based counterfactual error correction [57]. This is because a polynomial scaling of the number of photon sources (with the size of the cluster) is required in the asymptotic limit for the CSQC scheme to be scalable. The failure probability of every qubit needs to decrease exponentially with the size of the computation. Hence, the level of protection of each qubit must increase with the size of the problem, which implies a greater cluster creation time and hence a greater loss rate. Since there is a 50% ceiling on the tolerable photon loss, it is not possible to achieve the required level of protection for arbitrarily large computations, as discussed above for the case of an all-photon quantum repeater. Developing a scalable method for creation of arbitrarily large clusters in constant time would solve this problem and will also allow for a polynomial scaling of the number of photons with computation size. A recent paper proposes using counterfactual error correction to fault-tolerantly create surface code data qubits [58]. However, the resource requirements for this

scheme are found to be extremely high.

2.4 One-way repeater scheme based on logical Bell measurements

In this section, we present a one-way repeater scheme based on the quantum parity code. The scheme is called a *one-way* scheme because it only requires classical communication from the repeater stations to Bob to generate shared entanglement, unlike the previous scheme which required two way communication between adjacent repeater nodes. However, it should be noted that a QKD scheme using this shared entanglement would still require two-way communication between Alice and Bob.

2.4.1 One-way quantum repeaters

A one way quantum repeater uses teleportation based error correction at each repeater station, as opposed to multiplexing which was used in the previous section. The schematic of one way repeater scheme is shown in Fig. 2-9.

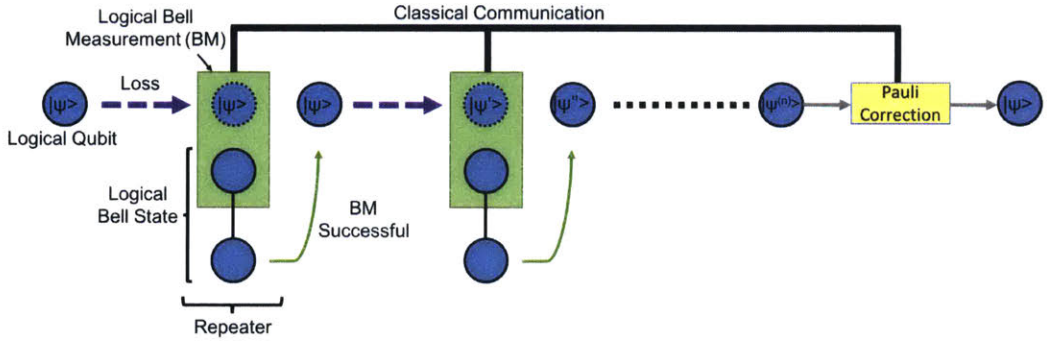

Figure 2-9: Schematic of the one-way repeater scheme

The blue circles in Fig. 2-9 represent logical qubits, which in this chapter will be qubits encoded in the quantum parity code, which is explained later. A repeater station in a one-way scheme cleans an incoming qubit i.e., it takes an imperfect logical input state and sends out a near perfect logical output state with a probability that

depends on the error rate and size of code. However, if the errors are below a certain threshold, with a large enough code, the success probability can be close to one. The repeater station consists of logical Bell pairs (depicted in Fig. 2-9 with two blue circles connected by a line), also in the quantum parity code. The imperfect input state and one half of the logical Bell measurement undergo a logical Bell measurement at the repeater and if the Bell measurement is successful, the second half of the logical Bell pair, which is the output of the repeater, is projected onto the input state with an additional Pauli rotation that depends on the result of the Bell measurement. Furthermore, if the logical Bell pair at the repeater was perfect, the output state is perfect when the Bell measurement succeeds. In the case of pure loss, the output state's transmissivity depends on the transmissivity of the logical Bell pair, which is generally higher that the transmissivity of the input state.

2.4.2 Quantum parity code and logical Bell measurement

The quantum parity code (QPC) [59, 60] is a generalization of Bacon-Shor code [61, 62], with encoded 0 and 1 states of the form $|0\rangle^{(p,q)} = (|+\rangle^{(p,q)} + |-\rangle^{(p,q)})/\sqrt{2}$ and $|1\rangle^{(p,q)} = (|+\rangle^{(p,q)} - |-\rangle^{(p,q)})/\sqrt{2}$ where $|\pm\rangle^{(p,q)} = \left(\frac{|0\rangle_L^{\otimes q} \pm |1\rangle_L^{\otimes q}}{\sqrt{2}}\right)^{\otimes p}$ where the subscript L signifies the qubit representation (as opposed to Fock space representation).

In addition to providing tolerance against loss (the quantum parity code also corrects for bit and phase flip, although we focus on loss in this chapter), the linear optic Bell measurement success probability on QPC encoded qubits goes as $1 - 1/2^p$ [52] in the absence of loss i.e., it approaches one as the size of the code increases. For dual rail qubits, in contrast, the Bell measurement success probability is limited to $1/2$ in the absence of ancilla inputs and even with single photon ancillas, the best known success probability is $25/32$ [36].

In this section, we outline the derivation of the QPC logical Bell measurement success probability when the states are lossy which was first presented in [52]. The two QPC qubits undergoing Bell measurement have different loss, which we label as γ_1(for the qubit which is half of the Bell pair created at the same repeater station) and γ_2(for the qubit received at the repeater station that has gone through fiber loss).

The dual rail qubits that constitute every QPC qubit undergo pairwise linear optic Bell measurement.

The (p, q) QPC can be visualized as having p blocks of q photons each. For the logical Bell measurement to succeed, every block must have at least one surviving pair. If this condition is met and there are l perfect blocks (without any lost photons), the success probability is $1 - 2^{-l}$. We calculate the Bell measurement success probability in terms of $\gamma = \gamma_1\gamma_2$:

- The probability of there being atleast one intact pair in a block is $1 - (1 - \gamma)^q$.
- The probability of the whole block being intact is $(\gamma)^q$.
- The probability of l blocks being perfect and the rest having 1 - $(q-1)$ photons is $\binom{l}{k}[\gamma^q]^k[1 - \gamma^q - (1 - \gamma)^q]^{p-l}$.
- When l blocks are perfect and every block has at least 1 pair intact, Bell measurement success probability is $1 - 2^{-l}$.

Hence, probability of successful Bell measurement is

$$\begin{aligned} P_B^{(p,q)} &= \sum_{l=0}^{p}(1 - 2^{-l})\binom{p}{l}[\gamma^q]^l[1 - \gamma^q - (1 - \gamma)^q]^{p-l} \\ &= [1 - (1 - \gamma)^q]^p - [1 - \gamma^q/2 - (1 - \gamma)^q]^p \end{aligned} \tag{2.17}$$

2.4.3 Key rate

.

With n repeater stations between Alice and Bob, the probability that every Bell measurement in the chain succeeds is $\left[P_B^{(p,q)}\right]^n$. Furthermore, the measurement of the single QPC qubit at Bob's end should succeed, which requires that all blocks have at least one photon and at least one block is perfect. Hence, the probability of the measurement at Bob's end succeeding is $P_{end} = [1 - (1 - \gamma_2)^q]^p - [1 - \gamma_2^q - (1 - \gamma_2)^q]^p$.

Suppose that each photon has a loss P_0 due to detector, source, fiber to chip coupling efficiency, and goes through k feed-forward steps which each have loss P_{ff}. The loss seen at the time of creation of the Bell pairs is $\gamma_1 = P_0 P_{\text{ff}}^k$. The loss seen after passing this state through fiber is $\gamma_2 = P_0 P_{\text{ff}}^k \eta^{\frac{1}{n+1}}$.

Based on our method for creating the logical Bell states (see 2.4.4) $k = \lceil \log_2(q+1-2) \rceil + \lceil \log_2(p) \rceil + 4$. The probability of successfully obtaining a Bell pair is $P_{end} \left[P_B^{(p,q)} \right]^n$ which gives us the bits per mode rate

$$S = P_{end} \left[P_B^{(p,q)} \right]^n / 2pq \tag{2.18}$$

where the factor of 2 in the denominator comes from the fact that every dual rail qubit requires two modes. Using the device parameters from Table 2.1, $P_0 = \eta_s \eta_d P_c / (2 - \eta_s \eta_d) = 0.9607$ and $P_{\text{ff}} = e^{-\beta \tau_s c_{\text{ch}}} = 0.9991$.

In Fig. 2-10, we plot the secret key rate for our repeater scheme with the $(p, q) = (12, 4)$ QPC for different values of n.

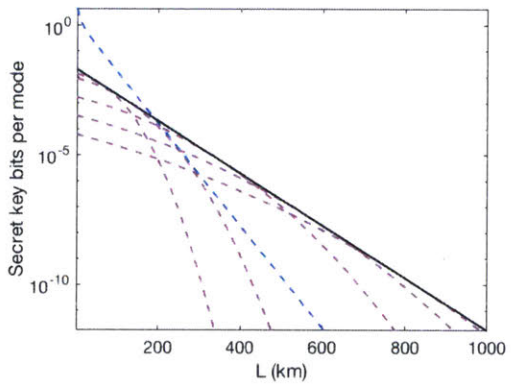

Figure 2-10: The key rate (in bits per mode) S achieved by an n-node repeater chain using the $(p, q) = (12, 4)$ QPC shown as a function of range L, for $n = 50, 100, 300, 500$ and 700. The envelope (optimizing over n) is shown in black. The PLOB bound is shown for comparison (blue dashed plot).

In Fig. 2-11, the envelope of the key rate, optimizing over n at every point, with different size of QPC codes is plotted. As in the previous section, for a fixed code size, the key rate drops exponentially with distance, but the exponent can be shallower than the PLOB bound. The smallest code that beats the PLOB bound is $(p, q) = (8, 3)$

which is a 24 photon code. However, the crossover with the PLOB bound for this code happens at a key rate of $\approx 10^{-20}$, which may have limited practical utility. The $(9, 3)$ QPC beats the PLOB bound at a more reasonable point. Going to larger code gives a better scaling of the rate with distance but this comes at the cost of a larger entangled state required at every repeater station. In the next section, we look at the resources required to create these entangled states at every repeater station.

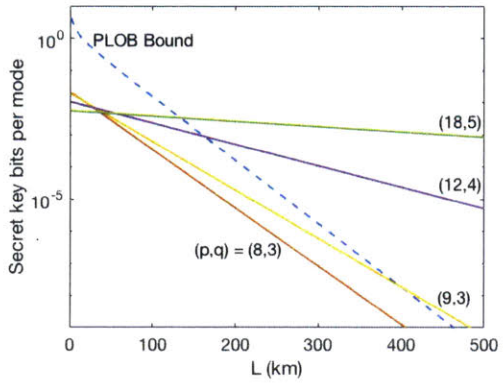

Figure 2-11: The key rate (in bits per mode) achieved for different sizes of the QPC, optimizing over n at every point. The PLOB bound is shown for comparison (blue dashed plot).

2.4.4 Creation of quantum parity encoded Bell pairs

In this section, we present a loss-tolerant scheme for the preparation of Bell pairs encoded in the quantum parity code from single photons. These logical Bell pairs are required at every repeater station at every clock cycle. The previously known scheme for the preparation of such a logical Bell pair required an linear optic entangling operation known as fusion-I and a KLM style CNOT gate [60], which are very expensive in terms of resource requirements and are also very sensitive to photon loss. In this section we present a scheme that is loss-tolerant because it uses only the fusion-II operation which is loss tolerant. Our objective is to obtain a state of the form $\left(|0\rangle^{(p,q)}|0\rangle^{(q,p)} + |1\rangle^{(p,q)}|1\rangle^{(p,q)}\right)/\sqrt{2}$ where $|\pm\rangle^{(p,q)} = \left(\frac{|0\rangle^{\otimes q} \pm |1\rangle^{\otimes q}}{\sqrt{2}}\right)^{\otimes p}$.

- Groups of six single photons are sent through GHZ factories which give three

photon GHZ states with probability $[\eta_s\eta_d(2-\eta_s\eta_d)]^3/32$. We write a three-photon GHZ state in the form

$$|+\rangle^{(3)} = \left(|0\rangle^{\otimes 3} + |1\rangle^{\otimes 3}\right)/\sqrt{2} \tag{2.19}$$

- A fusion-II (fII) gate acting on two qubits of GHZ states of size q_1 and q_2 results in a GHZ state with $q_1 + q_2 - 2$ photons. Starting with three- photon GHZstates, the fusion-II operation is used to grow the GHZ state to the size $q+1$, which can be written in the form

$$\begin{aligned}
|+\rangle^{(q+1)} &= \left(|0\rangle^{\otimes q+1} + |1\rangle^{\otimes q+1}\right)/\sqrt{2} \\
&= \left(|0\rangle^{\otimes q}(|+\rangle + |-\rangle) + |1\rangle^{\otimes q}(|+\rangle - |-\rangle)\right)/2 \\
&= \left(|+\rangle\left(|0\rangle^{\otimes q} + |1\rangle^{\otimes q}\right) + |-\rangle\left(|0\rangle^{\otimes q} - |1\rangle^{\otimes q}\right)\right)/2 \\
|+\rangle^{(q+1)} &= \left(|+\rangle|+\rangle^{(q)} + |-\rangle|-\rangle^{(q)}\right)/\sqrt{2}
\end{aligned} \tag{2.20}$$

Up to this point, the process is the same as Hayes et al. [60] From this point, we modify the protocol. In the following steps, we repeatedly make use of the entanglement swapping procedure

$$\frac{(|+\rangle_{b1}|+\rangle_{l1} + |-\rangle_{b1}|-\rangle_{l1})}{\sqrt{2}}\frac{(|+\rangle_{b2}|+\rangle_{l2} + |-\rangle_{b2}|-\rangle_{l2})}{\sqrt{2}} \rightarrow |+\rangle_{l1}|+\rangle_{l2} + |+\rangle_{l1}|+\rangle_{l2} \tag{2.21}$$

which requires Bell measurement on qubits b1 and b2 and logical Z and X operations on l1 and l2.

$$\begin{aligned}
&(|+\rangle_{b1}|+\rangle_{l1} + |-\rangle_{b1}|-\rangle_{l1})(|+\rangle_{b2}|+\rangle_{l1} + |-\rangle_{b2}|-\rangle_{l2})/2 \\
= \; &(|+\rangle_{b1}|+\rangle_{l1}|+\rangle_{b2}|+\rangle_{l2} + |-\rangle_{b1}|-\rangle_{l1}|-\rangle_{b2}|-\rangle_{l2}
\end{aligned} \tag{2.22}$$

$$+|+\rangle_{b1}|-\rangle_{l1}|+\rangle_{b2}|-\rangle_{l2}+|-\rangle_{b1}|+\rangle_{l1}|-\rangle_{b2}|+\rangle_{l2})/2$$

Rearranging

$$\begin{aligned}
&= (|+\rangle_{b1}|+\rangle_{b2}|+\rangle_{l1}|+\rangle_{l2}+|-\rangle_{b1}|-\rangle_{b2}|-\rangle_{l1}|-\rangle_{l2}+ \\
&\quad |+\rangle_{b1}|-\rangle_{b2}|+\rangle_{l1}|-\rangle_{l2}+|-\rangle_{b1}|+\rangle_{b2}|-\rangle_{l1}|+\rangle_{l2})/2 \\
&= (|\Psi_+\rangle+|\Psi_-\rangle)(|+\rangle_{l1}|+\rangle_{l2})/2\sqrt{2}+(|\Psi_+\rangle-|\Psi_-\rangle)(|-\rangle_{l1}|-\rangle_{l2})/2\sqrt{2}+ \\
&\quad (|\Phi_+\rangle+|\Phi_-\rangle)(|+\rangle_{l1}|-\rangle_{l2})/2\sqrt{2}+(|\Phi_+\rangle-|\Phi_-\rangle)(|-\rangle_{l1}|+\rangle_{l2})/2\sqrt{2} \qquad (2.23) \\
&= |\Psi_+\rangle(|+\rangle_{l1}|+\rangle_{l2}+|-\rangle_{l1}|-\rangle_{l2})/2\sqrt{2}+|\Psi_-\rangle(|+\rangle_{l1}|+\rangle_{l2}-|-\rangle_{l1}|-\rangle_{l2})/2\sqrt{2}+ \\
&\quad |\Phi_+\rangle(|+\rangle_{l1}|-\rangle_{l2}+|-\rangle_{l1}|+\rangle_{l2})/2\sqrt{2}+|\Phi_+\rangle(|+\rangle_{l1}|-\rangle_{l2}-|-\rangle_{l1}|+\rangle_{l2})/2\sqrt{2}
\end{aligned}$$

where

$$|\Psi_+\rangle = (|+\rangle_{b1}|+\rangle_{b2}+|-\rangle_{b1}|-\rangle_{b2})/\sqrt{2} \qquad (2.24)$$

$$|\Psi_-\rangle = (|+\rangle_{b1}|+\rangle_{b2}-|-\rangle_{b1}|-\rangle_{b2})/\sqrt{2} \qquad (2.25)$$

$$|\Phi_+\rangle = (|+\rangle_{b1}|-\rangle_{b2}+|-\rangle_{b1}|+\rangle_{b2})/\sqrt{2} \qquad (2.26)$$

$$|\Phi_-\rangle = (|+\rangle_{b1}|-\rangle_{b2}-|-\rangle_{b1}|+\rangle_{b2})/\sqrt{2} \qquad (2.27)$$

Hence, Bell measurement on b1, b2 followed by conditional X and Z operations on $l1$, $l2$ give the required state $|+\rangle_{l1}|+\rangle_{l2}+|+\rangle_{l1}|+\rangle_{l2}$

- While the state $|+\rangle^{(q+1)}$ is being created, some 3 photon GHZ states from Eq. 2.19 undergo single qubit rotations to give

$$|anc\rangle = \left(|+\rangle^{\otimes 3}+|-\rangle^{\otimes 3}\right)/\sqrt{2} \qquad (2.28)$$

- Combining states from Eq. 2.20 and 2.28 using Eq. 2.21 with the leftmost single qubit in each state as $b1$, $b2$, we obtain the state

$$|0\rangle_a^{(1,q)} = \left(|+\rangle^{\otimes 2}|+\rangle^{(1,q)} + |-\rangle^{\otimes 2}|-\rangle^{(1,q)}\right)/\sqrt{2} \tag{2.29}$$

where the subscript a signifies that the state has 2 "extra" qubits attached to it.

- By using Eq. 2.21 on 2 states of the form $|0\rangle_a^{(p_1,q)}$ and $|0\rangle_a^{(p_2,q)}$ with $b1$ being one of the two "extra" qubits of the state $|0\rangle_a^{(p_1,q)}$ and $b2$ being one of the two "extra" qubits of the state $|0\rangle_a^{(p_2,q)}$, it is possible to grow the second index of the logical code i.e

$$|0\rangle_a^{(p_1,q)}|0\rangle_a^{(p_2,q)} \rightarrow |0\rangle_a^{(p_1+p_2,q)} \tag{2.30}$$

- Once states $|0\rangle_a^{(p,q)}$ have been obtained, a measurement on one of the extra qubits in the 0, 1 basis with single qubit rotation reduces the state to

$$\left(|+\rangle|+\rangle^{(p,q)} + |-\rangle|-\rangle^{(p,q)}\right)/\sqrt{2} \tag{2.31}$$

- Taking two of these states and using 2.21 with b1, b2 being the remaining extra qubit, we get the required logical Bell state

$$\frac{\left(|+\rangle^{(p,q)}|+\rangle^{(p,q)} + |-\rangle^{(p,q)}|-\rangle^{(p,q)}\right)}{\sqrt{2}} = \frac{\left(|0\rangle^{(p,q)}|0\rangle^{(p,q)} + |1\rangle^{(p,q)}|1\rangle^{(p,q)}\right)}{\sqrt{2}} \tag{2.32}$$

Following the same improved multiplexing scheme for the probabilistic steps as 2.3.3, we calculate the number of single photon sources required at every repeater station to create the QPC encoded Bell pairs on every clock cycle. In Fig. 2-12, we plot the probability of successfully creating repeater states at $n = 300$ repeaters simultaneously as a function of the number of sources at every repeater station for the

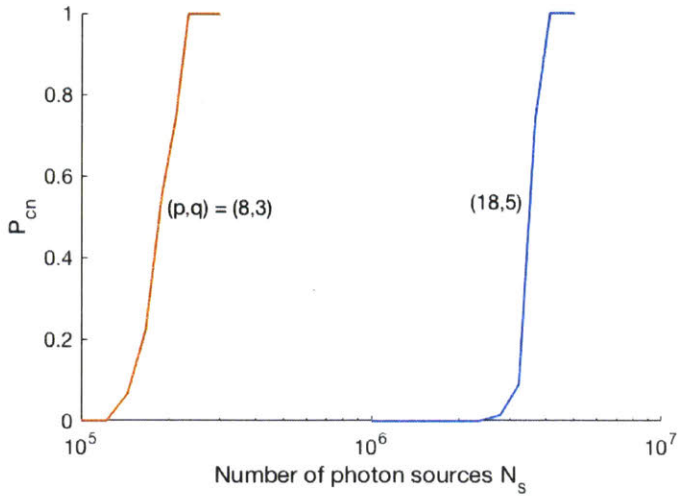

Figure 2-12: The probability that all $n = 300$ major nodes are simultaneously successful in creating encoded Bell pairs in quantum parity code of size $(p, q) = (8, 3)$ and $(18, 5)$

$(p, q) = (8, 3)$ and $(18, 5)$ QPC Bell pairs. Similar to Fig. 2-5, the probability of successfully creating the repeater state undergoes a sudden transition.

In Table 2.3, we summarize the number of number of photon sources N_S at each repeater required to create QPC (p, q) states (which have size N_q) across 300 repeater stations with 90% success probability. Furthermore, we also consider a scenario in which we start with sources 3-photon GHZ states, and evaluate the number of such sources, N_{GHZ}, required at every repeater station. The loss rate of each GHZ source is the same as obtained from linear optics using sources of the same efficiency as Table 2.1.

(p, q)	N_q	N_S	N_{GHZ}
$(8, 3)$	48	2×10^5	1000
$(9, 3)$	54	7×10^5	3500
$(12, 4)$	96	2×10^6	10000
$(18, 5)$	180	4×10^6	22000

Table 2.3: Number of photon sources N_S required at each repeater to create QPC (p, q) states across 300 repeaters with 90% probability. N_q is the number of qubits in the QPC Bell pair. Alternatively N_{GHZ} 3-photon GHZ states can be used. The loss rate of each GHZ source is the same as obtained from linear optics using sources of the same efficiency as table 2.1.

We find that the resource requirements with the third generation are an order of magnitude lower than the repeater in section 2.3. Furthermore, the requirements are significantly reduced if repeater stations are equipped with sources of small entangled states like 3-photon GHZ states.

However, as in section 2.3, the multiplexing based creation of entangled photonic states incurs a major overhead. Better schemes for entangled state creation, like those based on percolation will likely be needed, which are discussed in the next chapter.

2.5 Conclusions and open problems

In conclusion, we have performed a rigorous analysis of the resource requirements, and the achievable secret key rates of an all-optical repeater scheme that improves upon two recent proposal [51, 52], while taking into account all the losses in the system. While the all-optical repeater proposals of [51] and [52] present important conceptual advancements, we show that it may not be practically feasible given their astronomical resource requirements, both in terms of the number of photon sources and detectors needed at each repeater node, as well as the number of parallel optical fiber channels that must connect successive repeater nodes.

Our work improves the practicality in both of the aforementioned metrics, as well as the actual rate-vs.-distance performance achieved. Our improvements to the multiplexing based repeater reduce the number of photon sources required at each node by 5 orders of magnitude, and the number of parallel channels between repeater nodes required to beat the performance of a direct-transmission QKD scheme is brought down from more than two hundred, to 8. Furthermore, we find that moving to the one-way scheme further brings down the requirements by one order of magnitude. We also find that the resource requirements can be reduced by two more orders of magnitude if a source of 3-photon GHZ states is available.

These results suggest that further theoretical improvements on quantum photonic fault tolerant schemes may further improve the performance of all-optical quantum repeaters, as well as other applications of all-optical quantum processing. One of our

major contributions in this Chapter was to rigorously prove that the rate-loss scaling by the aforementioned genres of all-optical quantum repeaters with a fixed state size is given by $D\eta^s$ bits per mode, where D and s are constants that are functions of various device loss parameters, and that of design choices made (to choose the level of error protection). The fact that it is possible to achieve a value as the exponent $s < 1$ proves the fact that this scheme can outperform the key rates attainable by any QKD protocol that does not employ quantum repeaters, the rate performance of which are upper bounded by the PLOB bound, whose linear rate-transmittance decay implies $s = 1$.

In future work, it will be interesting to incorporate more realistic effects into the resource-performance tradeoff calculations of all-optical repeaters, in particular mode-mismatch errors in the passive interferometric manipulations on the photons held locally at the repeaters, and multi-photon errors arising from imperfect sources and noisy detectors.

This work shows that relatively small entangled states are capable of working as quantum repeaters but multiplexing based schemes for creating such states are inefficient. Better schemes for creating entangled photonic states would help make such schemes practical. We present one such scheme based on percolation theory in the next chapter.

Chapter 3

Percolation thresholds for linear optic quantum computing

3.1 Introduction

In *linear-optical quantum computing* (LOQC), a single photon in one of two orthogonal (spatial, temporal, or polarization) modes, i.e., $|10\rangle \equiv |0\rangle_L$ and $|01\rangle \equiv |1\rangle_L$ encodes a qubit, and passive linear optical interferometers and single-photon detectors implement gates and measurements. Since each qubit is encoded by one photon, we use *photon* and *qubit* synonymously. Gates and measurements in LOQC are inherently probabilistic even if all single-photon sources are ideal and all linear optical elements and detectors are lossless. Component losses further reduce success probabilities, leading to daunting requirements on number of devices (e.g., sources and detectors) to encode problems of practically-relevant size. Since the original Knill-Laflamme-Milburn (KLM) proposal for LOQC [28]—which was largely deemed unscalable due to the aforesaid reason—several variants have been proposed that use separately-prepared "ancilla" states and photon number resolving (PNR) detectors to boost the probabilities of nondeterministic operations.

A particularly promising variant is an LOQC architecture in the cluster-state model of quantum computing (QC) [1, 3], which was introduced by Kieling, Rudolph and Eisert [63, 64]. This scheme leverages percolation and renormalization, to (a)

probabilistically fuse many tiny microclusters (i.e., clusters of few entangled photons) using linear optical circuits into a randomly-grown large cluster, and (b) reinterpret the random instance of a large entangled cluster as a logical cluster state in the 2D square grid topology, which is a sufficient resource for universal QC [1]. Rudolph and colleagues subsequently showed constructions within the above framework, which they termed *ballistic* photonic QC, wherein they demonstrated that with 3-photon microclusters as an initial resource, one can create a percolated cluster with one-way transmission through a linear optical circuit, i.e., with no feedback [4, 65].

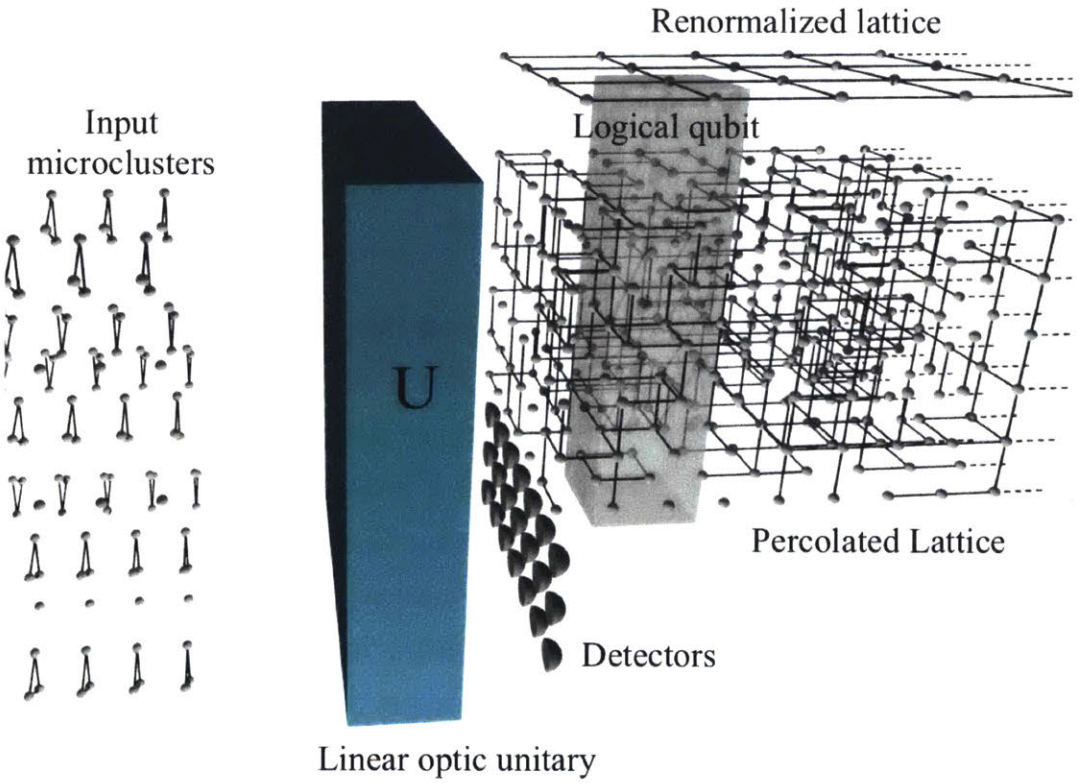

Figure 3-1: Ballistic photonic cluster state generation for quantum computing. A steady stream of entangled microclusters of size n-photons or less ($n = 3$ shown) is incident on a linear-optical interferometer (i.e., a multimode unitary transformation U), which produces an entangled cluster of photons at its output. If a percolation condition is met, the output can be renormalized into a fully-connected logical cluster in a topology universal for cluster-model quantum computing.

One can interpret the aforesaid feedback-free, or ballistic, framework of LOQC in the form visualized in Fig. 3-1, by 'pushing out' (postponing) the detections in-

volved in all the cluster-fusion operations to the very end. Consider an N-mode-input N-mode-output linear optical circuit—i.e., one that can be constructed using beam-splitters and phase-shifters [66]—and whose action on the input modes is described by a complex-valued unitary matrix U. At each time step, the linear-optical circuit is fed with several microclusters (of up to n entangled photons each) that occupy M of the N input modes. As we show, if a certain percolation threshold is exceeded, the spatio-temporal entangled sheet of photons that emerges at the output of U is a resource that is universal for cluster model quantum computing. This is true in the following sense. A fraction of the output modes is detected using PNR detectors at each time step. In the final "renormalization" step, the entangled state that emerges in the remaining output modes is broken up into logical blocks using information from the PNR detection outcomes. Exactly one representative photon is left unmeasured in each logical block while the rest of the photons are measured in appropriate bases to leave the representative photons in each logical block in a fully-connected 2D square grid cluster, into which one can encode any quantum algorithm. We emphasize that the detection outcomes are only used for the renormalization step, i.e., to figure out how to use the randomly connected output cluster for QC; they are *not* used to determine whether or not the unmeasured part of the output cluster is useful for universal QC (this is true with near certainty if the percolation condition is met).

A major open question which we address in this chapter is: if n-photon microclusters are the input resource, what is the minimum probability of success $\lambda_c^{(n)}$ with which each two-photon fusion attempt must succeed, such that one is guaranteed a percolated renormalizable cluster for universal quantum computing, assuming that the best possible spatio-temporal sequence of two-photon fusion attempts are employed on the input microclusters. Entangled microclusters can be used to increase the fusion success probability beyond 0.5, the highest value attainable with linear optics and photon detection alone [35, 36]. Therefore, some of the input microclusters can serve as building blocks for the percolated cluster while others can be used to boost the fusion operations. Therefore, a second important open question is: what is the maximum success probability $\lambda_{\max}^{(n)}$ attainable with n-photon microclusters used

as an ancillary resource? Clearly we need $\lambda_{\max}^{(n)} > \lambda_c^{(n)}$ for it to be possible to obtain a renormalizable percolated cluster. As n increases, $\lambda_{\max}^{(n)}$ and $\lambda_c^{(n)}$ are likely to increase and decrease respectively, driving the percolated cluster deeper into the supercritical-connected regime, which makes the construction more efficient by driving the dimensions of the renormalized blocks (and hence the number of photons that map to one logical lattice node in the renormalized lattice) to be smaller. Furthermore, if one allows for simultaneous fusion of three or more photons, very little is known about success probabilities of linear optical fusion and it is not clear if the thresholds and the efficiency of the above construction improves.

Another important practical question is the effect of losses and other device imperfections on the ballistic creation of resources for universal QC. If $\eta \in (0, 1)$ is the transmissivity each photon sees through its lifetime (including losses in the source, waveguides and detectors), as η decreases from one (the ideal lossless limit), $\lambda_{\max}^{(n)}(\eta)$ decreases while $\lambda_c^{(n)}(\eta)$ increases. There is a threshold on $\eta_c^{(n)}$ such that if $\eta < \eta_c^{(n)}$, $\lambda_{\max}^{(n)} > \lambda_c^{(n)}$ is no longer true. An open question therefore is whether this loss tolerance threshold improves with increasing size of input microclusters (i.e., $\eta_c^{(n)}$ decreases as n increases), and if so at what rate. Finally, in the presence of photon loss, since we don't know where losses occured, constructing a fully-connected universal logical renormalized cluster is non-trivial, and has not been addressed in the literature. To address this, one could modify the above scheme to start with the creation of logical photonic qubits that are tolerant to losses and other errors such as mode mismatch and detector noise, and thereafter do fusion, percolation and renormalization on these error-protected logical qubits.

When restricted to $n = 1$, i.e., only single photons as the initial resource, our setup in Fig. 3-1 resembles that of *Boson Sampling* (BS), a physics-based computation model introduced and analyzed by Aaronson and Arkhipov [67, 68]. If M photons are fed into a randomly chosen linear optical circuit U, and if *all* the output modes are detected using PNR detectors, the setup naturally samples from the induced N-mode M-photon joint probability mass function (pmf) at the output of U. It was shown that drawing samples from this particular joint pmf is very likely not possible

efficiently on a classical computer. However, it is also believed that BS does not have the computational power of universal quantum computing. The computational hardness of sampling from the output joint photon number distribution when $n \geq 2$ input clusters are employed, has not been analyzed. We emphasize however that the problem we described above (i.e., what conditions must be satisfied for the entangled state at the output of U to be a resource that is sufficient for universal QC) is distinct from the problem at the heart of BS (the computational hardness of sampling from the joint photon number distribution of the entangled state at the output of U). It will however be interesting to explore if there is a closer connection between the two problems, and whether a connectivity metric on the output entangled state can be mapped in a meaningful way to computational hardness of sampling from its joint photon number distribution.

3.2 Main results

Let us assume destructive two-photon fusion operations that succeed with probability λ. In other words, each fusion operation is assumed to act on two photons at a time, and regardless of whether the fusion succeeds or fails, those two photons are destroyed. With the optimal choice of sequence/pattern/algorithm to fuse the n-photon clusters, there exists an optimal (minimal) threshold $\lambda_c^{(n)}$, such that if all fusions succeed with probability $\lambda > \lambda_c^{(n)}$, the end product is a percolated cluster renormalizable for universal QC. These thresholds $\lambda_c^{(n)}$ for $n = 1, 2, \ldots$, and ways to achieve them, in particular for small values of n, are important questions that need to be answered in order to understand the resource-optimal way to realize photonic QC.

The results in this chapter can be summarized as follows:

1. **Converse**—We prove: $\lambda_c^{(n)} \geq 1/(n-1), \forall n \geq 2$, i.e., no matter how we choose to fuse n-photon clusters, if each two-photon fusion succeeds with probability less than $1/(n-1)$, the final cluster produced is fragmented with high probability, and not suitable for renormalization. This means that with $n = 3$ microclusters

(three-photon GHZ states) as the initial resource (as in [4, 65]), the minimum λ needed for percolation is 0.5. With $n = 2$ microclusters (Bell states) as the initial resource, if the fusions succeed with any probability less than one, the output cluster is not percolated. Hence, with pairwise destructive fusions, $n = 3$ microclusters are the minimum size needed for ballistic LOQC. However, our converse does not immediately tell us whether there exists a systematic prescription to achieve percolation *at* $\lambda_c^{(n)} = 1/(n-1)$. We also show that if $m \geq 2$ node fusion operations are employed to fuse n-qubit microclusters, the percolation threshold must satisfy: $\lambda_c^{(n,m)} \geq 1/[(n-1)(m-1)]$. However, very little is known about linear-optical circuits for $m > 2$ qubit fusion [69] (e.g., projecting 3 qubits to one of the 8 three-qubit GHZ states) and their associated success probabilities. Therefore, it remains unclear if the above bound on $\lambda_c^{(n,m)}$ is tight.

2. **New percolation framework**—We develop a new percolation framework to address the problem of assembling a large photonic cluster using cluster fragments, where the threshold on fusion success probability $\lambda_c^{(n)}$ maps on to the usual bond percolation threshold $p_c(G)$ of an appropriately defined logical graph G each of whose nodes corresponds to an n-photon microcluster. Each node in G is assigned a color based on how many of the n photons in the microcluster at that node are intended to be measured in fusion attempts, which is the node's degree, whereas each fusion attempt corresponds to a neighboring bond of a node in the logical graph.

3. **Improved achievable fusion thresholds**—Using our percolation framework, we present new constructions and associated fusion success thresholds for percolation. The lowest threshold we show achievable with $n = 3$ microclusters is ≈ 0.5898 which improves over a recently published threshold of 0.625 [4].

4. **Ballistic percolated cluster generation with a 2D graph**—We show a logical graph construction using a modified brickwork lattice, with which it is possible to fuse 3-photon microclusters in a 2D (planar) topology and achieve

percolation at $\lambda_c \approx 0.746$. This threshold being less than 0.78125 makes it possible to achieve using single-photon boosted linear optical fusion [36]. A planar architecture is very promising from an experimental standpoint because a planar integrated photonic waveguide can be used to weave such a cluster. This also shows it is possible to percolate a 2D lattice using single-photon-boosted fusion, a question left open by Rudolph [5].

5. **Conjectured achievable thresholds with two-photon fusion**—Finally, we conjecture, with compelling evidence, that if there is an infinite lattice G of maximum node degree n with bond percolation threshold p_c, it is possible to stitch together a giant percolated cluster renormalizable for QC using n-photon microclusters as long as the fusions succeed with probability $\lambda > p_c$. We show that the truth of this conjecture would imply that for $n = 3$, the lowest known achievable threshold would go down to 0.5, thereby proving $\lambda_c^{(3)} = 0.5$. We also conjecture, using an extension of the argument for $n = 3$, that the converse bound we prove is tight, i.e., it is possible to construct a logical graph that can be percolated with two fusion success probability, $\lambda_c^{(n)} = 1/(n-1)$.

6. **Loss tolerance of percolation thresholds**—Using a photon loss model inspired by a recent proposal to produce photonic microclusters using quantum dot emitters [70, 71], we prove an extension of our converse result, i.e., we show a lower bound on $\lambda_c^{(n)}$ that is a function of n and η (a parameter that quantifies the loss experienced by each photon). In other words, if the two-photon fusion success probability is less than this lower bound, for no sequence of fusing photons with a collection of n-photon microclusters, can one get a renormalizable percolated cluster. We discuss the implications of our results to the loss tolerance of photonic quantum computing using this scheme. We also discuss important open problems that need addressing, primarily that of renormalizing a cluster in the presence of photon loss and other device imperfections.

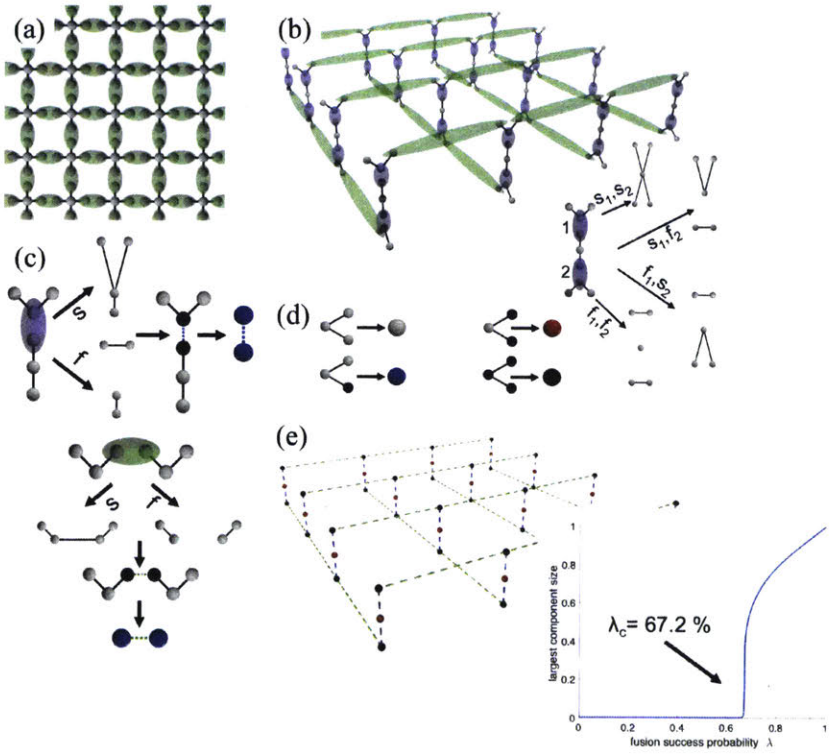

Figure 3-2: Different strategies and logical interpretations of piecing together a 2D square lattice by fusing microclusters: (a) 5-photon microclusters at each lattice node with fusion attempts on each lattice bond; (b) vertical arrangements of three 3-photon microclusters and 2 fusions create a 5-photon cluster if both fusions succeed; (c) interpreting fusion as coloring the measured nodes black and drawing a new bond between them if fusion succeeds, the linear optical circuits corresponding to the blue and green ellipses are shown in Fig. 2 and 3 of [4] respectively; (d) mapping microclusters to nodes in a logical graph and coloring them based on how many photons in the microcluster are left unmeasured; (e) pure bond percolation on the logical graph of colored nodes.

3.3 Revisiting ballistic cluster-state LOQC with a new approach

In this chapter, we develop a conceptually new way to construct percolated instances of renormalizable photonic clusters, and re-interpret recent results within our framework. We close the section with a conjecture. In Section 3.4 we use our percolation framework to develop new results on better achievable percolation thresholds, as well as general bounds on $\lambda_c^{(n)}$.

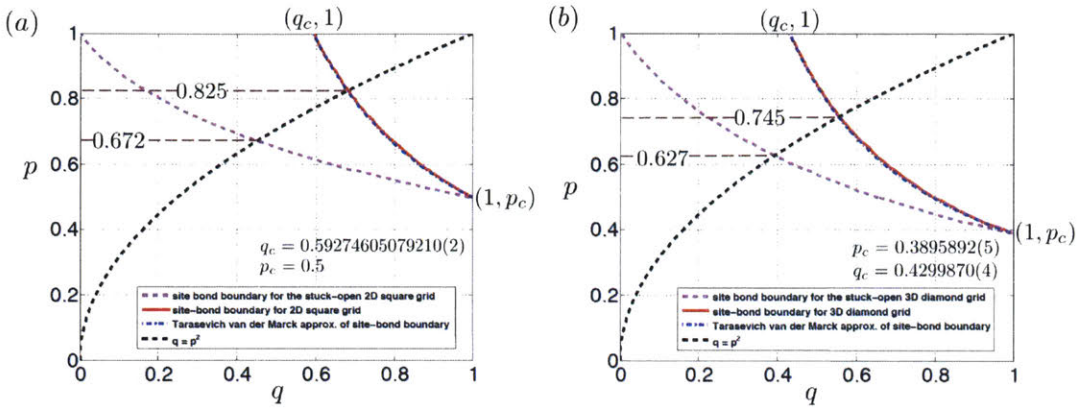

Figure 3-3: Site-bond percolation critical boundaries shown for the (a) 2D square and (b) 3D diamond lattices. The magenta curves correspond to a modified site-bond percolation problem described in the text where even if a site is not occupied, neighboring bonds can still be pairwise connected if occupied.

3.3.1 Graph states and linear optical fusion

We consider clusters of entangled photons that belong to a special class called *graph states* [72]. A cluster described by the graph $G(V, E)$ can be prepared by placing one qubit in the state $(|0\rangle_L + |1\rangle_L)/\sqrt{2}$ at every node in V and applying a two-qubit controlled phase (CZ) unitary operation across every edge in E. With single photons as the starting point, using passive linear optical circuits, a 2-qubit cluster can be generated with a success probability of 3/16 [53], and a 3-qubit cluster (in line or triangle topology) can be generated with a success probability of 1/32 [2], both assuming lossless linear optics and ideal detectors. The maximum success probability of linear-optical two-photon fusion, λ is 0.5 when no ancilla photons are used [73, 74]. Ancilla single photons can be used to achieve $\lambda = 0.78125$ [36].

3.3.2 Fusing microclusters on a regular lattice

We begin with an illustrative example of piecing together a large subgraph of the 2D square lattice by probabilistic fusion of microclusters using two-photon destructive fusion operations that succeed with probability λ.

A conservative approach: site-bond percolation

We begin by preparing 5-photon clusters in a star topology and placing them at each node of the lattice, as shown in Fig. 3-2(a) [63]. Suppose we succeed in preparing each of those clusters with probability q. We then attempt 2-photon fusions across each edge of the lattice, each of which succeeds with probability λ. The resulting graph state that is generated is a random instance of site-bond (mixed) percolation [75] where each bond is occupied with probability $p = \lambda$ and each site is occupied with probability q. The boundary in the (q, p) space that separates the percolating from the non percolating region is shown by the red solid plot in Fig. 3-3(a). We also show an analytical approximation of this critical boundary (blue dash-dotted plot), developed by Tarasevich and van der Marck [76]. If one had 3-photon clusters (GHZ states) as a starting resource, one can assemble a 5-photon star by attempting two fusions on three 3-photon clusters, as shown in Fig. 3-2(c). The probability of success in creating the 5-node star is thus $q = \lambda^2$, the probability that both fusions succeed. If either fusion fails, we call it a node failure. Therefore, per Kieling *et al.*'s recipe, the threshold value of λ beyond which one gets percolation is given by the intersection of the site-bond critical boundary and the line $q = p^2$, thereby obtaining $\lambda_c \approx 0.825$ (see Fig. 3-3(a)).

Exploiting failure modes: modified site-bond percolation with two stuck-open layers

It is too conservative to ask for both fusions to succeed at every node [4]. In other words, even if one or both fusions in creating the 5-node star were to fail, the leftover cluster fragments can still provide some connectivity. We illustrate this in Fig. 3-2(b), where we lay out the three 3-photon clusters at each node of the square lattice in the vertical arrangement shown, while the square lattice is divided into two crisscrossing layers of parallel 1D lattices. It is as if the lattice is stuck open at each node. If both fusions at a node—shown as light blue ovals—succeed (which happens with probability $q = \lambda^2$), the photon at the center of the vertical arrangement gets attached

to the two photons in the top layer as well as the two in the bottom layer, thus forming the 5-photon star. This has the effect of connecting the two layers at that node. If one or both fusions at a node fail (with probability $1 - q$), the node remains stuck open. But, even so, the two nodes in the top layer of the vertical arrangement remain connected to one another, and the same is true for the two nodes in the bottom layer. If one of the two fusions in the vertical arrangement succeeds (and the other fails), the two nodes in the layer closer to the successful fusion are connected via the center node, whereas the two nodes in the other layer (one closer to the failed fusion) are connected to one another directly. In all of these cases (i.e., if one or both fusions fail), the middle node plays no role in terms of providing long-range connectivity. The green ovals show fusion attempts between adjacent nodes in each of the two layers, the bonds of the square lattice.

The situation looks identical to (q, p) site-bond percolation with $q = \lambda^2$ and $p = \lambda$, except that even if a site is not active, the four neighboring bonds at that site can be pairwise connected to one another in the two stuck-open layers. We numerically evaluated the percolation region of this modified site-bond problem using the Newman-Ziff algorithm [77] on a square lattice of 1 million nodes. The resulting percolation boundary is shown in the magenta dashed plot in Fig. 3-3(a). This intersects with $q = p^2$ at $p = \lambda \approx 0.672$. This threshold is already below 0.78125, the success probability attainable by linear-optical fusion boosted with ancilla single photons [36].

Pure bond percolation on a logical graph

Let us revisit the picture in Fig. 3-2(b), and consider a new interpretation where each 3-photon cluster is thought of as a single (super) node in a logical graph shown in Fig. 3-2(e). We assign a color to the super node based on how many of its photons are intended to be measured (and hence destroyed) in the planned fusion attempts (Fig. 3-2(d)). The central photons in the 3-photon clusters at the centers of the vertical arrangements in Fig. 3-2(b) are not measured as part of a fusion. Hence, those 3-photon clusters map to a red node in the logical graph in Fig. 3-2(e). All other 3-photon clusters in Fig. 3-2(b) will have all their three photons measured in fusion

operations and so, all these 3-photon clusters are represented as black nodes in the logical graph. In the logical graph, a node represents an n-photon cluster, and a node's degree equals the number of its photons that will be measured in fusion attempts (and hence destroyed). A bond in the logical graph represents a fusion attempt, which is successfully activated with probability λ. With this new interpretation, the modified site-bond percolation discussed above can now be seen as simple single-parameter bond percolation on the logical graph, where each bond is independently activated with probability λ. It is simple to verify numerically (see the plot in Fig. 3-2(e)) that the bond percolation threshold equals $\lambda_c \approx 0.672$, as expected.

The black nodes disappear during the fusion attempts but help provide long-range connections. Only the red nodes, which in the example of Fig. 3-2(e) contain a single photon each after the fusion attempts, remain as part of the giant connected component, which is subsequently renormalized for quantum computing. Bond percolation guarantees that if N is the number of nodes in the logical graph G, and if $\lambda > \lambda_c$, the bond percolation threshold of G, then there is a unique giant connected component (GCC), i.e., a large cluster with $O(N)$ nodes. These $O(N)$ nodes have both red nodes and black nodes. However, it is simple to argue that there are $O(N)$ red nodes in the GCC.

Finally, note that in the example shown in Fig. 3-2(e), even though the logical graph—which describes how to lay out the microclusters prior to fusion attempts—is a non-planar two-layer graph, the physical giant cluster (of photons) obtained from percolation is a subgraph of the planar square lattice.

The diamond lattice and the (10,3)-b logical lattice

If one repeats the steps outlined in Sections 3.3.2, 3.3.2 and 3.3.2 for the 3D diamond lattice, i.e., lay out three 3-photon clusters in vertical arrangements as in Fig. 3-2(b) at each degree-4 node of the 3D diamond lattice—laid out in the layered 3D brickwork configuration as shown in [4]—and map it to a logical graph as described above, one obtains the logical lattice shown in Fig. 3-4. This is the (10,3)-b lattice [78] with one extra node inserted at the center of each of the vertical bonds. We refer to this

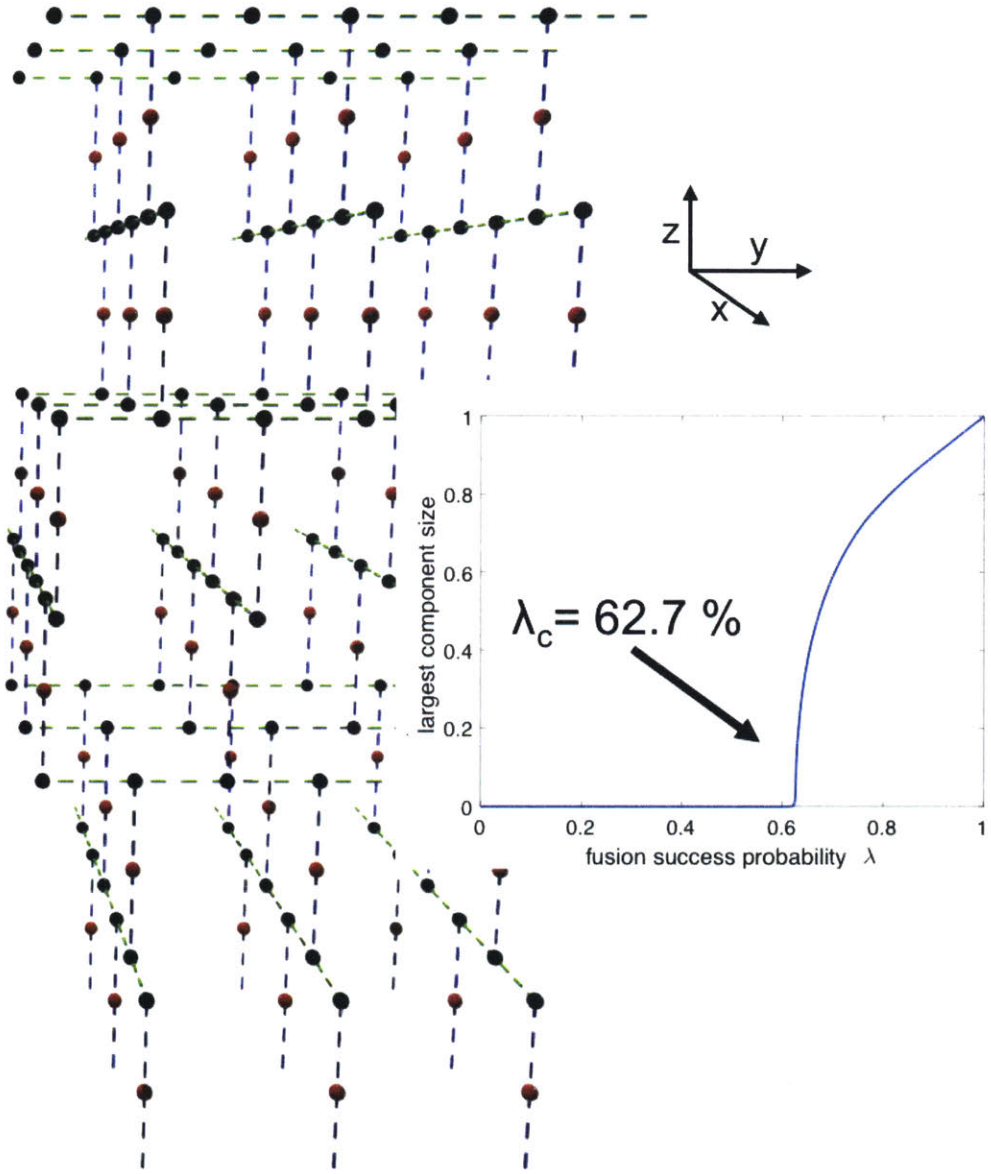

Figure 3-4: A 3D (10,3)-b lattice modified with additional nodes at the centers of each vertical bond. Pure bond percolation on this logical lattice corresponds to assembling the 3D diamond lattice using 3-photon microclusters discussed in [4]. Percolation threshold was evaluated by the Newman-Ziff method on a lattice with $\sim 10^6$ bonds.

as the 'modified" (10,3)-b lattice. The red nodes, as before, correspond to 3-photon microclusters with one unmeasured photon, whereas the black nodes correspond to 3-photon microclusters, all of whose photons will be measured in fusion attempts. We evaluated the bond percolation threshold of this modified (10,3)-b lattice using the Newman-Ziff algorithm, and obtained $\lambda_c \approx 0.627$, which agrees with, and sharpens the result of [4] (i.e., $\lambda_c \approx 0.625$); but is now interpreted as a standard bond percolation threshold.

3.3.3 General picture for ballistic LOQC

The discussion in Section 3.3.2 logically leads to a new approach to constructing a large percolated network of photons for ballistic LOQC. The basic optimization problem is to pick a *logical* graph G of maximum node degree n such that it has the lowest possible bond percolation threshold, $p_c^{(n)}(G)$. But there are additional conditions that G should satisfy in order for the GCC in the supercritical (percolated) regime to be useful for universal quantum computing.

In order to understand these additional conditions, consider that each node of the N-node graph G represents an n-photon microcluster—a connected graph of n nodes—and a bond between nodes u and v of G represents a fusion attempt between a photon in the microcluster at u and a photon in the microcluster at v. A node in G is assigned color zero only if the degree of the node is n. A bond-percolation instance on G, where every bond is activated independently at random with probability λ, results in $O(N)$ nodes of G in a unique connected component if $\lambda > p_c^{(n)}(G)$. This connected component in turn induces a (random) cluster state of the individual unmeasured photons in the original microclusters as nodes of a graph. In order for this percolated cluster to be useful for QC,

(a) The percolated cluster must have $O(N)$ photons. This is equivalent to the condition that the GCC of the logical graph resulting from bond percolation on G has $O(N)$ non-zero-color nodes.

(b) The percolated cluster of photons can be renormalized into a graph that is known

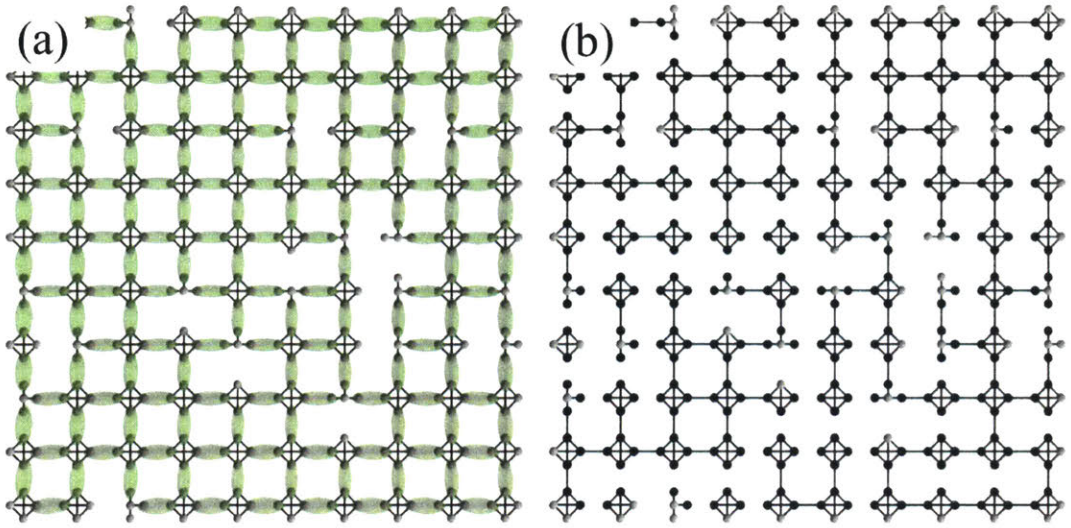

Figure 3-5: (a) 4-node microclusters laid out on nodes of a square lattice. A random $\alpha = 0.3$ fraction of microclusters are put in star configuration the central photon of which will not be measured in any fusion operation. All other photons are measured in fusion attempts. (b) A random instance after the fusion attempts, assuming that each fusion succeeds with probability $\lambda = 0.6$. The measured photons are colored black. The unmeasured photons (colored white) in the giant component of the percolated lattice form the backbone random graph that is renormalized into a fully connected 2D topology for universal cluster-state quantum computing.

to be a universal resource for quantum computing [79] (e.g., a square lattice).

Let us assume G is a regular lattice with uniform node degree d. Let us also assume that we have access to d-node microclusters. One strategy for selecting nodes in G designated to have non-zero-color is to pick a random fraction, α, of the N nodes in G as color-1 and populate them with d-photon star clusters. Clearly, these nodes will have one less degree $(d-1)$. We then populate d-photon clique clusters at the remaining $(1-\alpha)N$ degree-d nodes. These nodes have color 0 and hence all the photons in the cliques will be measured in the fusions. If α is small, then the fusion success probability exceeding the bond percolation threshold of G, i.e., $\lambda > p_c(G)$, should suffice to guarantee percolation. This would mean that $\lambda_c^{(n)} \leq \min_{G(V,E):deg(V)=n} p_c(G)$. In order to prove this formally, one needs to argue that conditions (a) and (b) in the previous paragraph are met. We leave this for future work. If this conjecture is correct, given that the bond percolation threshold of the degree-3 3D (10,3)-b lattice

is 0.546694 [78], it would mean that $\lambda_c^{(3)} \leq 0.546694$ for a 3D lattice. Furthermore, it is possible to generalize the (10,3)-b lattice to higher dimensions, following a procedure similar to the generalization of the "modified" (10,3)-b lattice in section 3.4.1. Under this construction, $\lambda_c \to 0.5$ as the number of dimensions $\to \infty$. Combined with our converse of $\lambda_c^{(3)} \geq 0.5$, this would imply that $\lambda_c^{(3)} = 0.5$. We conjecture that a higher dimensional construction with size $n \geq 3$ microclusters can saturate the converse bound which would imply that $\lambda_c^{(n)} = 1/(n-1)$, $\forall n \geq 3$. A schematic of the setup described in the discussion above, with G chosen as the 2D square lattice for illustration, is shown in Fig. 3-5.

3.4 Fundamental thresholds

We begin this section with new results on achievable fusion success thresholds using 3-photon microclusters in Section 3.4.1, i.e., tighter upper bounds on $\lambda_c^{(3)}$ compared to known results. In Section 3.4.2 we provide an intuitive proof of our general converse bound $\lambda_c^{(n)} \geq 1/(n-1)$, $\forall n \geq 2$. Finally, in Section 3.4.4, we discuss how losses in devices and inline losses affect the fusion thresholds, and discuss its implications for the resource overhead (number of sources and detectors) for ballistic LOQC in the presence of losses.

3.4.1 Achievable thresholds

Throughout this section, we take the size of our initial microcluster to be $n = 3$ photons. As in Fig. 3-2 (e), blue and green dashed lines correspond to the fusion operations represented by the blue and green ellipses in Fig. 3-2(c), respectively. The degree-3 nodes are color-0 (black) and hence have 3-photon clusters all whose photons will be measured in fusion attempts. The degree-2 nodes are color-1 (red) and have 3-photon clusters of which only two photons will be measured in fusion attempts.

Let us pick as the logical graph the modified 2D brickwork lattice shown in Fig. 3-6. The bond percolation threshold of this lattice is $\lambda_c \approx 0.746$, as shown in the inset of Fig. 3-6. It is simple to argue that conditions (a) and (b) discussed in

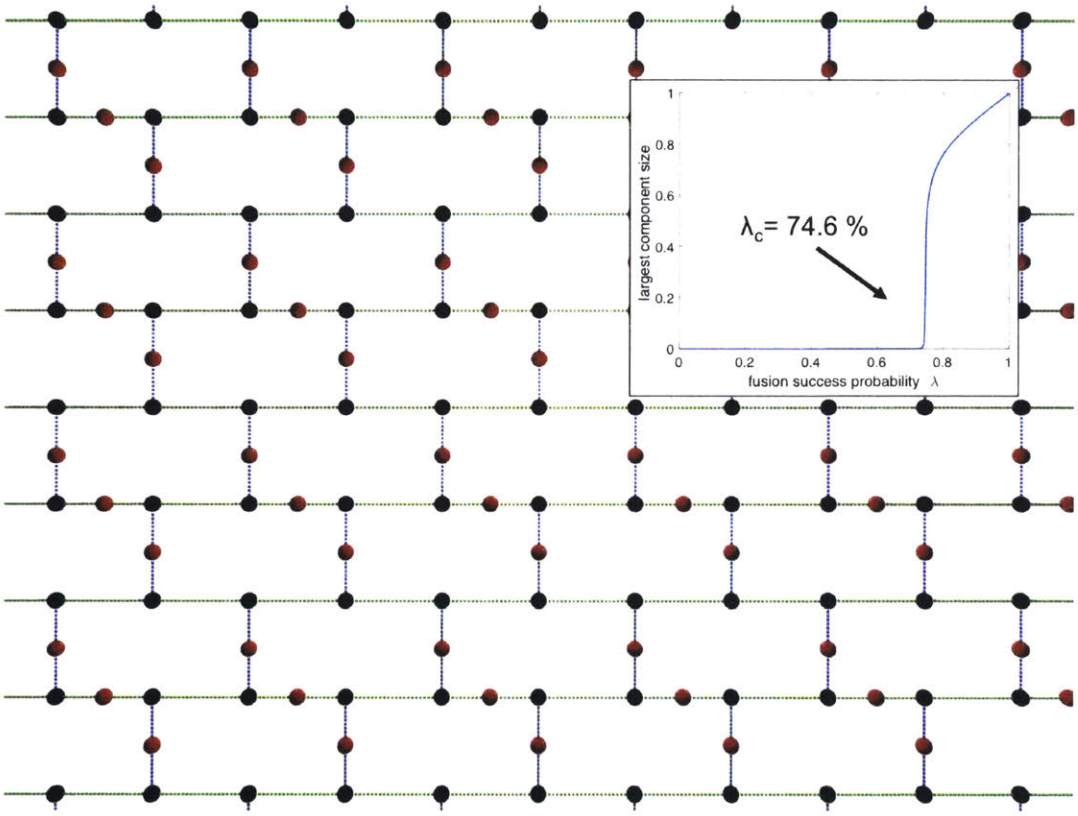

Figure 3-6: A modified 2D brickwork lattice used as logical graph with node colors as shown yields $\lambda_c \approx 0.746$, which settles an open question in [5] on whether it is possible to attain ballistic LOQC with 3-photon microclusters with a fully 2D architecture and $\lambda_c < 0.78125$, which is achievable with unentangled ancilla photons. Percolation threshold was evaluated by the Newman-Ziff method on a lattice with $\sim 10^6$ bonds.

Section 3.3.3 are met, and the resulting percolated cluster is renormalizable. Hence, we have shown that even with a 2D lattice, starting with three photon microclusters, it is possible to assemble a resource for universal QC, since $\lambda_c \approx 0.746 < 0.78125$ and two photon fusion with success probability 0.78125 is achievable with a linear optical circuit boosted with ancilla single photons [36]. Being able to percolate with a 2D lattice makes ballistic LOQC much easier from the experimental standpoint since a planar integrated photonic waveguide can be used to weave such a cluster. The existence of a 2D lattice with this property was posed as an open question by Rudolph [5].

In Section 3.3.2, we described a logical graph construction of the "modified" (10,3)-b lattice (Fig. 3-4), using which we reinterpreted the results of [4] as a pure bond percolation threshold, $\lambda_c \approx 0.627$. We now consider a 4D extension of the modified (10,3)-b lattice (Fig. 3-7) as the logical graph. The 3D lattice (Fig. 3-4) comprises (x, y)-plane layers of parallel 1D line lattices of black (degree-3) nodes stacked along the z direction. The layers alternate between their line lattices pointing in the x and y directions, while neighboring layers are straddled by a layer of red (degree-2) nodes. Each black node has two black-node neighbors on either side of the 1D lattice to which it belongs, connected via green bonds, and one red-node neighbor, connected via a blue bond. Along each line lattice of black bonds, the blue bonds alternate between the $+z$ and $-z$ directions. The adjective "modified" in the name of this lattice refers to the fact that in the standard (10,3)-b lattice, the red nodes are not there, and adjacent (x, y) planes of parallel lattices in alternating directions are directly connected via bonds. Our 4D generalization of the modified (10,3)-b lattice is shown in Fig. 3-7. It consists of a doubly infinite stacking of (x, y)-plane layers—of parallel 1D line lattices of black (degree-3) nodes—stacked along the z and w directions respectively. Of the three neighboring bonds of a black node, two (green) bonds—connecting to neighboring black nodes in the line lattice to which it belongs—are in the (x, y) plane, whereas one (blue) bond—connecting to a red node which in turns connects via another blue bond to a black node in a neighboring (x, y)-plane layer—points in either the z direction or in the w direction. Along each line lattice of

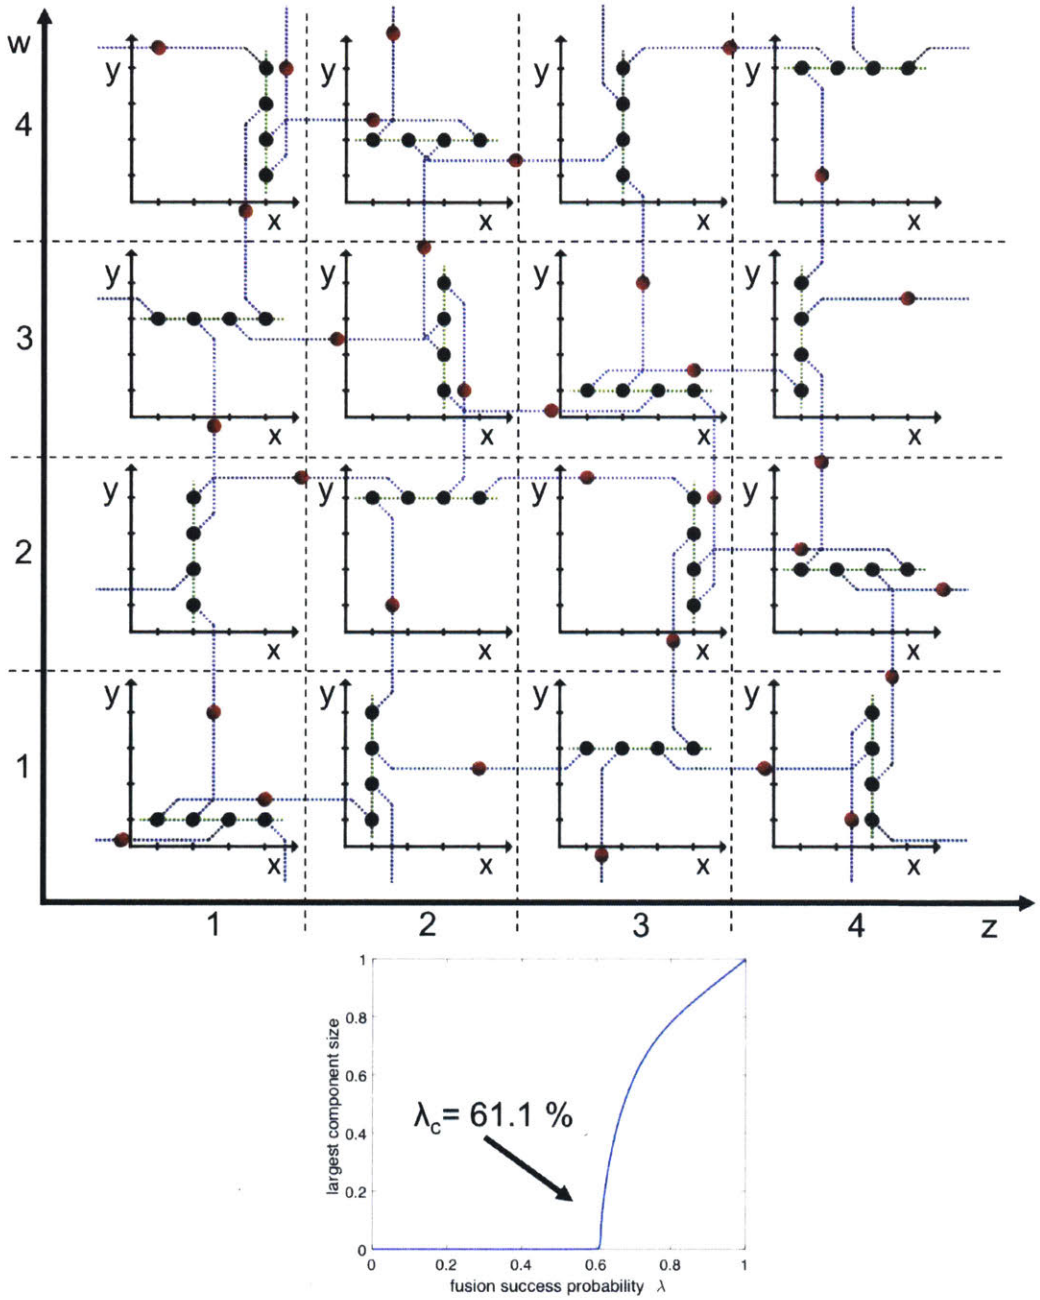

Figure 3-7: Schematic of the 4D extension of the (10,3)-b lattice, which when used as the logical graph with node colors as shown yields $\lambda_c \approx 0.611$. Percolation threshold was evaluated by the Newman-Ziff method on a lattice with $\sim 10^7$ bonds. The inner plots with x and y axes represent projections of the lattice on the (x, y) plane at the z and w values shown on the outer axes.

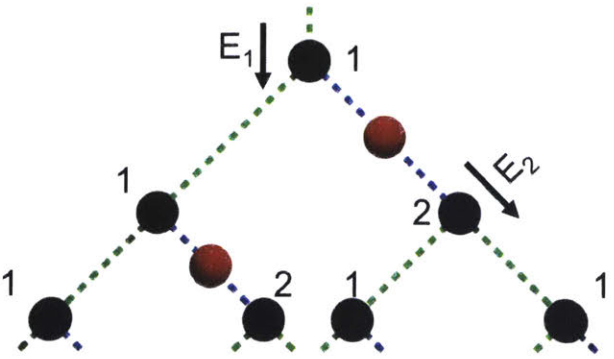

Figure 3-8: Schematic of the ∞-D extension of the (10,3)-b lattice, which when used as logical graph with node colors as shown yields $\lambda_c \approx 0.5898$. Percolation threshold was evaluated analytically.

black bonds, the blue bonds alternate between directions $+z$, $+w$, $-z$, $-w$, ..., and so on. The graph has a period of four in each of the x, y, z and w dimensions. One period of the lattice is depicted in Fig. 3-7. The inner axes represent an (x, y) plane at a given value of z and w. This construction results in longer loops compared to the 3D case while retaining the 3D graph's coordination number (average node degree), which in turn lowers the bond percolation threshold for the 4D logical graph. We find, using a Newman-Ziff simulation performed on a 4D modified (10,3)-b lattice of size $N \sim 10^7$ nodes, that $\lambda_c \approx 0.611$.

By adding more dimensions to the aforesaid logical lattice construction, the size of the loops is increased, hence progressively lowering λ_c. Finally, in the case of the ∞-dimensional modified $(10, 3)$-b lattice, the loops are infinitely far apart and hence the lattice is locally tree like. The local connectivity of this logical graph is depicted in Fig. 3-8. A simple analytical argument, explained below, shows that $\lambda_c \approx 0.5898$ for this limiting construction. This threshold, along with the converse proven in the next section, establishes that $0.5 \leq \lambda_c^{(3)} \leq 0.5898$, thereby improving upon ≈ 0.625, the lowest-known fusion probability threshold that is known to be achievable with 3-photon microclusters [4]. This also is the minimum $\lambda_c^{(3)}$ attainable from higher-dimensional logical lattices of the modified (10,3)-b lattice family. For the entire family of constructions, we argue that conditions (a) and (b) discussed in

Section 3.3.3 are met, and the resulting percolated cluster can be renormalized for QC.

The locally-tree-like structure of the ∞-dimensional modified $(10, 3)$-b lattice is shown in Fig. 3-8. Similar to the 3D and 4D modified (10,3)-b lattices, each black node has two green bonds and one blue bond (which leads to a black node via a red node and another blue bond). We denote the expected number of children of a node when approached via a green bond as E_1 and the expected number of children of a node when approached via a blue bond as E_2. When counting the number of children of a node, we only count red nodes since they are the only nodes with unmeasured qubits. Counting children from the top of the Fig 3-8, each black node is labelled as 1 or 2 depending on the bond from which it is approached. Counting children at the points labelled E_1 and E_2 yields the equations $E_1 = \lambda E_1 + \lambda + \lambda^2 E_2$ and $E_2 = 2\lambda E_1$ where λ is the bond probability. For percolation, $E_1 \to \infty$ and solving the equations with this condition, we find that $\lambda_c + 2\lambda_c^3 = 1$, which leads to $\lambda_c = 0.5898$.

A tree is known not to be a universal resource for QC [79]. However, entangled trees clusters can be used for other applications, e.g., as loss tolerant logical qubits [57], with applications to all-photonic quantum repeaters [51, 80]. We now show that with a degree-n Bethe Lattice (an infinite tree) as the logical graph, and with n-photon microclusters as the initial resource, we can get $\lambda_c^{(n)} = 1/(n-1)$, which saturates the converse bound we prove in the following section. Whether or not $\lambda_c^{(n)} = 1/(n-1)$ can be attained on a lattice whose percolated instance can be renormalized into a logical cluster universal for QC, remains open.

The logical graph that can be used to approach the $1/(n-1)$ limit is shown in Fig. 3-9 for $n = 3$. We start the depiction of our tree at a degree $n-1$ unmeasured node (i.e., a node with an unmeasured qubit), after which there are g generations of degree n black nodes, followed by a generation of unmeasured nodes, followed by g generations of black nodes and so on. In the tree depicted in Fig. 3-9, $g = 2$. Starting from an unmeasured node, given a bond probability of λ, the expected number of unmeasured nodes after $g+1$ generations is $\lambda(n-2)[\lambda(n-1)]^g$. Hence the critical bond percolation probability must satisfy $\lambda_c(n-2)[\lambda_c(n-1)]^g = 1$, which gives

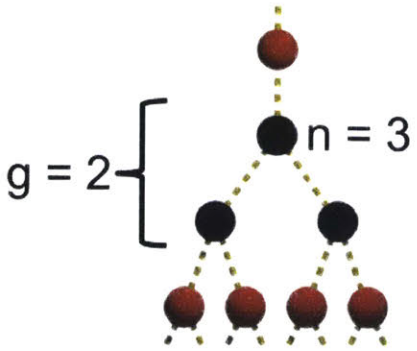

Figure 3-9: Schematic of the lattice construction used to approach the $\lambda_c = 1/(n-1)$ limit for the case of $n = 3$ and $g = 2$.

us $\lambda_c = (n-2)^{-1/(g+1)}(n-1)^{-g/(g+1)}$. As g increases, we approach the limit of $\lambda_c = 1/(n-1)$. In the argument above, we only count the number of unmeasured nodes and condition (a) of Section 3.3.3 is satisfied.

In the construction of the Bethe lattice above, the input states are n photon cliques, which are equivalent, up to local operations, to n photon GHZ states. The fusion operation used here (yellow dashed lines), acting on two qubits A and B, consists of a Hadamard gate on qubit A followed by Bell measurement of A and B in the $\{1/\sqrt{2}(|00\rangle \pm |11\rangle), 1/\sqrt{2}(|01\rangle \pm |10\rangle)\}$ basis (also described in [65]). Since the order of the Bell measurements is not important, we imagine first applying the fusion operations corresponding to successes. A successful fusion between two cliques removes qubits A and B from the graph and places the rest of the photons in a clique. Hence any two logical nodes that have an edge in Fig. 3-9 are part of the same clique and hence connected. A failed fusion results in an X measurement on A and a Z measurement on B. The Z measurement of a qubit simply removes the photon and all its edges. The X measurement of a qubit in a clique has the effect of a Z measurement followed by a Hadamard gate on one of the original neighbors of the qubit. Since a Hadamard gate followed by a Z (resp., X) measurement has the effect of an X (resp., Z) measurement, the result of the failed fusions is simply the removal of the corresponding nodes from the cliques without disturbing the connectivity between any other nodes. Hence the fusion operation described here can be used to create the

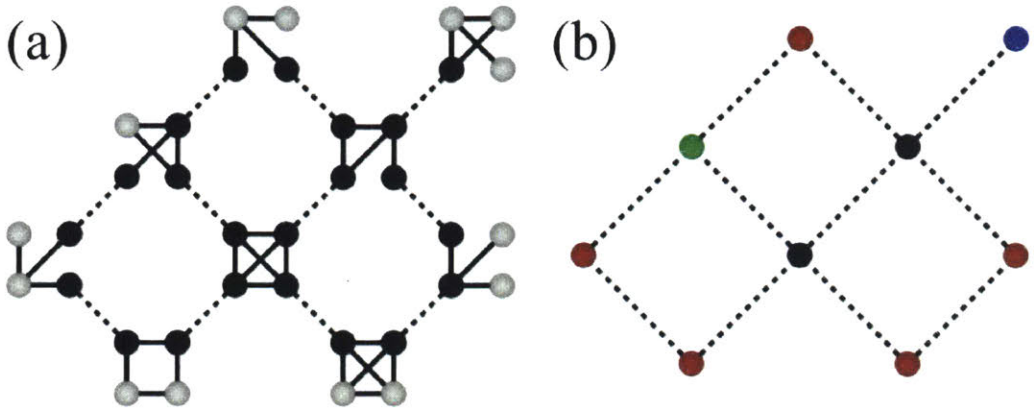

Figure 3-10: (a) An example of a series of two-node fusions on $n = 4$ sized microclusters. (b) Mapping of the microclusters to nodes in a logical graph. Logical nodes with one, two, three, and four measured physical nodes are colored as Blue, Red, Green, and Black, respectively.

logical graph in Fig. 3-9.

Finally, as discussed in Section 3.3.3, we conjecture that if there is an infinite lattice G of maximum node degree n with bond percolation threshold p_c, it is possible to assemble a giant percolated cluster renormalizable for universal QC using n-photon microclusters as long as the fusions succeed with probability $\lambda > p_c$. We also conjecture, using an extension of the argument for $n = 3$ using an infinite-dimensional modified (10,3)-b lattice, that the converse bound $\lambda_c^{(n)} \geq 1/(n-1)$ is tight for all n, i.e., it is possible to construct a logical graph that can be percolated with two-fusion success probability $= 1/(n-1) + \epsilon$, for any $\epsilon > 0$.

3.4.2 Intuitive converse

In this section we discuss the intuition behind our converse result: starting with N microclusters each of n photons and using any sequence of two-photon destructive fusion operations, the minimum fusion success probability λ_c sufficient to obtain a connected component of $O(N)$ unmeasured photons with high probability is $\geq 1/(n-1)$. A formal proof is provided in the next section. We sketch the intuition behind the proof below.

Fig. 3-10(a) illustrates an example with $n = 4$ photon microclusters and a set of two-photon fusion attempts shown as dashed lines each of which succeeds with probability λ, using the graphical interpretation of fusion presented in Fig. 3-2(c). Recalling our convention from Section 3.3, a black photon is one that gets measured in a fusion attempt, hence does not exist after the fusion attempt involving it has happened, regardless of its success. After all the fusions have been attempted, one obtains connected components involving only the white photons.

Given a large number N of n-photon microclusters, our objective is to pick a set of photon pairs on which to attempt fusions (each of which succeeds with probability λ), such that $\lambda > \lambda_c$ ensures a unique connected component of $O(N)$ white photons with the smallest possible λ_c. One can argue that the post-fusion connectivity graph that results between the surviving (white) photons, no matter what kind of destructive linear-optical fusion operation is used, can be no more connected than the connectivity between white photons in the graph shown in Fig. 3-10(a). In other words, if two white photons have a path connecting them (via black and white photons) in a random instance of the graph in Fig. 3-10(a), those two photons would also have a connected path in the actual post-fusion connectivity graph assuming the same success-failure fusion instances, if any linear optical circuit for fusion is employed.

Fig. 3-10(b) shows the mapping of Fig. 3-10(a) to a logical graph where each microcluster is replaced by a logical node, similar to Fig. 3-2(d). Here, dashed lines represent bonds in the logical graph that exist with probability λ. Logical nodes corresponding to microclusters with one, two, three, and four measured (black) photons are colored Blue, Red, Green and Black, respectively. Since the microclusters in Fig. 3-10(a) have $n = 4$ photons and each photon is associated with at most one fusion attempt, the maximum degree of each logical node in Fig. 3-10(b) is $n = 4$. Hence, the post-fusion instance of the logical graph in Fig. 3-10(b) represents an instance of bond percolation on some graph of degree four. In general, starting with n photon microclusters, and any sequence of fusion attempts, the resulting instance of the logical graph is a bond percolation instance (with bond success probability of λ) on some graph of maximum degree n. Of all infinite graphs of maximum degree n,

the minimum bond percolation threshold is that of the degree-n Bethe Lattice, and equals $1/(n-1)$ [81, 82]. Finally, since each of the N logical nodes in Fig. 3-10(b) maps to n photons in Fig. 3-10(a), with n finite, the absence of a connected component with $O(N)$ logical nodes in the post-fusion instance of the logical lattice of Fig. 3-10(b) implies the absence of a connected component of $O(N)$ white photons in the post-fusion instance of the physical photonic cluster shown in Fig. 3-10(a). This completes our proof that $\lambda_c^{(n)} \geq 1/(n-1)$. Since the Bethe Lattice is a tree, it cannot be renormalized into a logical cluster state that is needed for universal quantum computing [79]. Therefore, the above proof does not establish the tightness of the aforesaid bound. As explained however in the context of $n = 3$ in Section 3.3, we believe that the bound $\lambda_c^{(n)} \geq 1/(n-1)$ is tight, and is achievable by going to progressively higher dimensional equivalents of the (10,3)-b lattice for the logical graph, since the logical graph's local topology increasingly resembles that of the Bethe Lattice while retaining the renormalizability of its bond-percolated instance for any finite dimension. But a rigorous proof of the above and a fully-specified construction of the achievability of the $1/(n-1)$ threshold for universal QC is beyond the scope of this thesis.

Using similar reasoning, it is also possible to show that starting with N micro-clusters of size n and using any sequence of m-node (destructive) fusion operations, the fusion success probability threshold required to obtain a component of $O(N)$ unmeasured photons satisfies $\lambda_c^{(n,m)} \geq 1/[(n-1)(m-1)]$. Very little is known about linear-optical fusion of more than two photons at once and their associated success probabilities. Therefore, it is unclear whether the above lower bound on $\lambda_c^{(n,m)}$ is tight.

3.4.3 Formal proof of Converse

A graph $G = (V, E)$ is a collection of vertices V and edges E with each vertex having a color property that is a whole number. Two vertices in a graph are said to be connected if there are a path connecting them. $LCC(G)$ is the number of vertices in the largest connected component in G.

We define a function $G' = F_n(G, L, S)$ with $n \geq 1$ that has the following inputs: G is the input graph, $L = \{u, w\}$ is a set of two vertices in G and $S \in \{0, 1\}$. L is a valid input only if both u and w have nonzero color. G' is a graph that is identical to G except for the following differences: 1) The colors of u and w are incremented by one modulo $n+1$, 2) If $S = 1$ and $u \neq w$, then an edge is introduced between u and w. We use $F_n(G, \vec{L}, \vec{S})$ with input lists $\vec{L} = \{L_1, L_2, \ldots, L_k\}$ and $\vec{S} = \{S_1, S_2, \ldots, S_k\}$ as shorthand for $F_n(\ldots F_n(F_n(G, L_1, S_1), L_2, S_2), \ldots), L_k, S_k)$. Based on this definition of $F_n(G, \vec{L}, \vec{S})$, if all vertices in G have color 1, $\vec{L}$ is a valid list if and only if every vertex appears no more than n times in $\vec{L}$

We now define a class of functions that we call "destructive two node fusion". A "destructive two node fusion" is a function of the form $G' = D(G, L, S)$ where G is the input graph, $L = \{u, w\}$ is a set of two vertices in G and $S \in \{0, 1\}$. G' is the output graph which has the same set of vertices as G but different vertex colorings and a different set of edges. L is a valid input only if both u and w have nonzero color in G. U and W are the set of vertices in G connected to u and w respectively. The function $G' = D(G, L, S)$ is a "destructive two node fusion" if and only if it follows the following rules: *Rule 1* In G', u and w have color 0 and no edges attached to them. *Rule 2* If vertex $a \notin U \cup W$ was not connected to a vertex b in G, a and b should remain unconnected in G'. *Rule 3* For $S = 0$ and $U \neq W$, every vertex is U is disconnected from every vertex in W. As in the case for F_n, $D(G, \vec{L}, \vec{S})$ with $\vec{L} = \{L_1, L_2, \ldots, L_k\}$ and $\vec{S} = \{S_1, S_2, \ldots, S_k\}$ is shorthand for $D(\ldots D(D(G, L_1, S_1), L_2, S_2), \ldots), L_k, S_k)$. It should be noted that a set of valid inputs G, $\vec{L}$, $\vec{S}$ to $D(G, \vec{L}, \vec{S})$ is always a valid input to $F_1(G, \vec{L}, \vec{S})$ as well. Also, note that F_n is not in the class of "destructive two node fusion".

Lemma 1 $LCC(F_1(G, \vec{L}, \vec{S})) \geq LCC(D(G, \vec{L}, \vec{S}))$ for any "destructive two node fusion" D and any set of valid inputs to D.

Proof: Consider 2 graphs G_i and G_j with the same vertices and vertex colorings. Connectivity between any two vertices in G_i implies connectivity between the same vertices in G_j. u and w are vertices in G_i, G_j and U and W are the set of vertices in G_i connected to u and w respectively. D is a "destructive two node fusion". If u and

w are connected in $G_D = D(G_i, \{a,b\}, S)$, either u and w are connected in G_i or $(S = 1)\&((a \in U \& b \in W) \vee (a \in W \& b \in U))$. If u and w are connected in G_i, they must be connected in $G_F = F_1(G_j, \{a,b\}, S)$ since connectivity in G_i implies connectivity in G_j and F_1 does not remove edges. If $(S = 1)\&((a \in U \& b \in W) \vee (a \in W \& b \in U))$, there is an edge between a vertex of U and a vertex of W in G_F which implies that u and w are connected. Therefore, connectivity of u and w in G_D implies connectivity of u and w in G_F.

Consider a graph G with two vertices u, w. Using the result of the previous paragraph with $G_i = G_j = G$, connectivity of u, w in $G_{D1} = D(G, L_0, S_0)$ implies connectivity of u, w in $G_{F1} = F_1(G, L_0, S_0)$. Using the result from the previous paragraph again with $G_i = G_{D1}$ and $G_j = G_{F1}$, connectivity of u, w in $G_{D2} = D(G_{D1}, L_1, S_1)$ implies connectivity of u, w in $G_{F2} = F_1(G_{F1}, L_1, S_1)$. By repeating this process with lists $\vec{L} = (L_1, L_2, \ldots, L_k)$ and $\vec{S} = (S_1, S_2, \ldots, S_k)$, if two vertices are connected in $D(G, \vec{L}, \vec{S})$, they are connected in $F_1(G, \vec{L}, \vec{S})$. This leads to Lemma 1.

G_0 is a graph composed of nN vertices which are labelled as v_{lm} with $l \in 1, 2, \ldots, N$ and $m \in \{1, 2, \ldots, n\}$. Two vertices v_{lm} and $v_{l'm'}$ are unconnected if $l \neq l'$. We define C_k as the set of all vertices of the form v_{km} with $m \in \{1, 2, \ldots n\}$. Hence, G_0 is composed of N isolated clusters labelled as $C_1, C_2, \ldots, C_N$. All vertices of G_0 have color 1. Note that two vertices in the same cluster may or may not be connected but vertices in different clusters are unconnected.

G_1 is a graph that consists of N vertices labelled as v_l with $l \in \{1, 2, \ldots, N\}$. Every vertex v_k of G_1 maps to cluster C_k in G_0. There are no edges in G_1. All vertices of G_1 have color 1.

$\vec{L}_0$ is a list that is valid as input to $F_1(G_0, \vec{L}_0, \vec{S})$. This requires that every vertex appears no more than once in $\vec{L}_0$. $\vec{L}_1(\vec{L}_0)$ maps the list $\vec{L}_0$, which contains vertices from G_0 to a list that contains vertices in G_1 by removing the second subscript in every vertex e.g. if $\vec{L}_0 = (\{v_{11}, v_{22}\}, \{v_{12}, v_{31}\}, \ldots)$, $\vec{L}_1(\vec{L}_0) = (\{v_1, v_2\}, \{v_1, v_3\}, \ldots)$.

Lemma 2: $nLCC(F_n(G_1, \vec{L}_1(\vec{L}_0), \vec{S})) \geq LCC(F_1(G_0, \vec{L}_0, \vec{S}))$ for any valid list $\vec{L}_0$.

Proof: For $\vec{L}_0$ to be a valid input to $F_1(G_0, \vec{L}_0, \vec{S})$, every vertex appears at most

once in it. n vertices from G_0 in $\vec{L}_0$ map to 1 vertex from G_1 in $\vec{L}_1(\vec{L}_0)$. Hence, every vertex appears at most n times in $\vec{L}_1(\vec{L}_0)$ and therefore $\vec{L}_1(\vec{L}_0)$ is a valid input to $F_n(G_1, \vec{L}_1(\vec{L}_0), \vec{S})$.

If two vertices from different clusters C_a and C_b in $F_1(G_0, \vec{L}_0, \vec{S})$ are connected, there must exist at least one element $L_{0k} = \{v_{al}, v_{bl}\}$ in $\vec{L}_0$ with the corresponding entry in $\vec{S}$, $S_k = 1$. $l, m \in 1, 2, \ldots, n$. Hence the corresponding entry in $\vec{L}_1(\vec{L}_0)$ must have the form $L_{1k} = \{v_a, v_b\}$ with $S_k = 1$. Therefore, if two vertices in two different clusters C_a and C_b are connected in $F_1(G_0, \vec{L}_0, \vec{S})$, the corresponding vertices v_a and v_b in $F_n(G_1, \vec{L}_1(\vec{L}_0), \vec{S})$ are connected.

Suppose that the largest connected component of $F_1(G_0, \vec{L}_0, \vec{S})$ contains vertices from k different clusters. Since the size of every cluster is n, this implies that $LCC(F_1(G_0, \vec{L}_0, \vec{S})) \leq nk$ with the equality holding if any two vertices within the same cluster are connected. From the previous paragraph, if any vertex in $F_1(G_0, \vec{L}_0, \vec{S})$ is a part of the largest connected component, the corresponding vertex in $F_n(G_1, \vec{L}_1(\vec{L}_0), \vec{S})$ will be part of the largest connected component. Hence, $LCC(F_n(G_1, \vec{L}_1(\vec{L}_0)) = k$. Therefore, $nLCC(F_n(G_1, \vec{L}_1(\vec{L}_0))) = nk \geq LCC(F_1(G_0, \vec{L}_0, \vec{S}))$ and we obtain Lemma 2.

Assume that $\vec{L}_2$, a list of node pairs from G_1 is a valid input to $F_n(G_1, \vec{L}_2, \vec{S})$.

Lemma 3 For $\lambda < 1/(n-1)$, $LCC(F_n(G_1, \vec{L}_2, \vec{S}(\lambda))) = o(N)$ for any valid list $\vec{L}_2$ where $\vec{S}(\lambda)$ is a probabilistic function that returns a string of bits of the same size as the number of elements in $\vec{L}_2$. Each entry of $\vec{S}(\lambda)$ is independently chosen as 1 with probability λ and 0 with probability $1 - \lambda$.

Proof: The graph $F_n(G_1, \vec{L}_2, \vec{S}(\lambda))$ contains an edge between two vertices v_a and v_b if and only if $\{v_a, v_b\}$ if an element of $\vec{L}_2$ and the corresponding bit in $\vec{S}(\lambda)$ is 1. Since the entries of $\vec{S}(\lambda)$ are chosen independently as 1 with probability λ, $F_n(G_1, \vec{L}_2, \vec{S}(\lambda))$ represents an instance of bond percolation with bond probability λ over the set of vertices in G_1 with the set of attempted edges defined by the elements of $\vec{L}_2$. For $\vec{L}_2$ to be a valid list, a vertex can appear in $\vec{L}_2$ no more than n times which implies that the attempted graph of this bond percolation instance has maximum degree n. For bond percolation on a graph with maximum degree n, if the bond probability is

less than $1/(n-1)$, with high probability, the largest connected component is of size $o(N)$. This proves Lemma 3.

From Lemma 1, $LCC(D(G_0, \vec{L}_0, \vec{S}(\lambda))) \leq LCC(F_1(G_0, \vec{L}_0, \vec{S}(\lambda)))$ for any "destructive two node fusion" D. This implies, from Lemma 2 that $LCC(D(G_0, \vec{L}_0, \vec{S}(\lambda))) \leq LCC(F_1(G_0, \vec{L}_0, \vec{S}(\lambda))) \leq nLCC(F_n(G_1, \vec{L}_1(\vec{L}_0), \vec{S}(\lambda)))$ where G_0, G_1, $\vec{L}_0$ and $\vec{L}_1(\vec{L}_0)$ are described in the lead-up to Lemma 2. From Lemma 3, $LCC(F_n(G_1, \vec{L}_1(\vec{L}_0), \vec{S}(\lambda)))$ = o(N) if $\lambda < 1/(n-1)$, we obtain the following theorem.

Theorem 1 Starting with N isolated clusters of n vertices each where n is finite, the size of the largest connected component obtained from any set of attempted "destructive two-node fusions" that succeed with probability λ is $o(N)$ for $\lambda < 1/(n-1)$.

3.4.4 Photon loss

In this section, we present a lower bound on the percolation threshold $\lambda_c^{(n)}$ in the presence of photon loss. Our results suggest that in the presence of loss, there may be an optimum size of the input microclusters that achieves the lowest fusion success probability necessary for achieving percolation, and hence allows for the greatest tolerance to photon loss. We use a loss model inspired by a recently proposed method to produce photonic microclusters using quantum dot emitters [70, 71], a variant of which has also been proposed for nitrogen vacancy (NV) centers in diamond [83]. In this method, a quantum dot-confined electron is replaced by a confined dark exciton and this dark exciton subsequently interacts with a series of single photons that are initially unentangled. Using such interactions, the creation of an n photon microcluster in the line lattice graph state involves n entangling operations. We assume that each photon produced by the source experiences the same transmissivity η_0^n with $\eta_0 < 1$, and that detector and waveguide losses are lumped into the parameter η_0. The rationale behind this stems from the assumption that the exciton loss acts independently on each photon and that the entire microcluster needs to be produced at the same time: the transmissivity experienced by the k^{th} photon, $\eta_k = \eta_{\text{exciton}}^k \eta_{\text{waveguide}}^{n-k} = \eta_0^n r^k$, with $\eta_0 = \max(\eta_{\text{exciton}}, \eta_{\text{waveguide}})$ and $r = \min(\eta_{\text{exciton}}, \eta_{\text{waveguide}})/\max(\eta_{\text{exciton}}, \eta_{\text{waveguide}})$.

Since $r \leq 1$, $\eta_k \leq \eta_0^n$, $\forall k = 1, \ldots, n$. Therefore assuming that each photon in the n-photon microcluster experiences identical transmissivity η_0^n is an optimistic model which leads to a higher inferred graph connectivity in the post-fusion cluster compared to the true connectivity. Since we are seeking a lower bound on λ_c, this is acceptable.

In the absence of any photon loss, starting with N entangled microclusters of n photons each, the minimum value of two-photon fusion success probability *necessary* to obtain an $O(N)$ photon connected component satisfies the lower bound $\lambda_c \geq 1/(n-1)$, which if our conjecture explained above is true is also *sufficient* (achievable) for percolation. In the presence of photon loss, the above lower bound on λ_c remains a valid, yet trivial, lower bound. We would like a non-trivial lower bound on λ_c that is a function of η_0 and n, such that the lower bound increases with decreasing η_0.

Let us say the success probability of a two-photon fusion operation is λ. As discussed above, there are two types of photons, ones that are measured in fusion attempts and ones that are not. The latter type of photons constitute the renormalizable percolated giant component when $\lambda > \lambda_c$. In the presence of losses, both types of photons undergo loss. Loss of a photon that was measured in a (destructive) fusion attempt is *detected*, since the number of expected detector clicks at the output of the linear-optical circuit for fusion is lower than that is expected. On the other hand, the loss of the unmeasured photons cannot be detected (assuming we do not have access to a quantum non-demolition measurement). This results in the post-fusion cluster to be in a mixed state, a probabilistic mixture of all possible combinations of the unmeasured photons being lost or not. It is not known whether such a mixed state cluster (i.e., without the knowledge of which of the unmeasured photons were lost)—even if percolated—can be renormalized into a logical lattice or not, unless each photon (qubit) in the model considered in this chapter is replaced by a loss-protected logical qubit, e.g., a tree qubit [57]. However, since we are seeking a lower bound on λ_c, we only need to consider a pure graph state that is more connected that the true post-fusion cluster. The simplest way to do so is to pick the post-fusion cluster state where none of the unmeasured photons were lost. With these assumptions, each

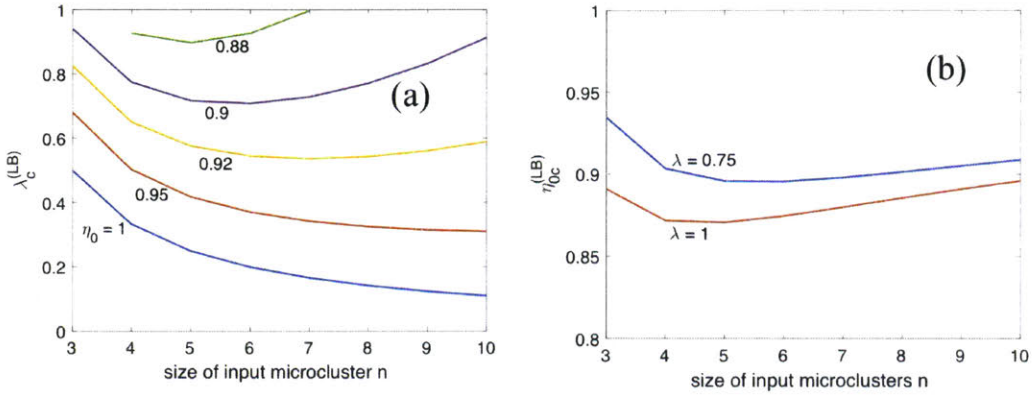

Figure 3-11: (a) A loss-dependent lower bound $\lambda_c^{(\mathrm{LB})}$ on the critical fusion probability λ_c as a function of the input microcluster size n for different values of η_0; (b) a loss-dependent lower bound $\eta_{0c}^{(\mathrm{LB})}$ on the critical loss parameter η_0 as a function of n for different values of fusion success probability λ.

fusion (a bond in the logical graph) succeeds with probability λ, η_0^{2n}. Therefore, following the arguments in Section 3.4.2, if $\lambda < 1/\left[(n-1)\eta_0^{2n}\right]$, the post-fusion cluster state cannot have a connected component with $O(N)$ unmeasured photons. Hence, we have the following loss-dependent lower bound: $\lambda_c \geq 1/\left[(n-1)\eta_0^{2n}\right] \equiv \lambda_c^{(\mathrm{LB})}$. This lower bound is plotted in Fig. 3-11(a) for different values of η_0. We find that while increasing the size of the input microclusters n in the lossless case ($\eta_0 = 1$) always results in a reduction in the necessary fusion success probability for percolation, in the presence of finite losses ($\eta_0 < 1$), there is an optimum value of n that gives the minimum fusion probability. For example, for $\eta_0 = 0.9$, $n = 6$ sized microclusters yield the lowest necessary fusion success probability threshold for percolation.

Conversely, for a given fusion success probability λ, there exists a threshold η_{0c}, s.t., if $\eta_0 < \eta_{0c}$, the post-fusion cluster cannot be percolated. We thus have a lower bound $\eta_{0c} \geq \eta_{0c}^{(\mathrm{LB})}$, where $\eta_{0c}^{(\mathrm{LB})} = [1/\lambda(n-1)]^{1/(2n)}$. In Fig. 3-11(b), we plot $\eta_{0c}^{(\mathrm{LB})}$ for different values of n, for $\lambda = 0.75$ and $\lambda = 1$. There is an optimum value of n which gives the best loss tolerance, e.g., for $\lambda = 0.75$, six photon microclusters gives the best loss tolerance of $\eta_{0c}^{(\mathrm{LB})} = 0.8957$ which corresponds to a loss of 48.36% seen by each photon. Furthermore, we find that going from $\lambda = 0.75$, which is attainable using four single ancilla photons and (lossless) linear optics [36] to deterministic fusion ($\lambda = 1$),

$\eta_{0c}^{(\mathrm{LB})}$ only decreases slightly, i.e., the equivalent per-photon loss threshold increases from 0.896 to 0.871. Hence, when losses are accounted for in ballistic cluster state creation, the advantage in having a fully deterministic fusion may be relatively small.

Finally, it may be possible to get a tighter lower bound on λ_c by using a more sophisticated loss model. For example, the assumption that the exciton loss acts independently on each photon is not entirely accurate, resulting in positively-correlated bonds within the microclusters. Further, the assumption we made about all unmeasured photon not being lost may affect the tightness of the lower bound. However, this last assumption may not have been that ominous, considering our conjectured tightness of the lower bound $\lambda_c \geq \lambda_c^{(\mathrm{LB})} = 1/(n-1)$ in the lossless case ($\eta_0 = 1$) was based on a construction where the fraction of photons α in the logical graph that are left unmeasured goes towards zero.

What we leave unaddressed in this section, are constructive solutions for ballistic photonic quantum computing, or achievability results (i.e., upper bounds on λ_c for a given η_0 or upper bonds on η_{0c} for a given λ) in the presence of photon loss. This will require one to determine how to construct a loss-error-corrected logical lattice fully ballistically (perhaps using tree error correction but with randomly-grown trees) the percolated instance of which can be provably renormalized into a logical cluster state, every node of which is an appropriately loss-protected photonic qubit. In recent work [4, 84], Gimeno-Segovia estimated loss tolerance of ballistic creation of certain percolated lattices by a strategy where one measures all the neighbors of lost photons in the Z basis. This method also accounts for losses in the photons undergoing fusion operations but not in the photons that remain unmeasured, thereby also not proving achievability results in the presence of photon loss. This is an important question that must be addressed systematically not only for photon loss, but for other forms of qubit error models specific to linear-optical qubits such as mode mismatch and detector dark clicks, for this scheme to become a practically feasible solution to scalable quantum computing.

3.5 Conclusions and open problems

In this chapter, we analyzed fundamental thresholds on the success probability of two-photon linear optical fusion operations for preparing large renormalizable photonic clusters for universal cluster-model quantum computing. We introduced a new percolation framework to study this problem, based on which we developed new constructions with improved thresholds and geometric properties over known results, and found fundamental bounds on the thresholds. We also discussed how losses—inline losses and losses at sources and detectors—affect the bounds on the percolation threshold, using loss models inspired from a recently-proposed method to produce photonic microclusters using quantum dots.

Many interesting open questions remain. One major fundamental open question is the minimum overhead required (i.e., number of physical photons in a logical qubit) to error correct for a given amount of loss rate on each physical photon. There has been considerable research to date on quantum error correction to tackle optical loss. This includes the work on tree codes [57, 56, 2] and surface codes [85]. The percolation approach discussed in this chapter is another way to code for optical loss, but as discussed in the previous section, more work needs to be done to design fully ballistic (feedback-free) constructions for renormalizing an error-free logical cluster for quantum computing. One way to do this would be to replace each physical photon in the construction discussed in this chapter by loss-protected photonic qubits, e.g., using tree clusters.

Furthermore, a big practical challenge in making scalable photonic quantum computing feasible is to develop error correction techniques to correct other (non loss) errors, the two most important being mode-mismatch errors and multi-photon events—both in the sources as well as in the detectors (i.e., dark clicks). The whole construction described here relies on a perfect interferometer processing many pairs of identical photons [17, 86, 87]. Mode mismatch can be caused due to the interfering photons not being in identical modes, or small errors and deviations in the splitting ratios of beamsplitters and phase errors. Our work reinforces the message from the recent work

of Rudolph and colleagues, that sources that can directly generate entangled clusters of a small number of photons would be a very valuable resource, and developing new ideas and designs of such photonic sources would be an extremely worthwhile pursuit.

Chapter 4

Percolation based architecture for cluster state quantum computation using photon-mediated entanglement between atomic memories

4.1 Introduction

The past years have seen rapid advances in controlling small groups of qubits encoded in atomic or atom-like quantum memories. An important question now concerns the development of architectures to efficiently combine these memories into large-scale systems capable of general-purpose quantum computing [88, 89, 90, 91], quantum simulation [92, 93], and measurement near the quantum limit [94]. A promising approach is entangling the atomic qubits with optical links to generate cluster states. Cluster states with adaptive measurements can perform general-purpose quantum computing [1]. A key challenge is to produce this cluster state fast enough to allow the one-way quantum computing and error correction within the finite coherence time of the memory. Large entangled states of this type also serve as resources for other quantum information processing applications, including quantum simulation

and quantum precision measurements.

Here, we show that percolation of heralded entanglement allows us to create arbitrarily large cluster states. This process is fast enough for implementation with device parameters that have been demonstrated; one does not need high cooperativity cavities, ancilla single photons, or time-consuming feed-forward operations. Furthermore, as opposed to previous schemes, we do not require error correction to account for missing bonds, and instead use renormalization [63, 4, 65] which can be done with constant overhead [95] if the bond probability exceeds the percolation threshold. The percolation approach also provides tolerance for site imperfections in several different lattice architectures. When combined with our novel transparent node architecture that allows long range connections, a further reduction in the percolation threshold can be achieved. We also found a theoretical limit of the percolation threshold across different geometries and found that our proposed geometries are within a factor 1.6 of the limit. Our approach applies to a number of leading physical qubit systems, including atomic gases [96, 97], ion traps [98], semiconductor quantum dots [99, 100, 101], or rare earth ions [102], though for clarity, we focus here on nitrogen vacancy (NV) centers in diamond [103, 104].

NV centers in diamond have many properties that make them favorable as a quantum memories. The NV^- charge state has a robust optical transition for heralded entanglement between distant NV centers [105, 106] and a long electronic spin (S=1) coherence time [107] for high-fidelity qubit operations. Recently, single qubit gates with fidelities up to 99% were achieved with optimal control techniques [108, 109]. NV centers can be coupled with nearby nuclear spins [110], which have coherence times exceeding one second even at room temperature [111]. The electronic spin state can be transferred to the nuclear spin, and measurement of nuclear spin with high fidelity is possible by repetitive measurement of the electron spin [112]. In addition, the memories are solid state and can be coupled with integrated photonic devices [113] making it a promising platform for large scale quantum computation.

Fig. 4-1 illustrates the percolation approach to cluster state generation with quantum memories. We work in the framework of cluster states where nodes represent

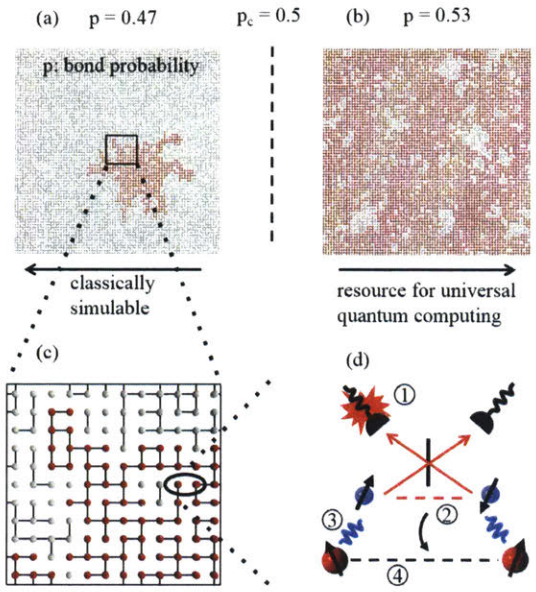

Figure 4-1: Cluster state generation by percolation. (a),(b) Transition in the size of the largest connected component (LCC) with increasing bond probability. Spheres and lines represent nodes and bonds respectively, and the red spheres represent the LCC. When the bond probability (p) goes above the percolation threshold (p_c), the size of the LCC suddenly increases and the cluster changes from being classically simulable to a resource for universal quantum computation. (c) Expanded view of (a). (d) Physical implementation of nodes and bonds with NV centers in diamond. ① Probabilistic Bell measurement (Barret-Kok protocol) is attempted on two nearest-neighbor broker qubits (electronic spins, blue spheres). ② Conditioned on one photon detection events, the two broker qubits are entangled onto a Bell state. ③ Hyperfine interaction between electronic spins and nuclear spins (client qubits, ^{15}N) mediates controlled-Z gates. ④ X-basis measurement of electronic spins projects nuclear spins into an entangled state heralded by the measurement results (entanglement swapping)

qubits in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ and edges/bonds represent controlled-Z (CZ) gates between neighboring nodes. Consider a square lattice where every edge exists with probability p as shown in Fig. 4-1(a)-(c). The computational power of such a system is related to the size of the largest connected component (LCC) in the cluster (shown in red). When $p < 0.5$, the clusters form small disconnected islands. In this regime, for a lattice with N nodes, the size of the LCC is $O(\log(N))$ [114]. Local measurements on such a lattice can be efficiently simulated classically and hence, the resource is not sufficient for quantum computing [64]. When the bond probability exceeds 0.5 there

is a sudden transition in the size of the LCC: the number of nodes in the LCC is now $\Theta(N)$. This is accompanied by a sudden transition in computational power; single qubit adaptive measurements on this cluster have the power of universal quantum computing [95]. The bond probability p_c at which the transition takes place is called the percolation threshold. Square, triangular and hexagonal lattice clusters above the percolation threshold are resources for universal quantum computation [63, 95], although there are examples of other lattices for which this is not true e.g. Bethe lattice clusters are not resources for universal quantum computation [79].

4.2 Creating photonic entanglement

Figure 4-1(d) shows the physical implementation of the link creation with NVs. The nuclear spins (red spheres) function as "client qubits" that store entanglement. They are coupled to the NV electronic spins — "broker qubits" — that are entangled remotely by Bell measurements mediated by photons. In each time step, we attempt to create one edge (entanglement) at each node by heralded entanglement mediated by photons. To be specific, we consider the Barret-Kok entanglement protocol [115] on the broker qubits of neighboring nodes/sites. If the probabilistic Bell measurement succeeds, the electron spins of the corresponding NVs are entangled. This entanglement is then transferred to the nuclear spin with an entanglement swapping procedure, as illustrated in Fig. 4-1(f) and described in detail in section 4.2.1 and Ref. [88]. If the Bell measurement fails, we wait for the nuclear spin and electronic spin to be decoupled, which happens after a time period of the hyperfine interaction, to avoid nuclear spin dephasing. The whole cycle from initialization to entanglement swapping takes approximately $t_0 = 5$ μs based on recent experimental demonstrations [106].

The Barret-Kok protocol to generate entanglement is advantageous because it does not require ancilla single-photons or high cooperativity cavities. Furthermore, photon loss in this scheme does not degrade fidelity, which is critical to the error correction overhead. This increased fidelity comes at the price of low bond success probability (detailed in Table 4.1) which is a problem for conventional architectures. This can be

overcome in our percolation based architecture. However, the imperfect/raw cluster produced in the percolation approach proposed here will need additional renormalization, which can be done with a constant overhead [95]. This point will be discussed in the conclusion in more detail.

The bond lengths could practically be very short, on the order of tens of microns, so that the entire cluster may be integrated on a chip, as illustrated in Fig. 4-3. Each node in the architecture requires an atomic memory and a $1 \times d$ switch, where d is the degree of the lattice being attempted (4 for the square lattice). Each edge in the lattice requires waveguides between the nodes, a beam-splitter, and two detectors to implement the Bell measurement.

At each time step in the operation of the system, each atomic memory generates a photon entangled with the electron spin that is routed towards one of the d neighbors for the Bell measurement, so that for n_{att} time steps for entanglement generation, entanglement with each of the d neighbors of a node is attempted n_{att}/d times. It is important to synchronize neighboring switches such that both photons required to attempt an edge arrive at the same time. For the hexagonal ($d = 3$), square ($d = 4$) and triangular ($d = 6$) lattices, synchronization is straightforward. Each switch only needs to be flipped $d-1$ times during entanglement creation, and hence the switching time is negligible, since electro-optic modulators can switch at sub-ns time-scales and the entanglement generation time varies from ~ 0.1 ms to ~ 100 ms, depending on the coupling scheme.

4.2.1 Probabilistic Bell Measurement for broker qubit entanglement

Broker qubits (electronic spin of NV centers) are entangled through Bell measurement (Barret-Kok protocol [115]), as used in [105]. The detailed pulse sequence is described in Fig. 4-2. First, $\pi/2$-pulses are applied to both broker qubits. These pulses should be strong enough to drive both hyperfine levels by power broadening, and are represented in light blue, in contrast to the selective pulses denoted in dark blue. An optical π-

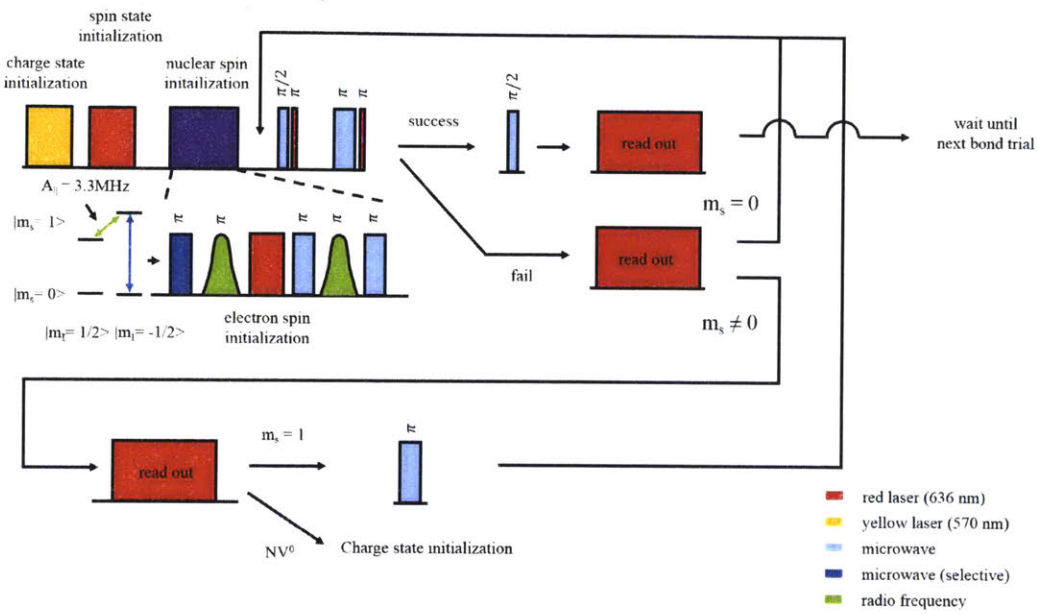

Figure 4-2: Pulse sequence for the broker qubit entanglement procedure.

Table 4.1: Performance metrics for different collection schemes

Collection	Bullseye	Waveguide	Cavity
Bond success prob. p_0	5×10^{-5}	2×10^{-4}	5×10^{-2}
Bond trial time t_0 [1]	5 μs	5 μs	5 μs
Detector jitter	$5 - 10ps$	$5 - 10ps$	$5 - 10ps$
Spectral diffusion	$\times$	$\times$	✓
Readout time	4μs	800 ns	400 ns
Demonstration	Achieved	Achievable	Long-term

pulse then reads the states of both spins. Emitted single-photons from the the two electronic spins are then detected after passing through a 50/50 beam-splitter as shown in Fig. 1(d). The beam splitter erases the information of which emitter the detected photon comes from. As a result, a one-photon detection event projects the two spin states into either $|\psi^+\rangle$ or $|\psi^-\rangle$ depending on which which detector fired. However, photon loss can cause a two-photon generation event, where the state of the emitters is projected onto $|00\rangle$, to be misinterpreted as a one-photon detection event. In order to overcome this issue, we apply microwave π pulses to each broker qubits. If the emitter was originally projected onto $|00\rangle$, this operation changes the emitter

state into $|11\rangle$. Emission from the two electronic spins is again detected after passing a 50/50 beam splitter. In the case of the desired $|\psi^+\rangle$ or $|\psi^-\rangle$ state and no photon loss, a one-photon detection event is registered. The unwanted $|11\rangle$ state or photon loss leads to no photon detection which we treat as a failed entanglement attempt. Overall, one-photon detection events in both steps makes sure that the two broker qubits are in either $|\psi^+\rangle$ or $|\psi^-\rangle$.

The success probability of the the broker qubit entangling operation p is determined by the photon collection efficiency η: $p = \frac{1}{2}\eta^2$ [115]. The entanglement probability per attempt is low, but success results in high-fidelity entangled states. This trade-off between success probability and fidelity makes this protocol appealing for our percolation-based architecture.

Indistinguishability of photons serves a critical role in Bell measurement. Thus, the wavepackets of the two photons should be perfectly overlapped in the time domain. However, two photons from different emitters should also be overlapped in the frequency domain. In the next paragraphs, we will investigate methods for overcoming two types of frequency mismatch [116]: static mismatch and dynamic mismatch.

Static mismatch results from the inhomogenous distribution of the ZPL of NVs, typically caused by local strain. This results in two nearby NVs being detuned upto $\sim$ 5 GHz [105]. A common misconception is that photon-mediated entanglement of two emitters requires that photons from the two emitters have exactly the same frequency. However, this may not be the case if the timing resolution of the photodetectors is much smaller than the inverse of the frequency difference between the emitters [116]. After detection, the photons are projected onto a narrow temporal wavepacket that depends on the detector jitter. This broadens the distribution in the frequency domain increasing indistinguishability. In other words, the phase difference originating from the difference in photon energy is heralded by the detection time [116], and can be corrected. In this work, a timing resolution of 5 $\sim$ 10 picoseconds is assumed, which has been demonstrated with superconducting nanowire single photon detectors (SNSPD) [117]. This detector can be both free-space coupled and waveguide-integrated. This detector resolution can easily handle $\sim$ 5 GHz static

mismatching. Alternatively, DC Stark shift can be used to match the ZPL frequency of two NVs [118, 105, 119].

Dynamic mismatch is the frequency mismatch originates from fluctuations in the time between different measurements. In contrast to static mismatch, one cannot unwind the phase difference due to dynamic mismatch with good timing resolution detectors, because the frequency difference is unknown. In the case of emitters, this spectral diffusion usually arises from the charge environment of etched surfaces of the nanostructures. As shown in Table 4.1, we have assumed that waveguide structures, in which an etched surface is a few hundred nanometers away from the emitter, do not produce spectral diffusion, while in cavity structures, etched surfaces that are only tens of nanometers from an emitter do.

Spectral diffusion can be mitigated by the Purcell effect. As analyzed in [116], the Purcell effect makes the emitter decay faster, so that the effects of spectral diffusion are minimized. A free space lifetime of $\sim$ 10 nsec can tolerate $\sim$ 100 MHz spectral diffusion. If we assume $C \sim 1$, the lifetime is reduced by half which can tolerate $\sim$ 200 MHz spectral diffusion.

Chemical surface treatment, such as high temperature annealing [120], for reducing spectral diffusion is an area of active research. Spectral diffusion of $\sim$ 1 GHz has been observed in NVs near the surface [121, 122]. We expect a rapid advance in this direction to reduce spectral diffusion under 200 MHz.

Alternatively, we could time-gate the photo-detection by post selection, to fill this gap ($\sim$ 200 MHz vs $\sim$ 1 GHz) e.g. we could only count single-photon detection events that occur within $\sim 1/5^{th}$ of the lifetime as successes and discard the rest. This will reduce the unknown phase evolution by a factor of 5. Although this will reduce collection efficiency by 80 %, we can still achieve percolation within the coherence time of the nuclear spin.

4.2.2 Entanglement swapping procedure

When entanglement between electronic spins is successful, the entanglement is transferred to the nuclear spin of NV centers (^{15}N), with a controlled-phase (CZ) gate

followed by X basis measurement [88]. In this way, one can make use of the long coherence time of the nuclear spin with the spin-photon interface of the electronic spin.

Physically, the CZ gate is implemented with a timed-hyperfine interaction with $A_{||}\tau = \pi$, where τ is the hyperfine interaction time. Since neighboring nodes should be synchronized, the ^{15}N ($A_{||} = 3.3$ MHz) nuclear spin is used instead of ^{13}C, because the ^{13}C nuclear spin is coupled to the electronic spin with a random strength. ^{15}NV can be produced by implantation [123], and its position can be controlled within 10 nm [124], which is desirable for integration with photonic circuits.

If the probabilistic electronic spin entanglement fails, the next electronic spin entanglement is attempted when $A_{||}\tau = 2n\pi$, where n is an integer, to ensure that the electronic spin is not entangled with the nuclear spin in the case of a failed attempt.

Both electronic spin entanglement and the swapping procedure heralds the exact state of the cluster. Since the heralded state only differs by local Cliffords, they can be compensated in the adaptive measurement steps in the actual computation.

4.2.3 Details on the timing of entangling operation

The detailed pulse sequence for the whole procedure is described in fig. 4-2. The charge state is prepared with yellow laser resonant excitation of the NV^0 ZPL, to minimizing spectral diffusion [120]. The electronic spin state is then initialized to $|m_s = 0\rangle$, by resonantly driving $|m_s = +1\rangle$ state through the A1 optical transition [116]. Alternatively, readout with a conditional $\pi/2$ pulse can be used. ^{15}N nuclear spin can be initialized to $|m_I = +1/2\rangle$, by swapping the nuclear spin with the electronic spin with a nuclear spin conditioned-electronic spin π pulse and electronic spin conditioned-nuclear spin π pulse, and electronic spin initialization. $|+\rangle_I = (|m_I = +\frac{1}{2}\rangle + |m_I = -\frac{1}{2}\rangle)/\sqrt{2}$ is then prepared by an electronic spin π pulse, nuclear spin $\pi/2$ pulse and electronic spin π pulse. This whole initialization process is tried only once at the very beginning of the cluster state generation. We assume that all the resonance frequencies are pre-characterized, so that there are negligible errors in

the driving.

As specified in the preceding sections, electronic spin entanglement is attempted by two Bell measurements. When the electronic spin entanglement fails, the subsequent electronic spin entanglement attempt should be delayed by $\tau = 2n\pi/A_{||}$, where n is an integer to prevent entanglement of the nuclear and electronic spin. When electronic spin entanglement is successful, the electronic entanglement is swapped to the nuclear spin, the electronic spin is initialized with a conditional π pulse to prevent further hyperfine interaction.

Most of the time, the Bell measurement fails because of low collection efficiency. In this case, electronic spin states are measured and initialized by a conditional optical π pulse. It cannot be directly initialized, as with the very first initialization, because the hyperfine interaction is always on, when electronic spin is non-zero. Therefore, only when the nuclear spin is disentangled with the electronic spin, i.e $\tau' = 2m\pi/A_{||}$, where m is an integer, spins are flipped by a timed-π-pulse.

In this work, we assumed that two-photon absorption is small enough so that we can ignore charge state conversion, or can account for it with error correction. However, this can be also overcome by (1) reducing resonant laser power and using a weaker pulse, or (2) a timed pulse sequence that can initialize the charge state, while decoupling NV^0 electronic spin with nuclear spins. This can be done by resonant yellow laser driving and 1064 nm pumping [125, 126] when $\tau'' = 2l\pi/A'_{||}$, where $A'_{||} = 35.7$ MHz, l is an integer. This procedure needs repetition because unlike the resonant excitation, pumping electrons from valence bands is probabilistic.

4.2.4 Collection efficiency and entanglement success probability

The probability of successfully heralding the entanglement of two NV centers is $p_0 = \eta^2/2$ [115], where η is the efficiency of emitting, transmitting, and detecting the photon entangled with the electronic spin (zero phonon line, ZPL) from the NV excited state. Table 4.1 summarizes p_0 for three representative types of NV-photon interfaces: low-

efficiency interfaces with $p_0 = 5 \times 10^{-5}$ representative of today's state of the art circular gratings or solid immersion lenses (SILs) [127, 106, 128]), medium-efficiency interfaces with $p_0 = 2 \times 10^{-4}$ for NV centers coupled to diamond waveguides [129, 113], and high-efficiency $p_0 = 5 \times 10^{-2}$ for nanocavity-coupled NV centers [130]. For all three coupling mechanisms, we assumed coupling efficiencies that are realistic today, which is discussed in detail below. After n_{att}/d entanglement attempts with a nearest neighbor, the probability of having generated a bond is $p = 1 - (1 - p_0)^{n_{\text{att}}/d}$.

The single-photons to be used for entanglement should be spectrally located in the zero phonon line (ZPL). The Debye-Waller factor, the ratio of ZPL emission rate to the total decay rate of the NV center varies from $3 \sim 5\%$ [105]. Here, we will use a conservative number: 3%. In addition to the low fraction of emission in the ZPL, diamonds suffers from low photon collection efficiency because of total internal reflection caused by the high refractive index of diamond. To overcome these issues, various nanophotonic structures for coupling to free space or photonic integrated circuits have been proposed. Here, we consider three collection schemes: free space collection with bullseye grating or solid immersion lens (SIL) [127], waveguide collection [113], and high-cooperativity cavity collection [130].

The Bullseye grating is a circular grating that maximizes free space coupling by interference. Record count rates ($\sim 2.7 \times 10^6$) have been reported with such structures, and simulations have shown $\sim 30\%$ collection efficiency at the ZPL wavelength. For shallow NV centers within $\sim$ 10 nm from the surface, $\sim 90\%$ collection efficiency has been estimated. However, we will assume a more conservative number of $\sim 30\%$ collection efficiency, which results in a ZPL photon collection efficiency of 1%.

Collection of photons from solid state emitters into waveguides coupled with photonic integrated circuits has been successfully demonstrated [131, 113]. These collection schemes are more scalable because single photon detectors, filters and beam splitters can be implemented on-chip. In addition, the well-confined optical field makes near unity detection efficiency possible with superconducting nanowire single-photon detectors (SNSPDs) as demonstrated in [132]. For NV centers, a hybrid approach with silicon nitride (SiN) waveguides successfully showed that the coupling efficiency

can be upto 83%. In this work, a diamond waveguide is *picked and placed* on a SiN waveguide [113]. Considering the errors in this process, we assumed a conservative photon collection efficiency of $\sim 66\%$ including detection efficiency.

Lastly, photonic cavities can re-distribute the spectrum of emitted photons, extract them out of the emitter and direct them into the waveguide. When the cavity is tuned to the ZPL wavelength, the Purcell effect enhances emission into the ZPL, resulting in higher conversion efficiency from emitter excitation to ZPL photon. This coherent interaction between the cavity and emitter has been successfully demonstrated in many platforms, and is an area of active research [130, 131, 133, 134]. We used a number for cooperativity already demonstrated in ref. [130]: $C \sim 1$. In this paper, a 1-dimensional photonic crystal cavity, which is a natural design for coupling light into a waveguide, was used. By removing a few holes in the Bragg reflector, one can direct photons from the cavity into the waveguide. Including this conversion efficiency, 32% overall collection efficiency is assumed in this chapter.

Table 4.2: ZPL Collection efficiencies of different schemes

Collection	Bullseye grating (or SIL)	Waveguide	Cavity
ZPL Enhancement	×	×	~33
Fraction of Emission in the ZPL (%)	3	3	50
Photon Collection Efficiency (%)	33	66	64
ZPL Collection Efficiency (%)	1	2	32

4.3 Percolation threshold

We performed simulations using the Newman-Ziff algorithm [77] with 9 million nodes to evaluate the growth of the clusters, based on the entanglement success probability per attempt. Fig. 4-4(a) plots the fraction of nodes that are within the largest cluster component (f_{LCC}), as a function of time from the start of the protocol for the three values of p_0, assuming $t_0 = 5$ μs. In Fig. 4-4(a), the underlying geometry is a square lattice. The lines represent different Bell measurement success probabilities corresponding to the coupling mechanisms in Table 4.1. Initially, f_{LCC} is $\mathrm{O}(\log(N)/N)$ [114] where N is the total number of nodes in the lattice. As the

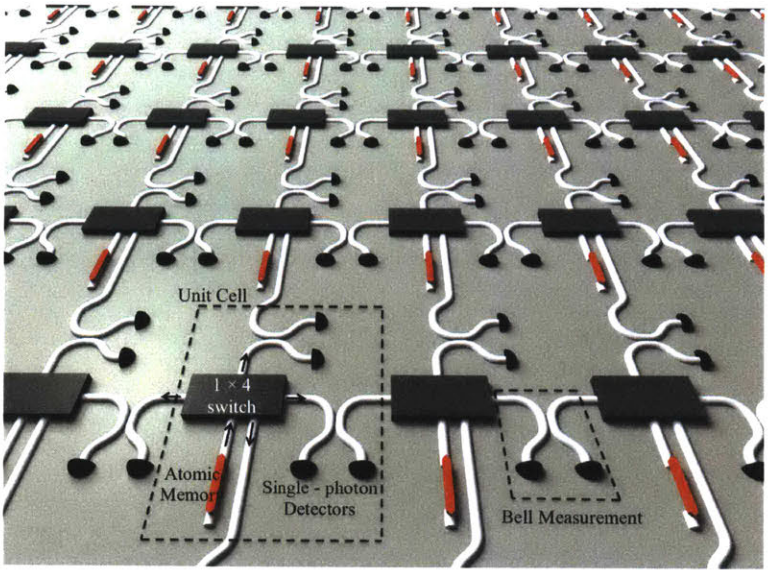

Figure 4-3: Physical implementation of the proposed architecture. A unit cell consists of an atomic memory, a 1×4 switch, waveguides and 4 single-photon detectors. Single-photons emitted from the atomic memory are coupled to the waveguide and directed to the switch. The switch chooses one of the nearest-neighbor nodes to be entangled with, and single-photons are interfered using a 50/50 beam-splitter. Single-photon detectors detect interfered photon projecting electronic spins onto an entangled state.

bond success probability passes the bond percolation threshold (p_c), f_{LCC} rapidly rises and becomes $\Theta(1)$. For a degree d lattice, the bond probability after time t is $p = 1 - (1 - p_0)^{t/t_0 d}$. From this, we can calculate that the time required to obtain a resource for universal quantum computation is $t_c = t_0 d \ln(1-p_c)/\ln(1-p_0)$, which is depicted with the vertical dashed lines in the figure. The transition becomes sharper as the number of nodes in the lattice (N) increases.

In all three collection schemes, the bond success probability exceeds the percolation threshold within 1 second, which is the experimentally demonstrated coherence time of a nuclear spin coupled to the electronic spin of the NV [111]. These simulations reveal, surprisingly, that even with free space coupling without any ZPL enhancement, an arbitrarily large cluster can be generated.

It is well known from percolation theory that higher connectivity between nodes can reduce the percolation threshold. Does the time to exceed threshold (t_c) change significantly with the lattice degree d? As shown in Figure 4-4(b), which plots f_{LCC}

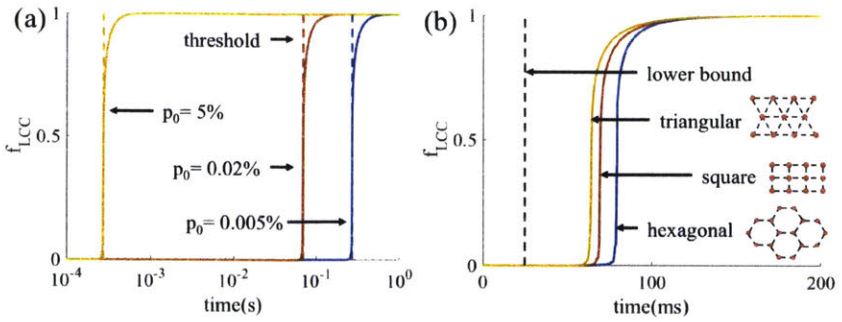

Figure 4-4: Size of the largest connected component vs time for (a) different values of p_0, the probability of successful Bell measurement in one attempt and (b) different underlying lattice geometries. A square lattice is used in (a) and $p_0 = 0.02$ % is used for (b).

for two additional lattice types, triangular ($d = 6$) and hexagonal ($d = 3$), t_c is nearly the same for the three lattice types. The reason is that increasing d lowers the bond percolation threshold, but it also decreases the number of entanglement attempts between NVs, which is n_{att}/d. This is because a single broker qubit per NV requires entanglement attempts to proceed serially. Increasing d would in fact substantially lower t_c if each site contained d broker qubits that could be entangled simultaneously.

4.4 Lower bound on the minimum time required

Let us consider the most general scenario of fully controllable connectivity in the graph, i.e., we can attempt Bell measurements on any pair of NVs in a time step. What is the minimum time, $t_c^{(LB)}$ ($\min(t_c)$), required to obtain a resource for universal quantum computation, optimizing over all lattice geometries, if the bonds are attempted without feed-forward? The bond probability after time t is $p = 1 - (1 - p_0)^{t/t_0 d}$. For percolation, $p \geq p_c$ i.e. $t_c \geq dt_0 \ln(1 - p_c)/\ln(1 - p_0)$. For a degree d lattice, $p_c \geq 1/(d - 1)$ [81, 82], with equality for a degree d Bethe lattice. This leads to $t_c \geq dt_0 \ln(1 - 1/(d - 1))/\ln(1 - p_0)$. t_c is minimized as $d \to \infty$ in which case we obtain $t_c^{(LB)} = -t_0/\ln(1 - p_0)$ which is the minimum possible time required to generate a resource for universal quantum computing without feed-forward. $t_c^{(LB)}$ is plotted as a black dashed line in Fig. 4-4(b). The lattice corresponding to this threshold is

the infinite-dimensional, infinite-degree Bethe lattice. Such infinite degree lattices are not a resource for universal quantum computing [79], yet, we find that the simple 2D lattices with nearest neighbor connectivity we chose in Fig. 4-4 are only a factor 3 above this limit and are resources for universal quantum computing.

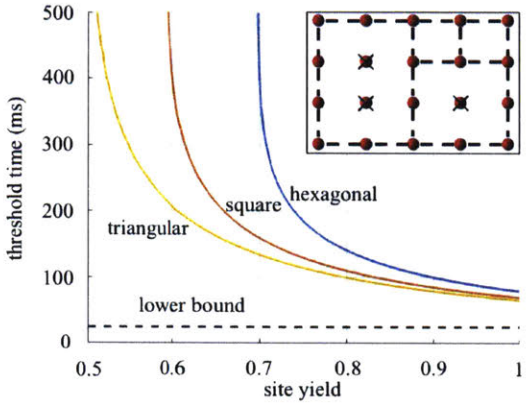

Figure 4-5: The minimum time required to obtain a percolated lattice with sub-unity site-yield. $p_0 = 0.02\%$. The inset shows the bonds that can be attempted in a square lattice if the sites marked with crosses are inactive

4.5 Faulty sites

For practical implementations, it is important to consider the effect of non-functional sites (e.g., a far-detuned NV center or a missing trapped ion). Even if all faulty nodes and their edges are removed, as illustrated in the Fig. 4-5 inset, the lattice can retain enough bonds to give a percolated cluster. In this case, the problem maps to site-bond percolation. This is quantified in Fig. 4-5 where we plot the minimum time required to obtain a percolated cluster as a function of the site-yield (q), assuming NVs coupled to diamond waveguides without a cavity (medium $p_0 = 0.02\%$). In general, a reduced site-yield can be compensated with a larger bond probability which would require a longer time (more attempts) to reach. While the hexagonal, square, and triangular lattices have a similar threshold time t_c when the site-yield is perfect, the tolerance to imperfect site-yield is different. Following the trend of the site percolation threshold (q_c), the triangular ($q_c = 0.5$) performs better than the square lattice($q_c \approx 0.593$),

which performs better than the hexagonal lattice ($q_c \approx 0.697$). The site percolation threshold corresponds to the minimum possible site-yield for percolation with all bonds having succeeded ($p = 1$).

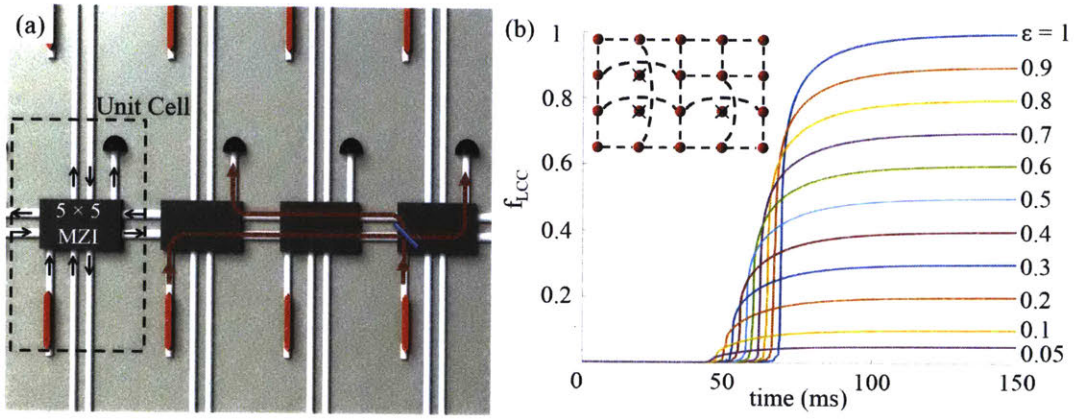

Figure 4-6: (a) A more general architecture with switches replaced by MZI arrays can allow for long-range entanglement as shown here (b) f_{LCC} as a function of time for different values of ϵ. $p_0 = 0.02\%$ is used here.

4.6 Transparent node architecture

The architecture that we have discussed thus far allows for only nearest neighbor connections between atomic memories. In Fig. 4-6, we present a modified architecture that can be used to make long range connections which can in turn decrease the threshold time and increase tolerance to imperfect site yield. Furthermore the architecture reduces the required number of detectors by a factor of 4. In this architecture (Fig. 4-6(a)), the 1×4 switch from Fig. 4-3 is replaced by a 5×5 beam-splitter array with the input and output ports depicted with arrows in the unit cell. The 5×5 beam-splitter array can be used to implement any linear optic unitary between the set of input and output modes [66], and compact on-chip versions have been demonstrated based on cascaded Mach-Zehnder interferometers (MZIs) [17]. The MZI arrays allow us to make long range connections by turning nodes "transparent", entangling distant nodes while maintaining a planar physical architecture.

One way to use this more general architecture is to randomly turn a fraction $1-\epsilon$ of the nodes transparent as shown in the inset of Fig. 4-6(b). The resulting plot of f_{LCC} vs time with $N = 9$ million nodes is shown in Fig. 4-6(b). As ϵ decreases starting from one, the maximum possible value of f_{LCC} is also reduced from one to ϵ because only a fraction ϵ nodes have active qubits. However, because the transparent nodes increase dimensionality while maintaining connectivity, reducing ϵ actually reduces the t_c at which we obtain a $\Theta(N)$ connected cluster for universal quantum computation. Therefore, for a given entanglement generation time, there is an optimum value of ϵ which gives us the largest LCC. We numerically found a minimum possible bond percolation threshold in the transparent architecture of 0.33 which is achieved when $1/N \ll \epsilon \ll 1$, i.e., when $\epsilon \to 0$ but the number of non-transparent nodes in the lattice is still $\Theta(N)$. Faulty sites can be incorporated into the fraction of transparent nodes as long as the yield far exceeds $1/N$ without affecting the minimum entanglement time t_c.

4.7 Conclusions and open problems

In conclusion, we proposed an architecture for quantum computing with atomic memories that uses the concept of percolation to produce a resource for universal quantum computing within the coherence time of the NV nuclear spins, even with imperfect site-yield. Compared to previous scheme which require repeating each bond a large number of times to obtain a bond probability of $p > 0.999$, we find that the required number of attempts is reduced by an order of magnitude, which lowers the requirement on memory coherence time by the same factor. Furthermore, the missing bonds are resolved using renormalization rather than error correction, which can be done with constant overhead. Our scheme does not need high cooperativity cavities or ancilla single photons, and minimizes the amount of feed-forward. The proposed blueprint is applicable to arbitrarily large numbers of qubits. For example, in the planar platform outlined in Fig. 4-6(a), with a realistic lattice spacing of 100-200 μm, a centimeter-scale chip could accommodate on the order of a million qubits.

Although we deal with the problem of ionization and imperfections in the cycling transition, accurate control of the microwave and optical transitions presents a challenge, and is an area of active research [135, 108]. In addition, the internal dynamics of the NV, such as transition from $|m_s = 0\rangle$ electronic excited state to the dark state has not been fully investigated, and could potentially degrade the fidelity of the cluster. Future work is required to make the percolation approach a truly scalable approach for universal quantum computing should focus on determining the most efficient algorithm and calculating the resource overhead required to renormalize a percolated lattice in the form of a quantum error correction code. Because of our architecture's natural tolerance to faulty sites and missing bonds, the size of the required error correction code can be expected to be smaller than conventional architectures. Even without a fault-tolerant method for renormalization, the percolation based approach could be a much easier path for demonstrations of cluster state creation for simulating many-body physics [92, 93], quantum metrology [94, 136] and quantum repeaters [48, 54, 51, 80, 137].

Chapter 5

Routing entanglement in the quantum Internet

A *quantum network* can generate, distribute and process quantum information in addition to classical data [138]. The most important function of a quantum network is to generate long distance quantum entanglement, which serves a number of tasks including the generation of multiparty shared secrets whose security relies only on the laws of physics [38, 39], distributed quantum computing [40], improved sensing [41, 42], blind quantum computing (quantum computing on encrypted data) [43], and secure private-bid auctions [44].

Recent experiments have demonstrated *entanglement links*, viz., long-range entanglement established between quantum memories separated by a few kilometers using a point-to-point optical link [106]. As illustrated in Fig. 5-1, measurements performed at nodes in a quantum network can be used to glue together small entanglement links into longer-distance clusters. The nodes contain quantum memories that store qubits up to their coherence time, sources that generate photons entangled with the quantum memory to be sent to neighboring nodes, and local quantum processors that can perform multi-qubit joint measurements. Entanglement attempts between neighboring nodes are synchronized on a global clock. The *quantum routing protocol* dictates the measurements to be performed locally at each node in order to obtain the desired entanglement topology. Possible goals of a routing protocol could

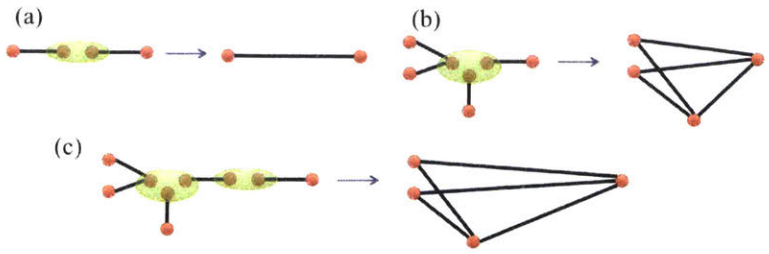

Figure 5-1: Examples of fusing small entangled clusters into larger ones using projective quantum measurements (green ovals) at nodes of a quantum network. Red circles represent qubits and black lines represent entanglement. (a) Two-qubit measurement (Bell state projection) used to connect two entangled links into a longer entangled link; (b) a three-qubit measurement (a GHZ state projection) fuses three clusters (two 2-qubit entangled links and one 3-qubit linear cluster) into one 4-qubit entangled cluster; (c) two adjacent nodes in a network performing a three-qubit GHZ measurement and a two-qubit Bell state measurement simultaneously. The measured qubits are lost, whereas the final entangled state of the unmeasured qubits upon successful completion of both measurements is the same regardless of the order or the simultaneity of the measurements. A quantum measurement at a node may succeed only with a probability, which is a function of the class of optical devices employed to realize the measurement (e.g., linear optics, single photon sources, and single photon detectors) and losses in devices. This figure does not show "failure outcomes", i.e., the resulting entangled state if one or both measurements fail.

be to enable high rate entanglement among multiple user-pairs simultaneously, or to generate multi-partite entanglement (entanglement between three or more parties).

The development of network algorithms and protocols for routing and scheduling information flows was critical for the creation of today's Internet. We expect a similar development in algorithm/protocol design to be critical to design a versatile and high performance quantum network. Some results and intuitions from the theory of classical networks carry over into quantum networking. However, many new challenges arise due to the idiosyncrasies of quantum information. Unlike classical communications, where the rate can be increased by increasing transmit power, photon loss fundamentally limits the entanglement rate over any single link, which must decay exponentially with the length of optical fiber, regardless of the choice of quantum source, the transmit power or the detection strategy [46, 139]. Whereas copying of bits at a network node is common in multipath routing in classical networks [140, 141], copying a qubit is impossible because of the quantum no-cloning theorem [142, 143]. Unlike

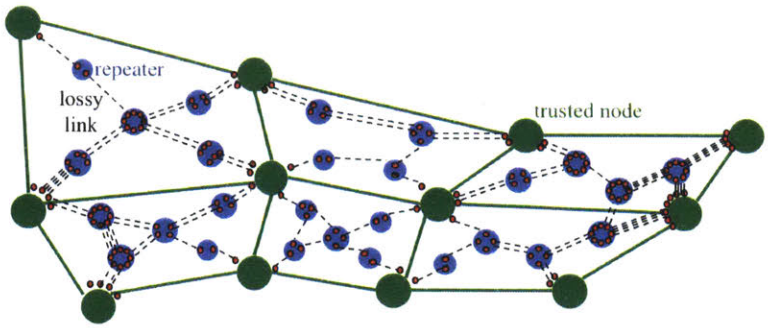

Figure 5-2: Schematic of a general quantum repeater network. The large (green) circles represent 'trusted' nodes, which are connected via a classical network. The blue circles denote repeater stations, and the red circles inside them represent quantum memories. The dashed lines connecting the red circles are independent lossy optical (fiber) channels. In principle, all nodes in the network could be equipped with quantum repeaters (i.e., no trusted nodes), in which case depending upon the need, a node can be a consumer of shared entanglement, or act a router to conduit entanglement flows between other nodes.

classical information flow, an entanglement flow does not have directionality. Rather, entanglement is generated across links all over the network and pieced together to form long-range entanglement. Quantum memories are much shorter lived and expensive compared to their classical counterparts making classical routing strategies such as disruption tolerant routing [144, 145]—where a packet is held by a node for until the desired next-hop link is up—much harder to mimic. Finally, distilling and shaping entanglement among a desired set of nodes from many copies of large (potentially random) entangled clusters is a purely quantum problem that has no classical analog.

In this chapter, we present protocols for repeater nodes to support multiple simultaneous entanglement flows when every node is limited to the same quantum processing used in repeater chains: (probabilistic) two-qubit Bell state measurement (BSM), also called entanglement swapping. BSMs have been experimentally demonstrated in many physical systems [146, 105, 147, 148, 96, 149]. Entanglement attempts between repeater nodes are probabilistic because they are connected via lossy optical links. In every clock cycle, pairs of neighboring repeater nodes attempt entanglement generation. The result of whether an entanglement generation was successful is trans-

mitted back to the corresponding pair of nodes. The repeater nodes then make their local BSM decisions based on this 'local' link-state knowledge, i.e., the successes and failures of entanglements established across their nearest neighbor links. Measurements at different nodes can all be done in parallel because BSMs commute with one another.

One of our interesting findings is that *multi-path routing*, i.e. using multiple paths for routing entanglement in a quantum network, can enable long distance entanglement generation with a superior rate-vs.-distance *scaling* than a single linear repeater chain along the shortest path connecting Alice and Bob [1]. The rate along a single repeater chain falls off exponentially with distance [55, 80]. Multi-path routing reduces the exponent, resulting in an exponential increase in rate with increasing distance. The quantum network we propose uses the same basic elements and operations (probabilistic BSMs) as a linear repeater chain but uses more repeaters. Note that increasing the number of repeaters in a repeater chain would not always increase the rate: given the total end-to-end distance, and given the losses at each node, there is an optimal number of repeaters between the end points of a flow that maximizes the rate, i.e., inserting more nodes along that linear path can actually diminish the rate [55, 80].

Another interesting result is that if the repeater nodes have 'global' link-state knowledge (knowledge of the state of all links in the network) and the entanglement generation probability is above a (percolation) threshold, multi-path routing enables long distance entanglement-generation at a rate that depends only linearly on the transmissivity η of a single link in the network, whereas the rate achieved by a linear repeater chain connecting Alice and Bob along the shortest path would decay as $\eta^{n_{\text{SP}}}$ where n_{SP} is the length of the shortest path. Pirandola recently showed that even a linear repeater chain can attain a rate that is proportional to η, but that requires repeater nodes equipped with error-corrected quantum processors [49]. We achieve

[1] Pirandola recently showed [49], for an information-theoretic description of repeaters that are ideal fully-error-corrected universal quantum processors, that the optimal rate attainable for multi-path entanglement routing using such ideal repeaters is superior to the rate of a linear chain of ideal repeaters.

the same feature (rate proportionality to η) by multipath routing with percolation, with a much simpler repeater. We also analyze repeater protocols to support multiple entanglement-generation flows. This analysis reveals that simple space-time-division multiplexing strategies that use local link-state knowledge at nodes can outperform the best rate-region (the set of rates simultaneously achievable by different flows) attainable by repeater nodes that simply time share among assisting individual flows.

Our work also opens up a number of new questions that remain unanswered. We abstract off the entanglement routing problem to the following parameters: G (network topology), p (probability of successful creation of an entangled pair across one link in a given time step), q (probability of a successful Bell measurement when attempted) and S (number of parallel links across a network edge). The protocols we develop here only require the quantum memory to hold a qubit for $T = 1$ time step.

Even in this simplified model, finding the rate-region optimizing routing protocol remains open. The aforesaid abstraction applies when the only source of imperfection at each component (including the quantum memories) is pure loss. Since our protocol only requires a quantum memory to hold a qubit for one entanglement attempt between neighboring stations ($T = 1$), photon loss would indeed be the major source of imperfection in many implementations of the protocol. Accounting for more general errors would require purification of entanglement [150, 151, 48], which will require us to introduce the fidelity of entanglement during intermediate steps of the routing protocol as an additional parameter, as was done by Jiang *et al.* [152]. Furthermore, we restricted our analysis only to 2-qubit measurements at repeater nodes. Multi-qubit unitary operations and multi-qubit measurements at repeater nodes (e.g., a 3-qubit GHZ projection across three locally held qubits) would require more complex repeaters than those in repeater chains, but may improve the achievable rates. Finally, it will be interesting to consider repeater protocols for distillation of multi-partite entanglement shared between more than two parties, and a repeater network that supports multiple simultaneous flows of multi-partite entanglement generation.

5.1 Background

Let us consider a quantum network with topology described by a graph $G(V, E)$. Each of the $N = |V|$ nodes is equipped with a quantum repeater, and each of the $M = |E|$ edges is a lossy optical channel of range L_i (km) and power transmissivity $\eta_i \propto e^{-\alpha L_i}$, $i \in E$. Consider K source-destination (Alice-Bob) pairs (A_j, B_j), $1 \leq j \leq K$, situated at (not necessarily distinct) nodes in V, each of which would like to generate maximally-entangled qubits (i.e., ebits) between themselves (and thus by definition not entangled with any other party, due to the monogamy property of entanglement), at the maximum rates possible R_j (ebits per channel use). The high-level objective is: *Given a class of quantum and classical operations at each of the repeater nodes of the underlying network, what operations should be performed at the repeater stations to maximize the rate region $(R_1, R_2, \ldots, R_K)$ simultaneously achievable by the entanglement flows?* More importantly, one would like to address networking questions such as: (a) what is the maximum rate-region attainable, (b) what is the tradeoff between sum throughput and latency of the K entanglement flows, and (c) where should repeater nodes be placed, with constraints on devices (e.g., memories, sources, and detectors), to maximize the attainable rate region; all being subject to various practical considerations. Ultimately one would like to develop explicit and efficient practical quantum routing protocols that employ quantum operations implemented via lossy and noisy devices, while only requiring local link-state knowledge and limited knowledge of the global network topology, analogous to the classical internet.

The entanglement-generation rate across a link of transmissivity η, in the absence of any repeater mediation, is limited to $-\log_2(1-\eta)$ ebits per mode, amounting to $\approx 1.44\eta$ ebits per mode when $\eta \ll 1$ [139][2]. The number of modes per second is a

[2]The achievability of $-\log_2(1-\eta)$ ebits per mode of secret communication rate over the lossy channel (with two way authenticated public classical communication) was first proven in 2009 by Pirandola *et al.* [47]. In 2014, Takeoka *et al.* proved an upper bound to the secret-key agreement capacity, $\log_2[(1+\eta)/(1-\eta)]$ ebits per mode [46], which equals $\approx 2.88\eta$ ebits per mode when $\eta \ll 1$, thereby establishing that the rate attained by *any* protocol must decay linearly with the channel's transmissivity and hence exponentially with distance L in optical fiber (since $\eta \sim e^{\alpha L}$). In 2015, Pirandola *et al.* proved an improved (weak converse) upper bound of $-\log_2(1-\eta)$ ebits per mode, which established that as the secret key agreement capacity of the pure loss bosonic channel [139]. In 2017, Wilde *et al.* proved $-\log_2(1-\eta)$ ebits per mode as a strong converse upper bound to the

device-dependent constant, upper bounded by the maximum of the optical bandwidth of the source and the electrical bandwidth of the detector. Since $\eta \sim e^{-\alpha L}$ where L is the length of optical fiber, the ebits-per-mode rate also falls off exponentially with range L. Most analyses of repeater networks have been limited to linear chains, with the objective of outperforming the repeater-less bound [48, 152, 154, 54, 55, 51, 80, 52]. Pirandola analyzed entanglement-generation capacities of repeater networks assuming ideal repeater nodes, i.e., those equipped with fully-error-corrected quantum processors and argued that for a single flow ($K = 1$), the maximum entanglement-generation rate R_1 reduces to the classical max-flow min-cut problem with edge e being associated with capacity $C(e) = -\log_2(1 - \eta(e))$ ebits per channel use [49], where $\eta(e)$ is the transmissivity of edge e. Pirandola subsequently argued that classical cut-set bounds with the above link capacity give outer bounds to the K-flow capacity region, but again, for ideal repeater nodes. Schoute and co-authors [155] developed routing protocols on specific network topologies and found scaling laws as functions of N, the number of qubits in the memories at nodes, and the time and space consumed by the routing algorithms, under the assumption that each link generates a perfect, lossless EPR pair in every time slot, and that the nodes' actions are limited to (perfect) Bell-state measurements (BSMs). Acín and co-authors [156] have considered the problem of entanglement percolation where neighboring nodes share a perfect, lossless pure state. Further, van Meter and co-authors developed explicit networking protocols also restricted to pair-wise EPR pair generation and BSMs, but accounting for imperfect fidelities of the EPR pairs (and thus requiring purification over multiple imperfect pairs), and finite coherence times of the qubit memories [157].

secret-key agreement capacity [153].

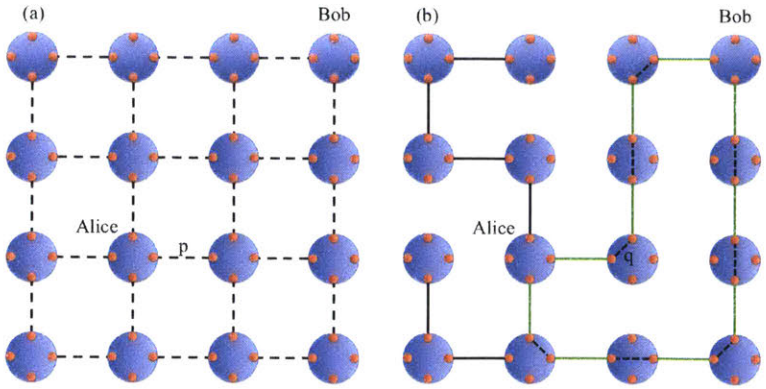

Figure 5-3: Schematic of a square-grid topology. The blue circles represent repeater stations and the red circles represent quantum memories. Every cycle (time slot) of the protocol consists of two phases. (a) In the first (external) phase, entanglement is attempted between neighboring repeaters along all edges, each of which succeed with probability p (dashed lines). (b) In the second (internal) phase, entanglement swaps are attempted within each repeater node based on the successes and failures of the neighboring links in the first phase—with the objective of creating an unbroken end-to-end connection between Alice and Bob. Each of these internal connections succeed with probability q. Memories can hold qubits for one time slot.

5.2 Entanglement routing protocols

5.2.1 Problem statement and notation

Consider a graph $G(V, E)$ that denotes the topology of the repeater network. See Fig. 5-2 for an illustration. Each node $v \in V$ is a repeater (blue circles), and each edge $e \in E$ is a physical link connecting two repeater nodes. $S(e) \in \mathbb{Z}^+$ is an integer edge weight, which corresponds to the number of parallel (single spatial, spectral or polarization mode) channels across the edge e (shown using blue lines). The number of memories at node v is $\sum_{e \in \mathcal{N}(v)} S(e)$ (see Fig. 5-2), where the sum is over $\mathcal{N}(v)$, the set of nearest neighbor edges of v, with $d(v) = |\mathcal{N}(v)|$ the degree of node v.

Time is slotted. We assume that each memory can hold a qubit perfectly for $T \geq 1$ time slots (T should be much smaller than the memory's coherence time). Each time slot t, $t = 1, 2, \ldots$, is divided into two phases: the "external" phase and the "internal" phase, which occur in that order. During the external phase, each of the $S(e)$ pairs of memories across an edge e attempts to establish a shared entangled (EPR) pair. An

entanglement attempt across any one of the $S(e)$ parallel links across edge e succeeds with probability $p_0(e) \sim \eta(e)$ [46, 139], where $\eta(e) \sim e^{-\alpha L(e)}$ is the transmissivity of a lossy optical channel of length $L(e)$. Using two-way classical communication over edge $e(u, v)$, neighboring repeater nodes u, v learn which of the $S(e)$ parallel links (if any) succeeded in the external phase.

Let us assume that neighboring repeaters pick upto one successfully-created ebit (i.e., ignore multiple successes if any) as in Ref. [55, 80], in which case the probability that one ebit is established successfully across the edge e during the external phase is given by: $p(e) = 1 - (1 - p_0)^{S(e)}$. Let us also assume $S(e) = S, \forall e \in E$, which in turn gives us $p(e) = p, \forall e \in E$. While the results in this chapter can be adapted to any network topology, we will henceforth use the 2D regular square grid topology (Fig. 5-3) to illustrate the performance of our routing algorithms.

One instance of the resulting external links created between repeater nodes after the external phase is shown in Fig. 5-3(b) using solid lines. In the internal phase, entanglement swap (BSM) operations are attempted locally at each repeater node between pairs of qubit memories. We associate these BSM attempts as *internal links*, i.e. links between memories internal to a repeater node, shown using dotted lines inside repeaters in Fig. 5-3(b). If $T > 1$, a repeater node can attempt a BSM between qubits held in two memories that were entangled with their respective neighboring node's qubits in two different time slots. For minimizing the demands on memory coherence time [55, 80], we will assume $T = 1$. So, BSMs will always be attempted between two qubits in distinct memories that were entangled with their respective counterparts at their respective neighboring nodes in the same time slot. Each of these internal-link attempts succeed with probability q. Therefore, after the conclusion of one time slot, along a path comprising k edges (and thus $k - 1$ repeater nodes, excluding Alice and Bob), one ebit is successfully shared between the end points of the path with probability $p^k q^{k-1}$. The maximum number of ebits that can be shared between Alice (say, node a) and Bob (say, node b) after one time slot is $\min\{d(a), d(b)\}$, assuming S is the same over all edges. For the square-grid topology shown, the maximum number of ebits that can be generated between Alice and Bob

in each time slot is 4.

The remainder of the chapter is dedicated to finding the optimal strategy for each repeater node in order to decide which locally-held qubits to attempt BSM(s) on during the internal phase of a time slot, based ideally only on knowledge of the outcomes (success or failure) of the nearest neighbor links, i.e., local link-state knowledge, during the respective preceding external phases. We will assume that each repeater node is aware of the overall network topology as well as the locations of the K Alice-Bob pairs. The goal of the optimal repeater strategy will be to attain the maximum entanglement-generation rate (if there is a single Alice and Bob, i.e., $K = 1$) or the maximum rate-region for multiple flows (i.e., $K > 1$).

5.2.2 Multipath routing of a single entanglement flow

Entanglement routing with global link-state information

We begin with the assumption that global link-state knowledge is available at each repeater node, i.e., the state of every external link in the network after the external phase is known to every repeater in the network and can be used to determine the choice of which internal links to attempt within the nodes. Each memory can only be part of one entanglement swap, i.e., each red node can only be part of one internal edge. Consider the following greedy algorithm to choose the internal links: consider the subgraph induced by the successful external links and the repeater nodes (at the end of the external phase), and find in it the shortest path connecting Alice and Bob. If no connected path between Alice and Bob exists, no shared ebits are generated in that time slot. If a shortest path of length k_1 is found, all internal links along the nodes of that path are attempted, and the probability a shared ebit is generated by this path is the probability that all $k_1 - 1$ internal link attempts were successful, i.e., q^{k_1-1}. We then remove all the (external and internal) links of the above path from the subgraph, and find a shortest path connecting Alice and Bob in the pruned subgraph. Note that instead of removing the links of the first path from the subgraph, we could simply search for a shortest path in the original subgraph but one that is edge disjoint

from the previous path. If such a path exists, we again attempt all internal links at the nodes of this path, so the probability the path contributes to the generation of an ebit between Alice and Bob is q^{k_2-1} where k_2 is the length of the second path; and so on.

The entanglement generation rate achieved using this greedy algorithm R_g is the sum of expected rates (in ebits per time slot) from these paths. Given the degree-4 nodes in a square grid topology, there can be a maximum of four edge disjoint paths between Alice and Bob. Fig. 5-3(b) illustrates our greedy algorithm. Given the set of external links created, the shortest path has length $k_1 = 4$, the next path has length $k_2 = 6$, and no further paths can be found. The two edge-disjoint paths are highlighted in green. Hence, the internal links depicted with the dashed lines in Fig. 5-3(b) are attempted and the expected number of shared ebits generated in this time cycle is: $q^{k_1-1} + q^{k_2-1}$. The net entanglement generation rate is the expectation of sums like the above (with up to four terms) over many random instantiations of the $(p, 1-p)$ external-link creations during the external phase of many time slots. Evaluating this expected rate $R_{\mathrm{g}}(p, q)$ achieved by the above routing strategy analytically as a function of the Alice-Bob separation (X, Y) is difficult, even for a square-grid topology.

The intuition behind this simple greedy algorithm is that the entanglement generation rate along a path of length k decays exponentially as q^{k-1}, suggesting that attempting internal links to facilitate connections along the shortest path first would optimize the expected rate. However, it is possible to draw random instances of successes of external links, where either one of the two possible options—(1) picking the shortest path (which disrupts all other paths) and (2) picking two edge disjoint (but longer) paths—could yield either a larger or a smaller expected rate than the other, depending upon the value of q. If q is larger than a threshold, option (2) would have a larger expected rate and vice versa. Finding the global optimal rule remains an open problem. It is easy however to prove that the greedy algorithm achieves a rate within a factor of four of the optimum algorithm employing global link-state knowledge, $R_{\mathrm{opt}}(p, q)$. Let us denote the length of the shortest path between Alice and Bob in the induced subgraph after the external phase, as $n_{\mathrm{SP}}(p)$. This quantity is of interest

in percolation theory, and is not completely understood analytically. It undergoes a sharp transition (i.e., starts out large and suddenly jumps to a much smaller value) as p crosses p_c from below to above. Clearly, $R_{\mathrm{g}}(p,q) \geq \mathbb{E}[q^{n_{\mathrm{SP}}(p)-1}]$ since using the shortest path is the first step of the greedy algorithm. Furthermore, since the optimal rule can create entanglement over a maximum of four edge-disjoint paths in each time step, each of which must have a length no less than the length of the shortest path, $R_{\mathrm{opt}}(p,q) \leq \mathbb{E}[4q^{n_{\mathrm{SP}}(p)-1}] \triangleq R_{\mathrm{opt}}^{(\mathrm{UB})}(p,q)$. Therefore, $R_{\mathrm{opt}} \geq R_{\mathrm{g}} \geq R_{\mathrm{opt}}/4$, i.e., the greedy rule will achieve the same rate-vs.-distance scaling as the optimal algorithm that employs global link-state knowledge, and can be worse only by a small constant factor.

In Fig. 5-4(a) we plot $R_{\mathrm{g}}(p,q)$ as a function of the Alice-Bob Manhattan distance (X,Y) on the square grid (measured in number of hops) with $q = 1$. When $p > p_c$, the bond percolation threshold of the underlying network ($p_c = 0.5$ for the square lattice), a *giant connected component* is formed by the external links alone at the end of the first (external) phase of a time slot. Recall that the rate along a length k path is $p^k q^{k-1}$, where $p \sim \eta$ is the transmissivity of each link. In the network case, when $p > p_c$ and $q = 1$, the p^k portion of the rate expression becomes immaterial for scaling with Alice-Bob distance, since percolation guarantees a connected path to exist between Alice and Bob along successful external links in each time slot. So, if $q = 1$, $R_{\mathrm{g}}(p,q)$ remains essentially distance invariant. When $p < p_c$, the rate falls off exponentially with distance (even when $q = 1$). It is instructive to note here that the optimal rate (entanglement-generation capacity) achievable on a single length k path does not depend on k, and only on the transmissivity of the lossiest link in the path, i.e., $C \sim \eta$ [49], but achieving this requires infinite-coherence-time quantum memories and ideal quantum operations at nodes. The multi-path gain in the $p > p_c$ regime lets us achieve a distance-independent rate, but with memories whose coherence times are no more than one time slot. The rates are calculated using monte-carlo simulations which results in some numerical noise that is insignificant compared to the difference between the plots, but is visible in $R_{\mathrm{g}}(0.45, 1)$.

A general upper bound on the entanglement generation rate is given by the min-

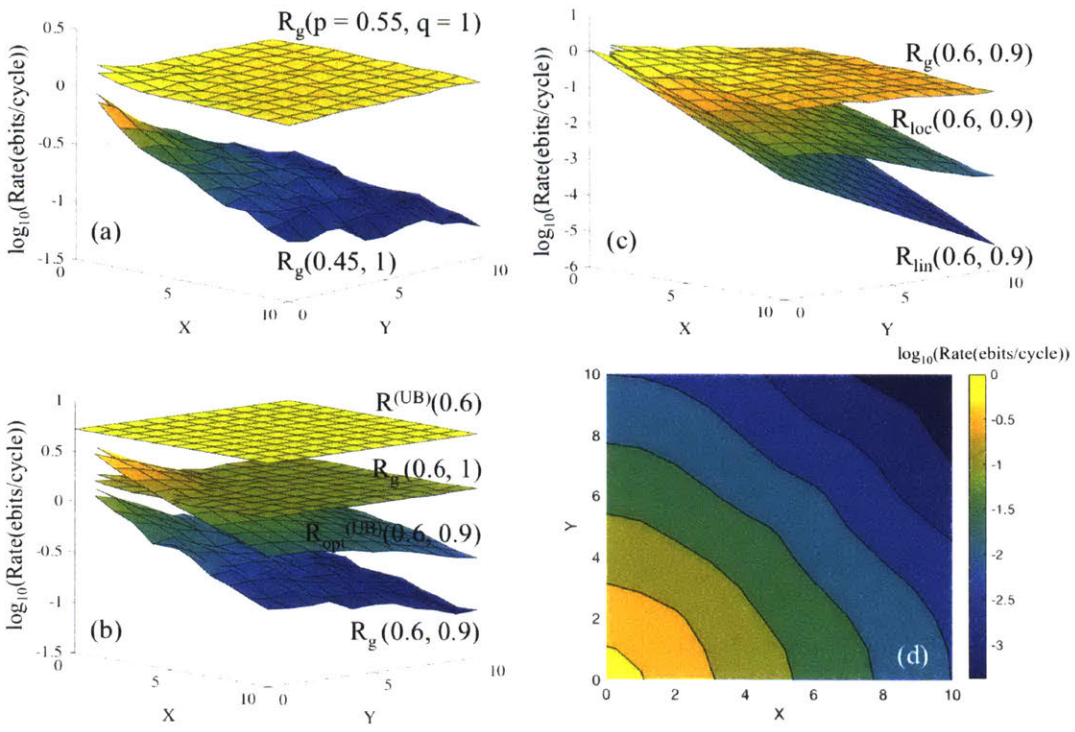

Figure 5-4: Entanglement generation rate as a function of the Alice-Bob separation along X and Y (on a square grid) as a function of (p, q); (a) $R_{\rm g}(p, q)$ is the rate attained by a global-knowledge-based protocol we propose where each node, in each time step, knows whether any link in the entire network succeeded or failed to establish entanglement. For the case of $q = 1$, $R_{\rm g}$ is distance independent when p is greater than the bond percolation threshold (0.5 for the square lattice) and decays exponentially if it is below the threshold. (b) $R^{(\rm UB)}(0.6)$ is the distance-independent Pirandola rate upper bound for $p = 0.6$, achieving which requires perfect quantum processing at repeater nodes. $R_{\rm g}(0.6, 1)$ is also distance independent, and within a factor 3.6 of $R^{(\rm UB)}(0.6)$. With $q < 1$, e.g., $R_g(0.6, 0.9)$, the rate decays exponentially with distance. $R^{(\rm UB)}_{\rm opt}$ is an upper bound on the rate attainable with global-knowledge by any protocol. (c) $R_{\rm loc}$ is attained by a protocol we propose where each node, in each time step, only needs to know the link state of neighboring edges. The rate-distance scaling exponent of $R_{\rm loc}$ is clearly worse than $R_{\rm g}$, but is significantly superior to that of a linear repeater chain along the shortest path, $R_{\rm lin}$, demonstrating multi-path routing advantage even with local link-state knowledge. (d) Contour plot of the entanglement generation rate with the local rule when $p = 0.6$ and $q = 0.9$. Although the Alice to Bob distance along the network links is $X + Y$, there is a noticeable enhancement in the rate along the $X = Y$ direction because of more Alice-Bob paths of similar length.

cut of the graph [49], and for a square lattice, is given by $R^{(\mathrm{UB})}(p) = -\log_2[(1-p)^4]$. $R^{(\mathrm{UB})}(0.6)$ is plotted in Fig. 5-4(b). The known methods for achieving $R^{(\mathrm{UB})}$ require infinite coherence time memories and error-corrected quantum processors at each node. For our implementation (assuming global link state knowledge), $R_{\mathrm{g}}(0.6, 1)$ is also plotted in Fig. 5-4(b). Although our protocol only requires memories to hold entanglement for one time step, the multi path advantage gives us the same constant rate-distance scaling and within a factor of ~ 3.6 of $R^{(\mathrm{UB})}(0.6)$. The assumption of perfect BSMs is unrealistic and $q < 1$, in which case $R_g(p, q)$ falls off exponentially with distance; even when $p > p_c$, as seen in the plot for $R_{\mathrm{g}}(0.6, 0.9)$. Finally, we plot the above discussed upper bound on R_{g}, $R^{(\mathrm{UB})}_{\mathrm{opt}}(0.6, 0.9)$, which as expected has the rate-distance scaling of R_{g}, but larger by a factor less than 4.

Entanglement routing with local link-state information

$R_{\mathrm{g}}(p, q)$, the rate attained by the protocol described in the previous subsection that employs global link-state knowledge, is re-plotted in Fig. 5-4(c). We also plot $R_{\mathrm{lin}}(p, q) = p^{n_{\mathrm{SP}}(1)} q^{n_{\mathrm{SP}}(1)-1}$, the rate attained by a single linear repeater chain, where $n_{\mathrm{SP}}(1)$ is the shortest-path length between Alice and Bob along the edges of the underlying square grid. The assumption of global link-state knowledge in large networks is unrealistic, as it requires memories whose coherence time increases with the network size due to the time required for the traversal of link-state information across the entire network. In this section, we describe a more realistic protocol in which knowledge of success and failure of an external link at each time slot is communicated only to the two repeater nodes connected by the link, as is the case in the analysis of many 'second-generation' linear repeater chains [154, 55, 152]. Repeater nodes need to decide on which pair(s) of memories BSMs should be attempted (i.e., which internal links to attempt), based only on information about the states of external links adjacent to them. We assume that network topology and positions of Alice and Bob are known to each repeater station, and communicated classically beforehand.

Let us consider a local repeater rule illustrated in Fig. 5-5. The repeater u inside the dotted square box has to make a decision regarding which internal edges to

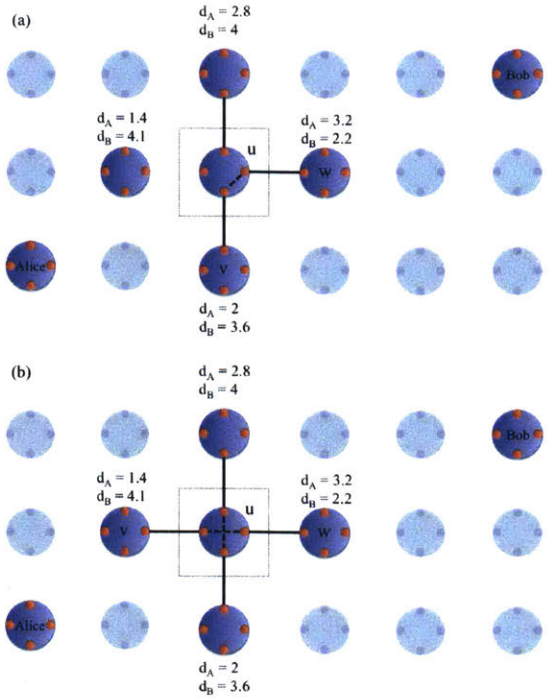

Figure 5-5: The entanglement swap rule used at the repeater u in the dotted box in the case of local link-state knowledge. v and w are the repeaters closest to Alice and Bob, respectively, with a direct edge to u. (a) If two or three links are up, the memories linked to v and w undergo an entanglement swap. (b) If four links are up, the remaining two memories also undergo an entanglement swap.

attempt based on the information of which of the four neighboring external edges have been successfully created in the external phase. We associate d_A and d_B as the distance to Alice and Bob, respectively, at every repeater node. We use the $\mathbb{L}^2$ norm distance to Alice (resp. Bob) as d_A (resp. d_B), which appears to be close to the optimal distance metric (discussed later). Of all the nearest neighbor nodes of u whose links to u were successful in that time slot, we label the one that has the minimum d_A as v. Similarly, the neighbor with a successful external link with u and the minimum d_B is labelled w. An internal link is attempted between the memories connected to v and w respectively, as shown in Fig. 5-5(a). If v and w are the same node, v (or w) is replaced by node u's nearest-neighbor node with the next smallest value of d_A (or d_B). The choice of whether to replace v or w is made in a manner that minimizes the sum of d_A and d_B from the eventually chosen two neighbors to connect. If all

four external links are successful, an additional internal link is attempted between the remaining two memories as shown in Fig. 5-5(b). If only one of the neighboring external links is successful, no internal links are attempted, since this repeater node cannot be part of a path from Alice to Bob in that time slot. If two neighbors have the same values of d_A and d_B, an unbiased coin is tossed to determine the choice of v and w, to preserve symmetry in the protocol.

The entanglement generation rate $R_{\text{loc}}(p, q)$ achieved by the above described local rule is plotted in Fig. 5-4(c) and compared to $R_{\text{g}}(p, q)$ and $R_{\text{lin}}(p, q)$. We use $p = 0.6$ and $q = 0.9$, the same values used for the global-information rate plots in Fig. 5-4(b). As one expects, the rate-distance scaling of R_{loc} is worse than that of R_{g}. However, the rate-distance scaling exponent achieved by the local rule is superior to that of a linear chain, even though the physical elements employed to build the repeaters are identical. This is proven analytically in section 5.2.3. Note that each of the three rates R_{g}, R_{loc}, R_{lin} fall exponentially with distance, but the exponents are different. The scaling advantage of R_{loc} over R_{lin} arises because the local rule allows the entanglement-generation flow between Alice and Bob to find different (and potentially simultaneously multiple) paths in different time slots, and does not have to rely on all links along a linear chain to be successful. This is analogous to multi-path routing in a classical computer network. The contour plot in Fig. 5-4(d) further illustrates this point: there is a noticeable enhancement of R_{loc} along the $X = Y$ line because the diagonal direction contains the largest spatial density of possible paths between Alice and Bob. The scaling advantage over R_{lin} persists in any direction, i.e., along $Y = 0$ as well.

Sweeping over different values of p and q, we find that the multi-path advantage relative to a linear repeater chain increases as p decreases from unity, but there is little relative improvement as q is varied (see section 5.2.3).

Clearly, other distance metrics (e.g., $\mathbb{L}^p$ norm for $p \geq 1$) can be used in lieu of the $\mathbb{L}^2$ norm in the algorithm described above. In section 5.2.4, we present a recursive numerical evaluation technique to find the rate-optimal distance metric, which can be applied to any network topology. For planar network topologies, the $\mathbb{L}^2$ norm appears

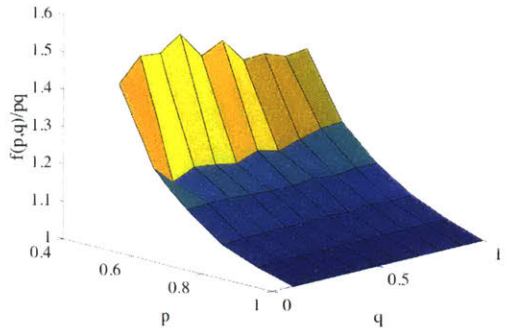

Figure 5-6: $f(p,q)/pq$ quantifies the improvement in the scaling of $R_{\text{loc}}(p,q)$ with respect to $R_{\text{lin}}(p,q)$ with respect to the Alice-Bob Manhattan distance, n. $f(p,q)/pq$ increases as p is reduced in $[1, p_c]$ but changing q has a negligible effect.

near-optimal for our local routing algorithm.

An analytical enumeration of the expected number of edge-disjoint paths as a function of p between Alice and Bob separated by a given distance (X, Y) in a bond-percolation instance (i.e., with $p > p_c$) of a network is an open question, the solution of which will enable a firmer quantitative understanding of the multi-path advantage in entanglement generation in a repeater network.

5.2.3 Multipath rate advantage

Numerical Evaluation

The goal this subsection is to quantify the improvement in the rate-vs.-distance exponent achieved by our local rule over that of a linear chain along the shortest path, for all possible pairs of values of p and q. Fig. 5-4(c) shows this improvement, i.e., that of $R_{\text{loc}}(p,q)$ compared to $R_{\text{lin}}(p,q)$, for $p = 0.6$ and $q = 0.9$. Clearly, $R_{\text{lin}}(p,q) = (pq)^n(p)/q \sim (1/q)[pq]^n$, where n is the Manhattan distance between Alice and Bob. We have numerically verified that $R_{\text{loc}}(p,q) \sim g(p,q)[f(p,q)]^n$ for n large. We hence quantify the rate improvement by numerically evaluating the ratio

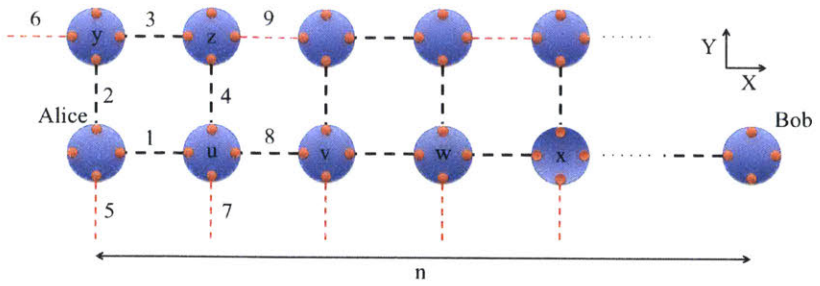

Figure 5-7: Network used to prove the lower bound on entanglement generation rate with our local routing rule which shows that scaling of the rate with Alice-Bob Manhattan ($\mathbb{L}^1$) distance for our rule is better than the scaling of the rate along a linear repeater chain along the shortest path between Alice-Bob.

$f(p,q)/(pq)$ exhaustively for all $(p,q) \in [0,1] \times [0,1]$, using Monte Carlo simulations. The results are shown in Fig. 5-6, for configurations of Alice and Bob located along 45° with respect to the grid axes. We see that $f(p,q)/pq$ increases as p decreases in $[p_c, 1]$, but changing q has a negligible effect on this ratio.

Analytical lower bound on the rate achieved by the local routing rule

In this subsection, we derive an analytical lower bound on the entanglement generation rate attained by our local routing rule (using the $\mathbb{L}^2$ norm as the distance metric), with the objective of demonstrating multi-path routing advantage, i.e., the rate-vs.-distance scaling attained by our local rule is strictly better than that attained by a linear repeater chain along the shortest path between Alice and Bob.

Consider routing entanglement between Alice and Bob located at (X, Y) and $(X+n, Y)$ respectively, i.e., n hops apart along the X dimension of the square lattice. We will evaluate a lower bound on R_{loc} by only evaluating the rate contributions from paths in which all the (external) links belong to the set of black dashed links shown in Fig. 5-7. The choice of internal links made at repeater nodes proceed as usual per our local rule. As a result, there are instances in which our local rule routes entanglement through paths comprising not just the black links, resulting in flows that do not contribute to our rate lower bound.

We will refer to Fig. 5-7 for the ensuing discussion. Recall that external links

succeed (are 'up') with probability p and fail (are 'down') with probability $1-p$, whereas internal links succeed with probability q. Consider $P(A \leftrightarrow v)$, the probability that there is a path between Alice A and repeater v that uses only black links. $P(A \leftrightarrow v)$ includes the probability of making the required internal links to create a path between A and v, but not the probability of any internal links at the end points A or v. It is easy to see that in any given time step, there can be no more than one edge-disjoint path between A and v along the black dashed links, since link 8 must be part of the path. Let l $\left(\widetilde{l}\right)$ be the event that the external link l is up (down). Further, note that at any given time step, of all the possible (0, 1 or 2) internal links attempted by our local rule at a repeater node, only one internal link, if successful, contributes to $A \leftrightarrow v$. Let $l-m$ be the event that the internal link attempted at a repeater node to connect external links l and m is successful. If links 1 and 8 are both up, node u attempts to connect those two links based on our local rule, regardless of the other links. If links 2, 3, 4 and 8 are up, but 1, 5, 6, 7 and 9 are down, u attempts to connect 4 and 8, z attempts to connect 3 and 4 and y attempts to connect 2 and 3. Considering these two possibilities, we have

$$\begin{aligned} P(A \leftrightarrow v) &> \Pr(1, 8, \text{1-8}) \\ &\quad + \Pr(2, 3, 4, 8, \widetilde{1}, \widetilde{5}, \widetilde{6}, \widetilde{7}, \widetilde{9}, \text{2-3}, \text{3-4}, \text{4-8}) \\ &= \left[p + p^3(1-p)^5 q^2\right] pq \\ &= p'pq, \end{aligned} \tag{5.1}$$

where $p' = p + p^3(1-p)^5 q^2 > p$.

$P(v \leftrightarrow x)$ is the probability that there is a path between v and x that uses only black links (the probability of internal link successes at the end points v and x are not included). $P(v \leftrightarrow x)$ and $P(A \leftrightarrow v)$ are not independent events because they both involve link 9. $P(v \leftrightarrow x | A \leftrightarrow v)$ is the probability that there exists a path along black dashed lines between v and X given that a path along black dashed lines exists between A and v. We now show that $P(v \leftrightarrow x | A \leftrightarrow v) > P(v \leftrightarrow x)$.

$$
\begin{aligned}
P(v \leftrightarrow x|A \leftrightarrow v) &= P(v \leftrightarrow x|A \leftrightarrow v, 9)\Pr(9|A \leftrightarrow v) + \\
&\quad P\left(v \leftrightarrow x|A \leftrightarrow v, \widetilde{9}\right) \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \\
&= P(v \leftrightarrow x|9)\Pr(9|A \leftrightarrow v) + \\
&\quad P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \\
&= P(v \leftrightarrow x|9)\left(1 - \Pr\left(\widetilde{9}|A \leftrightarrow v\right)\right) + \\
&\quad P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \\
&= \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \times \\
&\quad \left(P\left(v \leftrightarrow x|\widetilde{9}\right) - P(v \leftrightarrow x|9)\right) \\
&\quad + P(v \leftrightarrow x|9) \qquad (5.2)
\end{aligned}
$$

where $P(v \leftrightarrow x|A \leftrightarrow v, 9) = P(v \leftrightarrow x|9)$ and $P\left(v \leftrightarrow x|A \leftrightarrow v, \widetilde{9}\right) = P\left(v \leftrightarrow x|, \widetilde{9}\right)$ because link 9 being up or down is the only probabilistic event that influences both $P(A \leftrightarrow v)$ and $P(v \leftrightarrow x)$. Further,

$$
\begin{aligned}
P(v \leftrightarrow x) &= P(v \leftrightarrow x|9)\Pr(9) + P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}\right) \\
&= P(v \leftrightarrow x|9)\left(1 - \Pr\left(\widetilde{9}\right)\right) + \\
&\quad P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}\right) \\
&= \Pr\left(\widetilde{9}\right)\left(P\left(v \leftrightarrow x|\widetilde{9}\right) - P(v \leftrightarrow x|9)\right) \\
&\quad + P(v \leftrightarrow x|9). \qquad (5.3)
\end{aligned}
$$

Comparing 5.2 and 5.3, $\Pr\left(\widetilde{9}|A \leftrightarrow v\right) = \Pr\left(\widetilde{9}\right) \Pr\left(A \leftrightarrow v|\widetilde{9}\right) / \Pr\left(A \leftrightarrow v\right) >$ $\Pr\left(\widetilde{9}\right)$ because $\Pr\left(A \leftrightarrow v|\widetilde{9}\right) > \Pr\left(A \leftrightarrow v\right)$ following equation 5.1. Similarly, $\left(P\left(v \leftrightarrow x|\widetilde{9}\right) - P(v \leftrightarrow :\right.$ 0. Hence, $P(v \leftrightarrow x|A \leftrightarrow v) > P(v \leftrightarrow x)$.

From Fig. 5-7, we can see that in order to get a path along black dashed lines from A to x, there must be a path along black dashed lines from A to v and from v to x, and the internal link at v must succeed. Therefore,

$$
\begin{aligned}
P(A \leftrightarrow x) &= P(A \leftrightarrow v) q P(v \leftrightarrow x | A \leftrightarrow v) \\
&> P(A \leftrightarrow v) P(v \leftrightarrow x) q \\
&= \left(P(A \leftrightarrow v)\right)^2 q \\
&= (p'pq)^2 q, \qquad (5.4)
\end{aligned}
$$

where we use symmetry between $A \leftrightarrow v$ and $v \leftrightarrow x$ in the third line. Repeating this for all repeaters between Alice and Bob, it is easy to see that

$$
\begin{aligned}
R_{\text{loc}} > P(A \leftrightarrow B) &> p'^{\lceil n/2 \rceil} p^{\lfloor n/2 \rfloor} q^{n-1} \qquad (5.5) \\
&\geq \left(\sqrt{p'p}\right)^n q^{n-1} \\
&= \left[\left(\sqrt{p'p}\right) q\right]^n q^{-1} \\
&= (pq)^{\beta n} q^{-1},
\end{aligned}
$$

where $\lceil n/2 \rceil$ is the smallest integer greater than or equal to $n/2$ and $\lfloor n/2 \rfloor$ is the largest integer smaller than or equal to $n/2$. The second inequality uses the fact that $p' > p$ and $n > 0$. $\beta = \log\left[\left(\sqrt{p'p}\right) q\right] / \log\left[pq\right] < 1$ because $p < p' < 1$ and $q < 1$.

Therefore, since $R_{\text{loc}} > (pq)^{\beta n} q^{-1}$ with $\beta < 1$ and $R_{\text{lin}} = (pq)^n q^{-1}$, the exponent in the scaling with n is smaller in R_{loc} compared to R_{lin}, i.e. the rate-vs.-distance scaling is better with multi-path routing. Using a similar reasoning, it is easy to see that the same is true even when Alice and Bob are at located at different Y coordinates. It should be noted that the lower bound we derive here is not meant to be tight (see Section 5.2.3 for a full numerical evaluation of the exponents for R_{loc} and R_{lin}). The only purpose of this subsection was to prove that the rate-vs.-distance scaling for entanglement routing strictly benefits from multi-path routing.

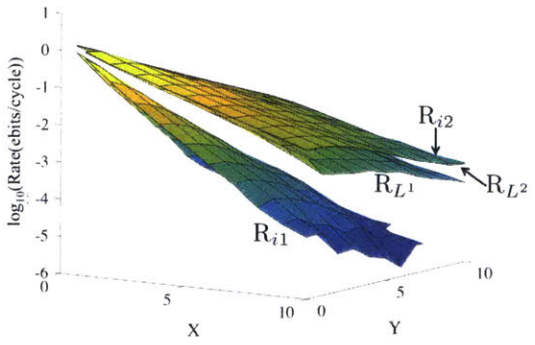

Figure 5-8: Entanglement generation rates with different distance metrics. $R_{\mathbb{L}^1}$ and $R_{\mathbb{L}^2}$ are evaluated using the $\mathbb{L}^1$ and $\mathbb{L}^2$ norms respectively. The distance metric for R_{i_1} (iteration 1) is calculated using $R_{\mathbb{L}^1}$, and R_{i_2} (iteration 2) is calculated using R_{i_1}. R_{i_2} and $R_{\mathbb{L}^1}$ are nearly indistinguishable as they almost coincide.

5.2.4 Distance metric for the local routing rule using $\mathbb{L}^1$ norm and recursion

Our entanglement routing protocol with local link-state information uses the 'distance' of neighboring repeater stations from Alice and Bob to decide which memories at a repeater should undergo entanglement swap attempts. The results presented in the chapter use the $\mathbb{L}^2$ norm as the distance metric. While the $\mathbb{L}^2$ norm can be easily calculated for the square grid, it may not be easily generalizable for other (e.g., non-planar) topologies. Further, even though we do not prove the rate optimality of our local link-state routing protocol, given a network topology, it is not clear whether or not the $\mathbb{L}^2$ norm is the optimal distance metric to be used in our protocol.

In order to adapt our algorithm for arbitrary network topologies, and also to find a near-optimal distance metric for our algorithm, we employ the following numerical recursive method. Our evaluation begins with calculating $R_{\mathbb{L}^1}(\mathbf{n}_1, \mathbf{n}_2)$, the entanglement generation rate achieved when our local rule is used to route entanglement between nodes $\mathbf{n}_1$ and $\mathbf{n}_2$, using the $\mathbb{L}^1$ norm as the distance metric. In Fig. 5-8, we plot $R_{\mathbb{L}^1}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y), where X and Y are the distance (in hops) between $\mathbf{n}_1$ and $\mathbf{n}_2$ along the horizontal and vertical dimensions of the square grid, respectively. The rate-distance scaling exponent for $R_{\mathbb{L}^1}$ is worse than that of $R_{\mathbb{L}^2}$,

the rate attained by our protocol, using the $\mathbb{L}^2$ norm as the distance metric. Next, for every repeater node $\mathbf{n}$, we define distances d_A and d_B to Alice $\mathbf{A}$ and Bob $\mathbf{B}$ respectively, with respect to the following new distance metric (let us name this metric $i1$): $d_A := 1/\mathrm{R}_{\mathbb{L}^1}(\mathbf{n}, \mathbf{A})$ and $d_B := 1/\mathrm{R}_{\mathbb{L}^1}(\mathbf{n}, \mathbf{B})$. We then calculate $\mathrm{R}_{i1}(\mathbf{n}_1, \mathbf{n}_2)$, the entanglement generation rate achieved when our local rule is used with the $i1$ distance metric to route entanglement between every pair of nodes $\mathbf{n}_1$ and $\mathbf{n}_2$. In Fig. 5-8, we plot $\mathrm{R}_{i1}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y). We see that the rate-distance scaling achieved by R_{i1} is even lower than that of $\mathrm{R}_{\mathbb{L}^1}$. However, when we go through the second iteration of the algorithm—i.e., define distance metric $i2$, under which $d_A = 1/\mathrm{R}_{i1}(\mathbf{n}, \mathbf{A})$ and $d_B = 1/\mathrm{R}_{i1}(\mathbf{n}, \mathbf{B})$, and use our local rule to evaluate $\mathrm{R}_{i2}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y)—we find that the resulting rate R_{i2} is almost the same (visually indistinguishable in the plot) as $\mathrm{R}_{\mathbb{L}^2}$, the rate we obtained directly when using the $\mathbb{L}^2$ norm as the distance metric. This suggests that: (a) for the square grid (and presumably for any planar network topology) the $\mathbb{L}^2$ norm metric might be near-optimal for use within our local rule, and that (b) for any given network topology, one could potentially pre-compute the optimal distance metric by a recursive strategy on the given topology using the $\mathbb{L}^1$ norm as the starting point. However, there are instances where our local rule does not give the rate-optimal local routing rule. As an example, when $p = 1$ and $q = 1$, it is possible to find four disjoint paths without any link-state knowledge (the links are all deterministic) and the optimal rate is four ebits/cycle for any location of Alice and Bob. However, the fact that we are trying to route every flow through the best possible path without any coordination between different flows leads to collisions, which results in a rate that is below the optimal rate of four ebits/cycle. Finding the rate-optimal local routing rule across different parameter values is left for future research.

5.2.5 Simultaneous entanglement flows

In this section, we consider simultaneous entanglement-generation flows between two Alice-Bob pairs, using local link state knowledge at all repeater nodes. Consider two pairs Alice 1 - Bob 1 (red nodes) and Alice 2 - Bob 2 (green nodes) as shown in

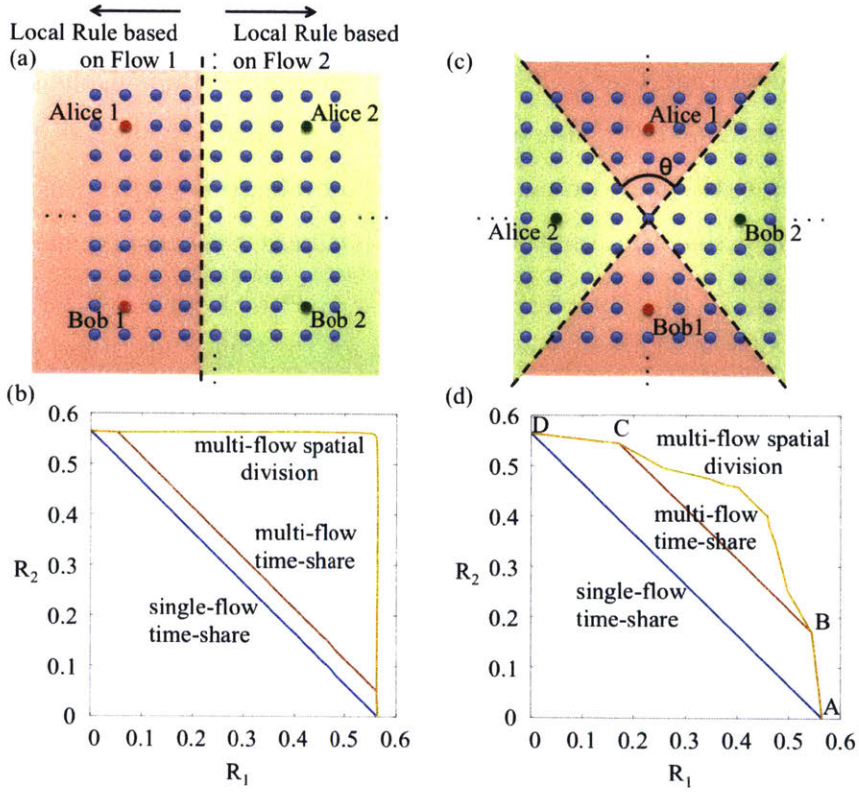

Figure 5-9: (a) Multi-flow routing for two Alice-Bob pairs that lie along the sides of a 6×6 square, embedded in a 100×100 grid; (b) rate region (R_1, R_2) with different rules at repeater nodes, each employing local link-state knowledge, for $p = q = 0.9$. (c) Multi flow routing when the Alice-Bob paths cross (d) multiflow rate region for two local-knowledge rules.

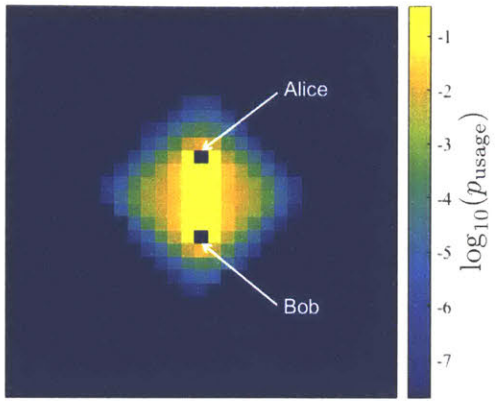

Figure 5-10: A heat map plotting p_{usage}, the probability that a given repeater node is involved in a successful creation of a shared ebit generated between Alice and Bob, separated by 6 hops in an underlying square grid topology, when our local rule is employed. We assume $p = 0.9$ and $q = 0.9$.

the two scenarios in Fig. 5-9. In Fig. 5-9(a), the shortest paths connecting the two Alice-Bob pairs do not cross, but they do in Fig. 5-9(b). In both cases, they are placed at the four corners of a 6×6 square grid, embedded within a large square grid network. Denote by R_1 and R_2 the entanglement generation rates achieved by the two Alice-Bob pairs respectively. We first consider the case of non-intersecting flows shown in Fig. 5-9(a). A simple strategy is for every single repeater node (including the nodes labeled as the two Alices and Bobs) to use the local rule described in the previous section tailored to support the Alice 1-Bob 1 flow for a fraction, λ, of the time slots and to support the Alice 2-Bob 2 flow for the remaining $1 - \lambda$ fraction. For $p = q = 0.9$, the rate region attained by varying $\lambda \in [0, 1]$ is depicted with the blue line in Fig. 5-9(b), which we refer to as single-flow time-share. However, if every repeater with the exception of the Alices and Bobs carry out the above time-sharing strategy, even when all repeater nodes support flow 1, there is still some 'left-over' non-zero R_2 that is attained. This multi-flow time-share rate region is shown using the red line in Fig. 5-9(a).

In Fig. 5-10, for the case that Alice and Bob are separated by 6 hops on the square grid, we plot a color map of p_{usage}, the probability a given repeater node is involved in a successful creation of a shared ebit generated between Alice and Bob when our

local rule is employed. We observe that only the repeaters lying in a small spatial region surrounding the straight line joining Alice and Bob are used significantly. This observation motivates a multi-flow spatial-division rule, in which we divide the network between the two flows, as shown in Fig. 5-9(a). Any repeater in the red shaded region follows the local rule tied to the Alice 1 - Bob 1 flow while repeaters in the green region operate with the local rule tied to the Alice 2 - Bob 2 flow. The placement of the boundary determines the rates R_1 and R_2. The rate region attained is plotted with the yellow line in Fig. 5-9(b). This significantly outperforms time sharing. The two flows can co-exist and operate with a very small reduction from their individual best rates, because the repeaters they respectively benefit from the most form almost disjoint sets.

In the other extreme, we consider two Alice-Bob pairs, still separated by six hops, but with their shortest paths crossing as shown in Fig. 5-9(c). The rate region attained by multi-flow time sharing, shown by the line segment BC, still provides an improvement over single-flow time-sharing, shown by the line segment AD, as shown in Fig. 5-9(d). It is interesting to note that the maximum R_1 under multi-flow time sharing (point B) is slightly lower than maximum R_1 with the single-flow time-share rule (point A). This happens because unlike in single flow time-share, the nodes at Alice 2 and Bob 2 do not contribute to R_1 under multi-flow time-share. A point along AB represents time sharing between the strategies at points A and B. To further increase the rate, we adopt a multi-flow spatial division strategy in which nodes in the red region are configured to assist flow 1 and nodes in the green region are configured to assist flow 2. Varying the angle θ demarcating those regions results in the rate region shown by the yellow line in Fig. 5-9(d). This time, the improvement due to the spatial-division rule is not as pronounced, since the spatial regions corresponding to 'useful' repeater nodes for the two flows are not disjoint.

5.3 Conclusions and open problems

We proposed and analyzed quantum repeater protocols for entanglement generation in a quantum network in an architecture that uses the same elements as in linear repeater chains. We accounted for channel losses between repeater nodes and the probabilistic nature of entanglement swaps at each repeater stemming from device inefficiencies as well as the probabilistic nature of Bell-state measurements (e.g., due to inherent limitations of using linear optics and lossy detectors). The rate attained for a single entanglement-generation flow can far outperform that attainable over a linear repeater chain, even when the nodes only have local link-state knowledge, due to the multi-path routing advantage. We also proposed a modified version of our routing protocol for supporting simultaneous entanglement generation flows between multiple Alice-Bob pairs. We found multi-flow entanglement routing strategies that outperform the rate region attained when each repeater simply time shares among each flow. Our results suggest that building and connecting quantum repeaters in non-trivial network topologies could provide a substantial benefit over linear repeater chains alone. Seen another way, given constraints on the number and quality of quantum memories, link losses between nodes, and limited and imperfect processing capabilities at repeater nodes, a 2D network topology can outperform the repeaterless rate-vs.-distance upper limits [46, 139] more easily than a linear repeater chain connecting the communicating parties.

Our work has also opened a number of new questions. Even in our simplified model—an abstraction that applies when the only source of imperfection at each component (including the quantum memories) is pure loss—the rate-optimal protocol remains open. Since our protocol only requires a quantum memory to hold a qubit for one entanglement attempt between neighboring stations, photon loss would indeed be the major source of imperfection in many implementations of the protocol. Accounting for more general errors would require purification of entanglement [150, 151, 48], i.e., converting several poorer-quality EPR pairs into a few good ones using local quantum operations and classical communication, accounting which will require us

to introduce the fidelity of shared entanglement at intermediate steps of the protocol. Furthermore, we restricted our analysis to the same operation used in the nodes of a linear repeater chain: 2-qubit measurements. Being able to perform multi-qubit unitary operations and multi-qubit measurements at repeater nodes (e.g., a 3-qubit GHZ projection across three locally held qubits) may improve the achievable rate regions. The idea of using a distance metric to choose the measurements at the repeater station could be used in protocols that use measurements of more than two qubits as well. Finally, it will be interesting to consider repeater protocols for the distillation of multi-partite entanglement shared between more than two parties, and a repeater network that can support multiple simultaneous flows of generation of multi-partite entanglement.

Chapter 6

Devices

6.1 Introduction

In this section I briefly present an overview of experimental and theoretical work on devices for the architectures discussed in this thesis that I have been involved in. While I have been the lead author in the work presented in the chapters of this thesis up to this point, the work presented here has been led by my colleagues and their names have been stated at the beginning relevant sections. Section 6.2 presents a theoretical proposal for an on-demand single photon source [158] and an experimental demonstration of a spontaneous four wave mixing based heralded photon source [159]. Section 6.3 presents a theoretical proposal for a cavity mediated optical photodetection scheme that may be useful for implementing feed-forward.

6.2 Single photon sources

6.2.1 Temporally and frequency multiplexed single photon source using quantum feedback control for scalable photonic quantum technologies

Lead Author: Mikkel Heuck

In this work [158], we investigate the feasibility of single photon sources that meet the requirements of scalable photonic quantum technologies: near-unity purity single photons produced in a reproducable chip-integrated photonic circuit. Our proposal uses temporal multiplexing of parametrically produced signal-idler photon pairs and we explain how additional multiplexing of the frequency degree of freedom may lead to significantly improved performance.

As illustrated in Fig. 6-1, our proposed device consists of an ultrahigh-Q microring resonator (Q of 10-100M) consisting of a material, such as silicon, with a $\chi^{(3)}$ nonlinearity for photon pair generation by sFWM. This storage ring is coupled to photon number resolving detectors (PNRDs) through Mach-Zehnder interferometer (MZI) filters [160]. The filters enable decoupling of certain frequencies from the waveguide by controlling the path-imbalance of the MZI relative to the length of the ring (see inset in Fig. 6-1). Idler photons and the pump field couple out of the storage ring within a single time bin, whereas the signal can be stored for up to M bins. If M is large, the state of the cavity can be pushed close to a single photon state by probabilistic addition of photons into the cavity (by pumping the ring) and the probabilistic subtraction of photons (by coupling the signal photons out), based on the detection of the idler photons in previous time bins. We consider a control protocol based on Bayesian inference with both idler and signal photon detection to optimize the signal-photon state. This approach shows the trade-off between heralding the generation of a single photon state and its purity. Our study reveals that, for near-term realistic device parameters, highly efficient ($\sim$99%) sources of single photons could be possible in scalable nanophotonic platforms.

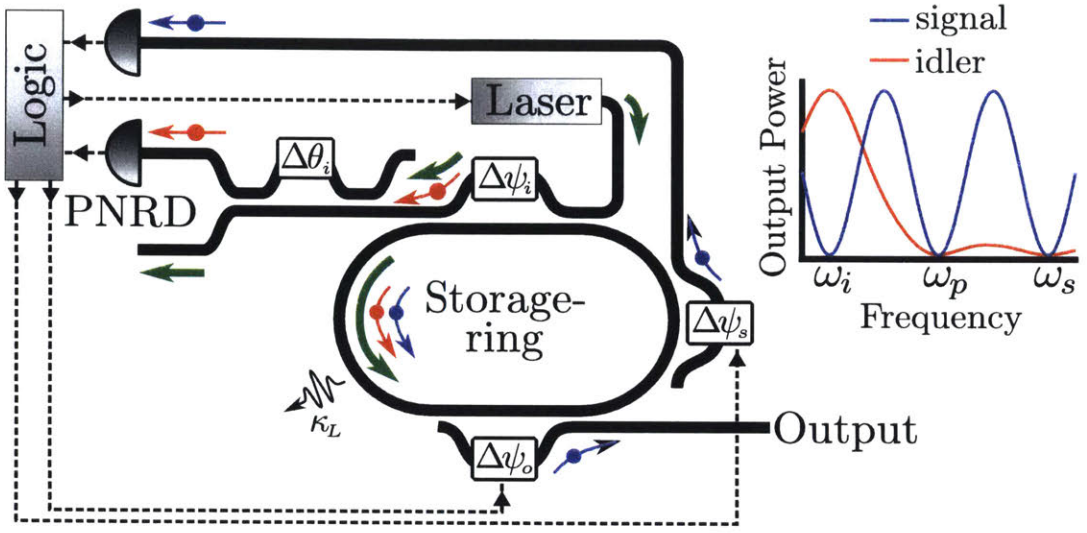

Figure 6-1: Storage-and release design. Solid lines are optical waveguides, while dashed lines represent electrical control signals. PNRD: photon number resolving detector. The inset illustrates the power spectrum coupled out of the signal filter in its closed configuration and the spectrum arriving at the idler detector.

6.2.2 Integrated source of spectrally filtered correlated photons for large-scale quantum photonic systems

Lead Authors: Nicholas Harris and Davide Grassani

Stimulated four wave mixing (sFWM) is a popular method for generating photon pairs on a photonic chip, which can be used as a heralded single photon source. However, generation of heralded photon pairs on chip using sFWM uses a pump field that is $\approx 10^{10}$ times larger than the desired pairs. Achieving such large isolation on-chip is a challenging, but important step, because any unwanted pump field would lead to extraneous counts, and could also lead to the creation of additional unwanted photon pairs. In this work [159], we demonstrate the generation of quantum-correlated photon pairs combined with the spectral filtering of the pump field by more than 95 dB on a single silicon chip using electrically tunable ring resonators and passive Bragg reflectors. Moreover, we perform the demultiplexing and routing of signal and idler photons after transferring them via an optical fiber to a second identical chip. Nonclassical two-photon temporal correlations with a coincidence-to-accidental ratio

of 50 are measured without further off-chip filtering. Fig. 6-2 shows the schematic of the source and the experimental setups used to characterize the source. Fig. 6-3 shows an optical micrograph of the chip.

6.3 A cavity-enabled technique for optical, quantum limited photodetection

Lead Authors: Christopher Panuski

Traditional photodetectors, which rely upon the electronic amplification of photogenerated charge carriers, are often the performance-limiting component of both classical and quantum optical systems. To circumvent the fundamental limitations associated with these devices, we introduce an alternative, semiconductor-based photodetection architecture, a schematic for which is shown in Fig. 6-4, in which signal amplification is achieved in the optical domain through the incorporation of a high quality photonic cavity. We demonstrate the feasibility of realizing single photon amplification at room temperature with all-optical readout, which inherently affords a significantly enhanced detection bandwidth, orders of magnitude lower thermal noise, and minimal channel attenuation as compared to classical electronic detectors. Such a system may be especially promising for implementing optical feed-forward with optical modulation of phase shifters, without ever requiring conversion into the electronic domain.

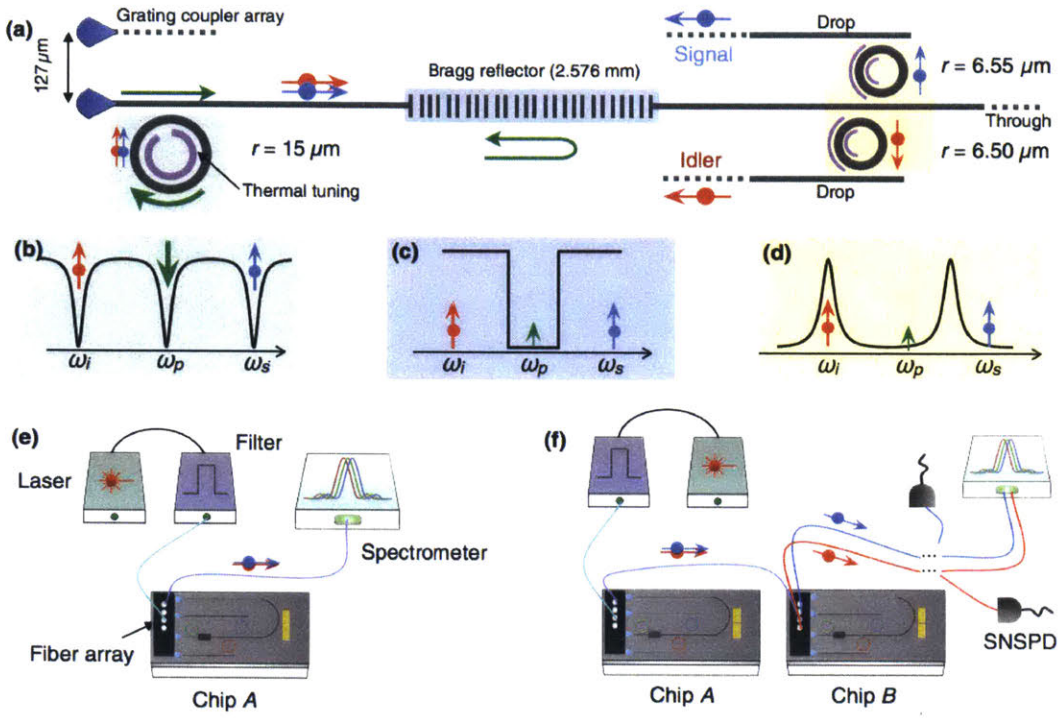

Figure 6-2: (a) Schematic layout of the photonic integrated circuit composed of a high-Q thermally tunable ring for efficient pair generation by spontaneous four-wave mixing, followed by a DBR for pump rejection and the add-drop ring-resonator filters for the demultiplexing of signal and idler photons. Convenient optical coupling to a single-mode polarization-maintaining fiber array is achieved via focusing grating couplers separated by a 127μm pitch. (b) Schematic transmission spectrum of the first ring around the pump wavelength ω_p. When one of the ring resonances is tuned to the laser at ω_p, signal and idler photons are produced in correlated pairs at neighboring resonance wavelengths ω_s and ω_i, respectively. (Pairs are also generated at wavelengths spaced by multiple free spectral ranges.) (c) Schematic transmission spectrum of the DBR with the stop band overlapping with the pump wavelength ω_p. (d) Add-drop filter spectrum tuned to route idler photons to the drop port. (e) First experimental setup: single-chip pump rejection. The add-drop rings are both tuned on resonance with the pump. Light is collected from the common throughport. (f) Second experimental setup: Correlated photon pairs generated in chip A are sent via a fiber to chip B where further pump rejection and signal or idler demultiplexing are performed before spectral characterization or coincidence measurements with off-chip SNSPDs.

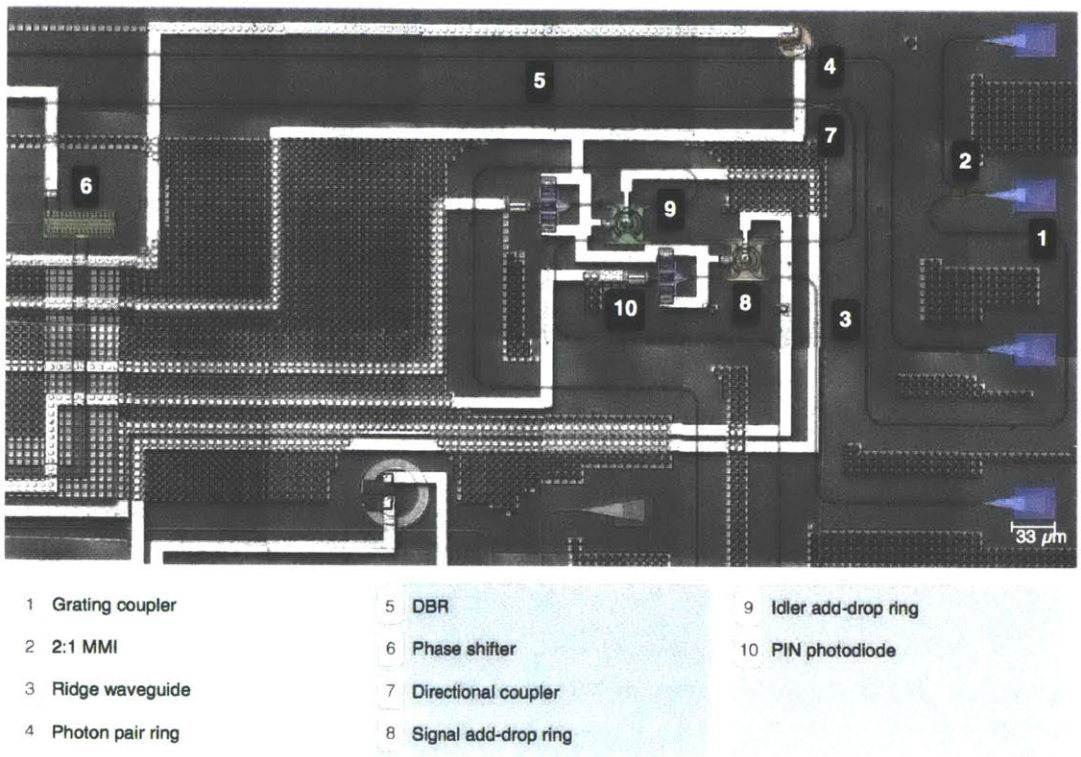

Figure 6-3: Optical micrograph of the source (one of four on the chip in an area of 2.4 by 1.36mm^2). Two grating couplers are not shown. (1) Grating couplers used to couple (collect) light to (from) the system are shown on the right. The input light is split by (2) a 2:1 multimode interferometer for optical alignment. Pump light is then routed via (3) a 500-by-220-nm ridge waveguide to (4) the pair-generation ring. The pump is removed with (5) the DBR, which is divided into two sections [(6) Fabry-Perot resonances due to the division can be controlled with the thermo-optic phase shifter [6]]. The multiplexed signal and idler photon combs are then split off for spectral monitoring at (7) the directional coupler before demultiplexing and/or further filtering with the (8) signal and (9) idler add-drop rings. (10) The p-doped/intrinsic/n-doped germanium photodiodes were not used during the experiment; however, they could be used to monitor the add-drop ring alignment.

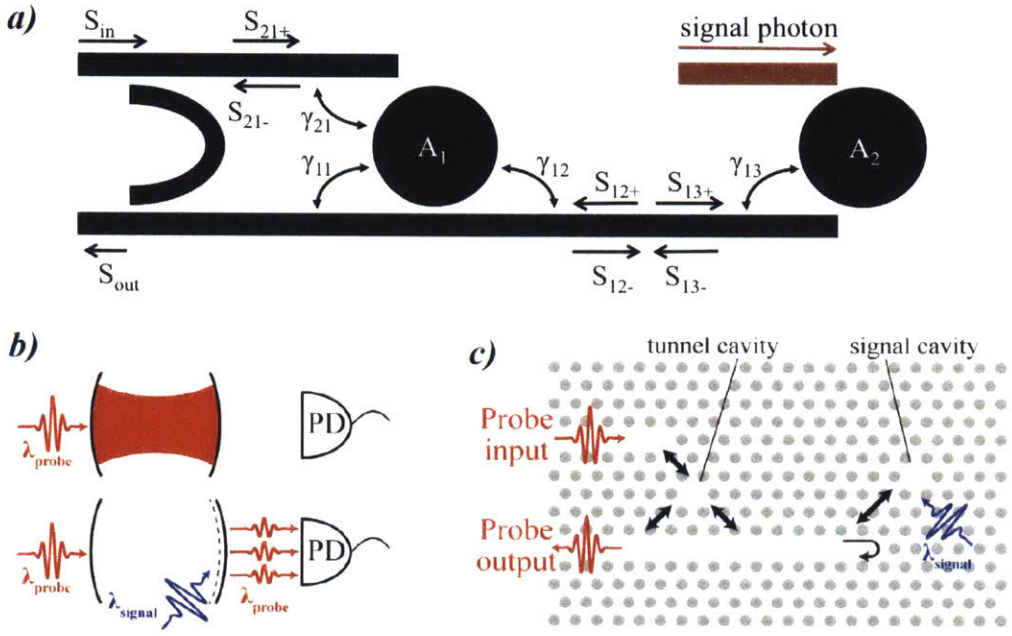

Figure 6-4: a) Schematic of the coupled cavity detector. Similar to a Q-switched laser, cavity A_1 is in a high-Q state due to destructive interference between its output decay channel and light returning from cavity A_2. Injecting a signal photon into A_2 disrupts this interference condition, creating a low-Q configuration which quickly flushes out (b) the energy stored in the composite system. c) A possible experimental implementation of the coupled cavity configuration using a photonic crystal. Here A_1 is the tunnel cavity, and A_2 is the signal cavity.

Chapter 7

Conclusion

There have been many small scale demonstrations of experimental quantum information processing in different physical systems. There have also been impressive theoretical breakthroughs in quantum algorithms and quantum error correction. Given the inherently better scaling of some quantum algorithms, it has been known for some time that a large enough quantum computer will surpass any classical computer. However, moving from proof of concept demonstrations to useful quantum processors that are large enough to outperform classical computers presents a number of challenges. Actually building a useful quantum information processor will require experimental progress towards scalable platforms capable of housing large numbers of qubits with a low error rate. Furthermore, realistic theory will be required which goes beyond just scaling and thresholds, and quantitatively lays out device requirements for experimentalists, while using the simplest possible resources.

The focus of this thesis was on photonic quantum information processing which has several advantages including scalability and negligible decoherence rate, but must contend with the probabilistic nature of linear optic gates. In Chapter 2, we studied the resource requirements for building an all-optical repeater that is "useful" i.e. capable of beating the repeaterless bound, even when the sources, waveguides and detectors used to create the required entangled states are lossy. Device parameters that are aggressive but achievable in the near future were assumed. A 48 photon entangled state is capable of working as a useful quantum repeater. However, making

this resource state with probabilistic linear optic gates requires 10^5 photon sources. While this number is still extremely large, the improvements presented in this thesis reduced the resource requirements by six orders of magnitude in this thesis, which suggests that there is a lot of room for improvement in cluster state creation and error correction. Furthermore, we found that if sources of three-photon Greenberger-Horne-Zeilinger(GHZ) states were available, the resource requirements would go down to 1000 such sources.

To make the cluster state creation process more efficient and reduce the amount of feed-forward required, we studied a more efficient way of generating photonic cluster states using percolation theory in Chapter 3. We presented a mapping from a set of fusion operations to a logical graph: the bond percolation threshold of the logical graph gives us the required success probability of each entangling operation in order to obtain a resource for universal quantum computation. This allowed us to find configurations in which low success probability entangling gates can still allow for universal quantum computation. Furthermore, we were settled a hitherto unsolved problem settle by showing that a renormalizable cluster can be created with 3-photon microclusters over a 2D graph without feedforward, which is attractive for an integrated photonic realization. We proved a lower bound on the required success probability of each fusion operation for scalable, feed-forward free creation of large cluster states, starting with n photon microclusters: $1/(n-1)$.

Chapter 4 presented an architecture for cluster state quantum information processing with nitrogen vacancy centers in diamond which also uses percolation theory. The architecture reduces the required coherence time for the qubit and has a natural tolerance to faulty sites.

We then presented protocols for a quantum network architecture in Chapter 5 which only require entanglement swaps at repeater stations and knowledge of neighboring links. Even with such limitations, the network provides an improved scaling of the entanglement generation rate and is capable of connecting multiple parties simultaneously.

Finally, Chapter 6, presented experimental and theoretical work on devices for

photonic quantum information processing that I was involved in.

Recent work has shown how percolated lattice can be renormalized to a fault tolerant Raussendorf lattice [161, 162]. Future work should focus on calculating the resource requirements for fault tolerant quantum computing with the percolation based approach which should have lower resource requirements. The error models used in this thesis have been mostly limited to photon loss. The codes used here do provide tolerance against bit flip and phase flip errors, but the exact mapping of the error channel induced by photon distinguishability and impurity should be studied in future work to determined to determine the corresponding experimental requirements.

Bibliography

[1] R. Raussendorf and H. J. Briegel, "A One-Way Quantum Computer," *Physical Review Letters*, vol. 86, pp. 5188–5191, may 2001.

[2] M. Varnava, D. E. Browne, and T. Rudolph, "How good must single photon sources and detectors Be for efficient linear optical quantum computation?," *Physical Review Letters*, vol. 100, p. 060502, feb 2008.

[3] D. E. Browne and T. Rudolph, "Resource-Efficient Linear Optical Quantum Computation," *Physical Review Letters*, vol. 95, p. 010501, jun 2005.

[4] M. Gimeno-Segovia, P. Shadbolt, D. E. Browne, and T. Rudolph, "From Three-Photon Greenberger-Horne-Zeilinger States to Ballistic Universal Quantum Computation," *Physical Review Letters*, vol. 115, p. 020502, jul 2015.

[5] T. Rudolph, "Why I am optimistic about the photonic route to quantum computing," *arXiv:1607.08535*, jul 2016.

[6] N. C. Harris, Y. Ma, J. Mower, T. Baehr-Jones, D. Englund, M. Hochberg, and C. Galland, "Efficient, compact and low loss thermo-optic phase shifter in silicon," *Optics Express*, vol. 22, p. 10487, may 2014.

[7] P. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pp. 124–134, IEEE Comput. Soc. Press, 1994.

[8] N. Wiebe, D. Braun, and S. Lloyd, "Quantum Algorithm for Data Fitting," *Physical Review Letters*, vol. 109, p. 050505, aug 2012.

[9] F. G. S. L. Brandao and K. Svore, "Quantum Speed-ups for Semidefinite Programming," sep 2016.

[10] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96*, (New York, New York, USA), pp. 212–219, ACM Press, jul 1996.

[11] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum algorithms for supervised and unsupervised machine learning," jul 2013.

[12] S. Lloyd, "Universal Quantum Simulators," *Science*, vol. 273, pp. 1073–1078, aug 1996.

[13] A. Vergis, K. Steiglitz, and B. Dickinson, "The complexity of analog computation," *Mathematics and Computers in Simulation*, vol. 28, pp. 91–113, apr 1986.

[14] S. Aaronson and Scott, "Guest Column: NP-complete problems and physical reality," *ACM SIGACT News*, vol. 36, p. 30, mar 2005.

[15] D. Gottesman, "Fault-tolerant quantum computation with constant overhead," *Quantum Information & Computation*, vol. 14, pp. 1338–1372, nov 2014.

[16] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, "Blueprint for a microwave trapped ion quantum computer," *Science Advances*, vol. 3, no. 2, 2017.

[17] J. Carolan, C. Harrold, C. Sparrow, E. Martin-Lopez, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh, G. D. Marshall, M. G. Thompson, J. C. F. Matthews, T. Hashimoto, J. L. O'Brien, and A. Laing, "Universal linear optics," *Science*, vol. 349, no. 6249, pp. 711–716, 2015.

[18] N. C. Harris, G. R. Steinbrecher, M. Prabhu, Y. Lahini, J. Mower, D. Bunandar, C. Chen, F. N. C. Wong, T. Baehr-Jones, M. Hochberg, S. Lloyd, and D. Englund, "Quantum transport simulations in a programmable nanophotonic processor," *Nature Photonics*, vol. 11, pp. 447–452, jun 2017.

[19] N. C. Harris, D. Bunandar, M. Pant, G. R. Steinbrecher, J. Mower, M. Prabhu, T. Baehr-Jones, M. Hochberg, and D. Englund, "Large-scale quantum photonic circuits in silicon," *Nanophotonics*, vol. 5, pp. 456–468, jan 2016.

[20] T. Rudolph, "Why I am optimistic about the silicon-photonic route to quantum computing," *APL Photonics*, vol. 2, p. 030901, mar 2017.

[21] I. Chuang and Y. Yamamoto, "Simple quantum computer," *Physical Review A*, vol. 52, pp. 3489–3496, nov 1995.

[22] N. K. Langford, S. Ramelow, R. Prevedel, W. J. Munro, G. J. Milburn, and A. Zeilinger, "Efficient quantum computing using coherent photon conversion.," *Nature*, vol. 478, pp. 360–3, oct 2011.

[23] M. Y. Niu, I. L. Chuang, and J. H. Shapiro, "Universal Quantum Computation using Coherent χ (2) Interactions," 2017.

[24] J. Shapiro, "Single-photon Kerr nonlinearities do not help quantum computation," *Physical Review A*, vol. 73, p. 062305, jun 2006.

[25] J. Dove, C. Chudzicki, and J. H. Shapiro, "Phase-noise limitations on single-photon cross-phase modulation with differing group velocities," *Physical Review A*, vol. 90, p. 062314, dec 2014.

[26] S. Xu, E. Rephaeli, and S. Fan, "Analytic Properties of Two-Photon Scattering Matrix in Integrated Quantum Systems Determined by the Cluster Decomposition Principle," *Physical Review Letters*, vol. 111, p. 223602, nov 2013.

[27] D. J. Brod and J. Combes, "Passive CPHASE Gate via Cross-Kerr Nonlinearities," *Physical Review Letters*, vol. 117, no. 8, 2016.

[28] E. Knill, R. Laflamme, and G. J. Milburn, "A scheme for efficient quantum computation with linear optics.," *Nature*, vol. 409, pp. 46–52, jan 2001.

[29] P. Kok, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, "Linear optical quantum computing with photonic qubits," *Reviews of Modern Physics*, vol. 79, pp. 135–174, jan 2007.

[30] D. Browne, E. Kashefi, and S. Perdrix, "Computational depth complexity of measurement-based quantum computation," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6519 LNCS, pp. 35–46, 2011.

[31] R. Raussendorf, J. Harrington, and K. Goyal, "A fault-tolerant one-way quantum computer," *Annals of Physics*, vol. 321, pp. 2242–2270, sep 2006.

[32] R. Raussendorf and J. Harrington, "Fault-Tolerant Quantum Computation with High Threshold in Two Dimensions," *Physical Review Letters*, vol. 98, p. 190504, may 2007.

[33] R. Raussendorf, J. Harrington, and K. Goyal, "Topological fault-tolerance in cluster state quantum computation," *New Journal of Physics*, vol. 9, pp. 199–199, jun 2007.

[34] D. M. Greenberger, M. A. Horne, and A. Zeilinger, "Going Beyond Bell's Theorem," in *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, pp. 69–72, Dordrecht: Springer Netherlands, 1989.

[35] W. P. Grice, "Arbitrarily complete Bell-state measurement using only linear optical elements," *Physical Review A*, vol. 84, p. 042331, oct 2011.

[36] F. Ewert and P. van Loock, "3 / 4 -Efficient Bell Measurement with Passive Linear Optics and Unentangled Ancillae," *Physical Review Letters*, vol. 113, p. 140403, sep 2014.

[37] H. A. Zaidi and P. van Loock, "Beating the One-Half Limit of Ancilla-Free Linear Optics Bell Measurements," *Physical Review Letters*, vol. 110, p. 260501, jun 2013.

[38] A. K. Ekert, "Quantum cryptography based on Bell's theorem," *Physical Review Letters*, vol. 67, pp. 661–663, aug 1991.

[39] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," *Theoretical Computer Science*, vol. 560, pp. 7–11, 2014.

[40] J. Cirac, A. Ekert, S. Huelga, and C. Macchiavello, "Distributed quantum computation over noisy channels," *Physical Review A*, 1999.

[41] D. Gottesman, T. Jennewein, and S. Croke, "Longer-Baseline Telescopes Using Quantum Repeaters," *Physical Review Letters*, vol. 109, p. 070503, aug 2012.

[42] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, "A quantum network of clocks," *Nature Physics*, vol. 10, pp. 582–587, oct 2014.

[43] A. Broadbent, J. Fitzsimons, and E. Kashefi, "Universal Blind Quantum Computation," in *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pp. 517–526, IEEE, oct 2009.

[44] S. Guha, T. Hogg, D. Fattal, T. Spiller, and R. G. Beausoleil, "Quantum Auctions using Adiabatic Evolution : The Corrupt Auctioneer and Circuit," *International Journal of Quantum Information*, vol. 06, pp. 815–839, aug 2008.

[45] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, "The Ultimate Rate of Quantum Communications," *arXiv preprint*, oct 2015.

[46] M. Takeoka, S. Guha, and M. M. Wilde, "Fundamental rate-loss tradeoff for optical quantum key distribution.," *Nature communications*, vol. 5, p. 5235, jan 2014.

[47] S. Pirandola, R. García-Patrón, S. L. Braunstein, and S. Lloyd, "Direct and reverse secret-key capacities of a quantum channel.," *Physical review letters*, vol. 102, p. 050503, feb 2009.

[48] H.-J. Briegel, W. Dür, J. Cirac, and P. Zoller, "Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication," *Physical Review Letters*, vol. 81, pp. 5932–5935, dec 1998.

[49] S. Pirandola, "Capacities of repeater-assisted quantum communications," *arXiv preprint*, jan 2016.

[50] K. Azuma, A. Mizutani, and H.-K. Lo, "Fundamental rate-loss trade-off for the quantum internet," *Nature Communications 2016 7:null*, vol. 7, p. ncomms13523, nov 2016.

[51] K. Azuma, K. Tamaki, and H.-K. Lo, "All-photonic quantum repeaters.," *Nature communications*, vol. 6, p. 6787, jan 2015.

[52] F. Ewert, M. Bergmann, and P. van Loock, "Ultrafast Long-Distance Quantum Communication with Static Linear Optics," *Physical Review Letters*, vol. 117, p. 210501, nov 2016.

[53] Q. Zhang, X.-H. Bao, C.-Y. Lu, X.-Q. Zhou, T. Yang, T. Rudolph, and J.-W. Pan, "Demonstration of a scheme for the generation of âĂIJevent-readyâĂİ entangled photon pairs from a single-photon source," *Physical Review A*, vol. 77, p. 062316, jun 2008.

[54] N. Sinclair, E. Saglamyurek, H. Mallahzadeh, J. A. Slater, M. George, R. Ricken, M. P. Hedges, D. Oblak, C. Simon, W. Sohler, and W. Tittel, "Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control.," *Physical review letters*, vol. 113, p. 053603, aug 2014.

[55] S. Guha, H. Krovi, C. A. Fuchs, Z. Dutton, J. A. Slater, C. Simon, and W. Tittel, "Rate-loss analysis of an efficient quantum repeater architecture," *Physical Review A*, vol. 92, p. 022357, aug 2015.

[56] M. Varnava, D. E. Browne, and T. Rudolph, "Loss tolerant linear optical quantum memory by measurement-based quantum computing," *New Journal of Physics*, vol. 9, pp. 203–203, jun 2007.

[57] M. Varnava, D. Browne, and T. Rudolph, "Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction," *Physical Review Letters*, vol. 97, p. 120501, sep 2006.

[58] Y. Li, P. C. Humphreys, G. J. Mendoza, and S. C. Benjamin, "Resource Costs for Fault-Tolerant Linear Optical Quantum Computing," *Physical Review X*, vol. 5, p. 041007, oct 2015.

[59] A. Gilchrist, A. Hayes, and T. Ralph, "Efficient parity-encoded optical quantum computing," *Physical Review A*, vol. 75, p. 052328, may 2007.

[60] A. J. F. Hayes, A. Gilchrist, and T. C. Ralph, "Loss-tolerant operations in parity-code linear optics quantum computing," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 77, no. 1, 2008.

[61] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," *Physical Review A*, vol. 52, pp. R2493–R2496, oct 1995.

[62] D. Bacon, "Operator quantum error-correcting subsystems for self-correcting quantum memories," *Physical Review A*, vol. 73, p. 012340, jan 2006.

[63] K. Kieling, T. Rudolph, and J. Eisert, "Percolation, Renormalization, and Quantum Computing with Nondeterministic Gates," *Physical Review Letters*, vol. 99, p. 130501, sep 2007.

[64] K. Kieling and J. Eisert, "Percolation in quantum computation and communication," *Lecture Notes in Physics*, vol. 762, pp. 287–319, 2009.

[65] H. A. Zaidi, C. Dawson, P. Van Loock, and T. Rudolph, "Near-deterministic creation of universal cluster states with probabilistic Bell measurements and three-qubit resource states," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 91, p. 042301, apr 2015.

[66] M. Reck and A. Zeilinger, "Experimental realization of any discrete unitary operator," *Physical Review Letters*, vol. 73, pp. 58–61, jul 1994.

[67] S. Aaronson and A. Arkhipov, "The computational complexity of linear optics," in *Proceedings of the 43rd annual ACM symposium on Theory of computing - STOC '11*, (New York, New York, USA), p. 333, ACM Press, jun 2011.

[68] S. Aaronson and A. Arkhipov, "The Computational Complexity of Linear Optics," *Theory of Computing*, vol. 9, pp. 143–252, 2013.

[69] J.-w. Pan and A. Zeilinger, "Greenberger-Horne-Zeilinger-state analyzer," *Physical Review A*, vol. 57, pp. 2208–2211, mar 1998.

[70] N. H. Lindner and T. Rudolph, "Proposal for pulsed On-demand sources of photonic cluster state strings," *Physical Review Letters*, vol. 103, p. 113602, sep 2009.

[71] I. Schwartz, D. Cogan, E. R. Schmidgall, Y. Don, L. Gantz, O. Kenneth, N. H. Lindner, and D. Gershoni, "Deterministic generation of a cluster state of entangled photons.," *Science*, vol. 49, no. 7285, pp. 1804–1807, 2016.

[72] H. J. Briegel and R. Raussendorf, "Persistent Entanglement in Arrays of Interacting Particles," *Physical Review Letters*, vol. 86, pp. 910–913, jan 2001.

[73] M. Michler, K. Mattle, H. Weinfurter, and A. Zeilinger, "Interferometric Bell-state analysis," *Physical Review A*, vol. 53, pp. R1209–R1212, mar 1996.

[74] J. Calsamiglia and N. Lütkenhaus, "Maximum efficiency of a linear-optical Bell-state analyzer," *Applied Physics B*, vol. 72, pp. 67–71, jan 2001.

[75] J. M. Hammersley, "A generalization of McDiarmid's theorem for mixed Bernoulli percolation," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 88, p. 167, jul 1980.

[76] Y. Y. Tarasevich and S. C. van der Marck, "An investigation of site-bond percolation on many lattices," *International Journal of Modern Physics C*, vol. 10, no. 07, p. 14, 1999.

[77] M. E. J. Newman and R. M. Ziff, "Fast Monte Carlo algorithm for site or bond percolation," *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, vol. 64, pp. 1–16, jun 2001.

[78] J. Tran, T. Yoo, S. Stahlheber, and A. Small, "Percolation thresholds on three-dimensional lattices with three nearest neighbors," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2013, p. P05014, may 2013.

[79] M. Van Den Nest, A. Miyake, W. Dür, and H. J. Briegel, "Universal resources for measurement-based quantum computation," *Physical Review Letters*, vol. 97, no. 15, 2006.

[80] M. Pant, H. Krovi, D. Englund, and S. Guha, "Rate-distance tradeoff and resource costs for all-optical quantum repeaters," *Physical Review A*, vol. 95, no. 1, p. 012304, 2017.

[81] S. R. Broadbent and J. M. Hammersley, "Percolation processes," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 53, no. 53, pp. 629–641, 1957.

[82] M. Sahimi, B. D. Hughes, L. E. Scriven, and H. T. Davis, "On Polya random walks, lattice Green functions, and the bond percolation threshold," *J. Phys. A: Math. Gen.*, vol. 16, no. 16, pp. 67–72, 1983.

[83] D. D. B. Rao, S. Yang, and J. Wrachtrup, "Generation of entangled photon strings using NV centers in diamond," *Physical Review B - Condensed Matter and Materials Physics*, vol. 92, p. 081301, aug 2015.

[84] M. Gimeno-Segovia, *Towards practical Linear Optical Quantum Computing.* PhD thesis, Imperial College London, 2015.

[85] S. D. Barrett and T. M. Stace, "Fault Tolerant Quantum Computation with Very High Threshold for Loss Errors," *Physical Review Letters*, vol. 105, p. 200502, nov 2010.

[86] J. Mower, N. C. Harris, G. R. Steinbrecher, Y. Lahini, and D. Englund, "High-fidelity quantum state evolution in imperfect photonic integrated circuits," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 92, jun 2015.

[87] N. C. Harris, G. R. Steinbrecher, J. Mower, Y. Lahini, M. Prabhu, T. Baehr-Jones, M. Hochberg, S. Lloyd, and D. Englund, "Bosonic transport simulations in a large-scale programmable nanophotonic processor," *arXiv: 1507.03406v2*, vol. 24, pp. 1–8, jul 2015.

[88] K. Nemoto, M. Trupke, S. J. Devitt, A. M. Stephens, B. Scharfenberger, K. Buczak, T. Nöbauer, M. S. Everitt, J. Schmiedmayer, and W. J. Munro, "Photonic architecture for scalable quantum information processing in diamond," *Physical Review X*, vol. 4, p. 031022, aug 2014.

[89] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, "Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects," *Physical Review A*, vol. 89, p. 022317, feb 2014.

[90] N. Y. Yao, L. Jiang, A. V. Gorshkov, P. C. Maurer, G. Giedke, J. I. Cirac, and M. D. Lukin, "Scalable architecture for a room temperature solid-state quantum information processor," *Nat Commun*, vol. 3, p. 800, dec 2012.

[91] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, "Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links," *Physical Review X*, vol. 4, no. 4, 2014.

[92] J. Cai, A. Retzker, F. Jelezko, and M. B. Plenio, "Towards a large-scale quantum simulator on diamond surface at room temperature," *Nature Physics*, vol. 9, pp. 168–173, 2013.

[93] J.-B. You, W. L. Yang, Z.-Y. Xu, A. H. Chan, and C. H. Oh, "Phase transition of light in circuit-QED lattices coupled to nitrogen-vacancy centers in diamond," *Physical Review B*, vol. 90, p. 195112, nov 2014.

[94] V. Giovannetti, S. Lloyd, and L. Maccone, "Advances in Quantum Metrology," *Nature Photonics*, vol. 5, pp. 222–229, 2011.

[95] D. E. Browne, M. B. Elliott, S. T. Flammia, S. T. Merkel, A. Miyake, and A. J. Short, "Phase transition of computational power in the resource states for one-way quantum computation," *New Journal of Physics*, vol. 10, p. 023010, feb 2008.

[96] C. W. Chou, H. de Riedmatten, D. Felinto, S. V. Polyakov, S. J. van Enk, and H. J. Kimble, "Measurement-induced entanglement for excitation stored in remote atomic ensembles," *Nature*, vol. 438, pp. 828–832, dec 2005.

[97] L.-m. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics," *Nature*, vol. 414, no. 6862, pp. 413–418, 2001.

[98] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Entanglement of single-atom quantum bits at a distance," *Nature*, vol. 449, pp. 68–71, sep 2007.

[99] M. Reed, J. Randall, R. Aggarwal, R. Matyi, T. Moore, and A. Wetsel, "Observation of discrete electronic states in a zero-dimensional semiconductor nanostructure," *Physical Review Letters*, vol. 60, pp. 535–537, feb 1988.

[100] A. Delteil, S. Zhe, W.-B. Gao, E. Togan, S. Faelt, and A. Imamoglu, "Generation of heralded entanglement between distant hole spins," *Nature Physics*, vol. 12, no. 3, pp. 218–223, 2015.

[101] D. Sridharan and E. Waks, "Generating entanglement between quantum dots with different resonant frequencies based on dipole-induced transparency," *Physical Review A*, vol. 78, p. 052321, nov 2008.

[102] R. Kolesov, K. Xia, R. Reuter, R. Stöhr, A. Zappe, J. Meijer, P. Hemmer, and J. Wrachtrup, "Optical detection of a single rare-earth ion in a crystal," *Nature Communications*, vol. 3, p. 1029, aug 2012.

[103] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. Hollenberg, "The nitrogen-vacancy colour centre in diamond," feb 2013.

[104] I. Aharonovich, A. D. Greentree, and S. Prawer, "Diamond photonics," *Nature Photonics*, vol. 5, no. 7, pp. 397–405, 2011.

[105] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, "Heralded entanglement between solid-state qubits separated by three metres," *Nature*, vol. 497, no. 7447, pp. 86–90, 2013.

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

[107] N. Bar-Gill, L. Pham, A. Jarmola, D. Budker, and R. Walsworth, "Solid-state electronic spin coherence time approaching one second," *Nature Communications*, vol. 4, p. 1743, apr 2013.

[108] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, J. Biamonte, and J. Wrachtrup, "High-fidelity spin entanglement using optimal control," *Nature Communications*, vol. 5, pp. 345–380, feb 2014.

[109] R. S. Said and J. Twamley, "Robust control of entanglement in a nitrogen-vacancy center coupled to a C13 nuclear spin in diamond," *Physical Review A*, vol. 80, p. 032303, sep 2009.

[110] L. Childress, M. V. Gurudev Dutt, J. M. Taylor, A. S. Zibrov, F. Jelezko, J. Wrachtrup, P. R. Hemmer, and M. D. Lukin, "Coherent Dynamics of Coupled Electron and Nuclear Spin Qubits in Diamond," *Science*, vol. 314, no. 5797, 2006.

[111] P. C. Maurer, G. Kucsko, C. Latta, L. Jiang, N. Y. Yao, S. D. Bennett, F. Pastawski, D. Hunger, N. Chisholm, M. Markham, D. J. Twitchen, J. I. Cirac, and M. D. Lukin, "Room-Temperature Quantum Bit Memory Exceeding One Second," *Science*, vol. 336, no. 6086, pp. 1283–1286, 2012.

[112] L. Robledo, L. Childress, H. Bernien, B. Hensen, P. F. Alkemade, and R. Hanson, "High-fidelity projective read-out of a solid-state spin quantum register," *Nature*, vol. 477, no. 7366, pp. 574–578, 2011.

[113] S. L. Mouradian, T. Schröder, C. B. Poitras, L. Li, J. Goldstein, E. H. Chen, M. Walsh, J. Cardenas, M. L. Markham, D. J. Twitchen, M. Lipson, and D. Englund, "Scalable Integration of Long-Lived Quantum Memories into a Photonic Circuit," *Physical Review X*, vol. 5, p. 031009, jul 2015.

[114] M. Z. Bazant, "Largest cluster in subcritical percolation," *Physical Review E*, vol. 62, pp. 1660–1669, aug 2000.

[115] S. D. Barrett and P. Kok, "Efficient high-fidelity quantum computation using matter qubits and linear optics," *Physical Review A*, vol. 71, p. 060310, jun 2005.

[116] J. Metz and S. D. Barrett, "Effect of frequency-mismatched photons in quantum-information processing," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 77, no. 4, 2008.

[117] N. Calandri, Q. Y. Zhao, D. Zhu, A. Dane, and K. K. Berggren, "Superconducting nanowire detector jitter limited by detector geometry," *Applied Physics Letters*, vol. 109, no. 15, 2016.

[118] P. Tamarat, T. Gaebel, J. R. Rabeau, M. Khan, A. D. Greentree, H. Wilson, L. C. L. Hollenberg, S. Prawer, P. Hemmer, F. Jelezko, and J. Wrachtrup, "Stark shift control of single optical centers in diamond," *Physical Review Letters*, vol. 97, no. 8, 2006.

[119] V. M. Acosta, C. Santori, A. Faraon, Z. Huang, K.-M. C. Fu, A. Stacey, D. A. Simpson, K. Ganesan, S. Tomljenovic-Hanic, A. D. Greentree, S. Prawer, and R. G. Beausoleil, "Dynamic Stabilization of the Optical Resonances of Single Nitrogen-Vacancy Centers in Diamond," *Physical Review Letters*, vol. 108, p. 206401, may 2012.

[120] Y. Chu, N. de Leon, B. Shields, B. Hausmann, R. Evans, E. Togan, M. J. Burek, M. Markham, A. Stacey, A. Zibrov, A. Yacoby, D. Twitchen, M. Loncar, H. Park, P. Maletinsky, and M. Lukin, "Coherent Optical Transitions in Implanted Nitrogen Vacancy Centers," *Nano Letters*, vol. 14, pp. 1982–1986, apr 2014.

[121] D. Riedel, I. Söllner, B. J. Shields, S. Starosielec, P. Appel, E. Neu, P. Maletinsky, and R. J. Warburton, "Deterministic enhancement of coherent photon generation from a nitrogen-vacancy center in ultrapure diamond," mar 2017.

[122] T. Ishikawa, K.-M. C. Fu, C. Santori, V. M. Acosta, R. G. Beausoleil, H. Watanabe, S. Shikata, and K. M. Itoh, "Optical and Spin Coherence Properties of Nitrogen-Vacancy Centers Placed in a 100 nm Thick Isotopically Purified Diamond Layer," *Nano Letters*, vol. 12, pp. 2083–2087, apr 2012.

[123] J. R. Rabeau, P. Reichart, G. Tamanyan, D. N. Jamieson, S. Prawer, F. Jelezko, T. Gaebel, I. Popa, M. Domhan, and J. Wrachtrup, "Implantation of labelled single nitrogen vacancy centers in diamond using N15," *Applied Physics Letters*, vol. 88, p. 023113, jan 2006.

[124] I. Bayn, E. H. Chen, M. E. Trusheim, L. Li, T. Schröder, O. Gaathon, M. Lu, A. Stein, M. Liu, K. Kisslinger, H. Clevenson, and D. Englund, "Generation of Ensembles of Individually Resolvable Nitrogen Vacancies Using Nanometer-Scale Apertures in Ultrahigh-Aspect Ratio Planar Implantation Masks," *Nano Letters*, vol. 15, pp. 1751–1758, mar 2015.

[125] P. Siyushev, H. Pinto, M. Voros, A. Gali, F. Jelezko, and J. Wrachtrup, "Optically controlled switching of the charge state of a single nitrogen-vacancy center in diamond at cryogenic temperatures," *Physical Review Letters*, vol. 110, no. 16, p. 167402, 2013.

[126] P. Ji and M. V. G. Dutt, "Charge state dynamics of the nitrogen vacancy center in diamond under 1064-nm laser excitation," *Physical Review B*, vol. 94, p. 024101, jul 2016.

[127] L. Li, E. H. Chen, J. Zheng, S. L. Mouradian, F. Dolde, T. Schröder, S. Karaveli, M. L. Markham, D. J. Twitchen, and D. Englund, "Efficient Photon Collection from a Nitrogen Vacancy Center in a Circular Bullseye Grating," *Nano Letters*, vol. 15, pp. 1493–1497, mar 2015.

[128] J. P. Hadden, J. P. Harrison, A. C. Stanley-Clarke, L. Marseglia, Y.-L. D. Ho, B. R. Patton, J. L. O'Brien, and J. G. Rarity, "Strongly enhanced photon collection from diamond defect centers under microfabricated integrated solid immersion lenses," *Applied Physics Letters*, vol. 97, p. 241901, dec 2010.

[129] B. J. Hausmann, M. Khan, Y. Zhang, T. M. Babinec, K. Martinick, M. McCutcheon, P. R. Hemmer, and M. Lončar, "Fabrication of diamond nanowires for quantum information processing applications," *Diamond and Related Materials*, vol. 19, pp. 621–629, may 2010.

[130] L. Li, T. Schr\"oder, E. H. Chen, M. Walsh, I. Bayn, J. Goldstein, O. Gaathon, M. E. Trusheim, M. Lu, J. Mower, M. Cotlet, M. L. Markham, D. J. Twitchen, and D. Englund, "Coherent spin control of a nanocavity-enhanced qubit in diamond.," *Nature Comm.*, vol. 6, p. 6173, jan 2015.

[131] A. Sipahigil, R. E. Evans, D. D. Sukachev, M. J. Burek, J. Borregaard, M. K. Bhaskar, C. T. Nguyen, J. L. Pacheco, H. A. Atikian, C. Meuwly, R. M. Camacho, F. Jelezko, E. Bielejec, H. Park, M. Lončar, and M. D. Lukin, "An integrated diamond nanophotonics platform for quantum-optical networks," *Science*, vol. 354, no. 6314, 2016.

[132] F. Marsili, V. B. Verma, J. A. Stern, S. Harrington, A. E. Lita, T. Gerrits, I. Vayshenker, B. Baek, M. D. Shaw, R. P. Mirin, and S. W. Nam, "Detecting single infrared photons with 93% system efficiency," *Nature Photonics*, vol. 7, no. 3, pp. 210–214, 2013.

[133] A. Faraon, P. E. Barclay, C. Santori, K.-M. C. Fu, and R. G. Beausoleil, "Resonant enhancement of the zero-phonon emission from a color center in a diamond cavity," *Nature Photonics*, vol. 5, no. 5, p. 5, 2010.

[134] K. Srinivasan and O. Painter, "Linear and nonlinear optical spectroscopy of a strongly coupled microdiskâĂŞquantum dot system," *Nature*, vol. 450, pp. 862–865, dec 2007.

[135] A. Reiserer, N. Kalb, M. S. Blok, K. J. M. van Bemmelen, T. H. Taminiau, R. Hanson, D. J. Twitchen, and M. Markham, "Robust Quantum-Network Memory Using Decoherence-Protected Subspaces of Nuclear Spins," *Physical Review X*, vol. 6, p. 021040, jun 2016.

[136] P. Cappellaro and M. D. Lukin, "Quantum correlation in disordered spin systems: Applications to magnetic sensing," *Physical Review A*, vol. 80, p. 032311, sep 2009.

[137] L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, "Fault-Tolerant Quantum Communication Based on Solid-State Photon Emitters," *Physical Review Letters*, vol. 96, p. 070504, feb 2006.

[138] H. J. Kimble, "The quantum internet.," *Nature*, vol. 453, pp. 1023–30, jun 2008.

[139] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, "Fundamental Limits of Repeaterless Quantum Communications," *Nature Communications*, vol. 8, p. 15043, oct 2017.

[140] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," *Technical report number CS-200006, Duke University*, no. CS-200006, pp. 1–14, 2000.

[141] S. Biswas and R. Morris, "ExOR: Opportunistics Multi-Hop Routing for Wireless Networks," *Sigcomm*, 2005.

[142] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," *Nature*, vol. 299, pp. 802–803, oct 1982.

[143] D. Dieks, "Communication by EPR devices," *Physics Letters A*, vol. 92, pp. 271–272, nov 1982.

[144] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications - SIGCOMM '04*, vol. 34, (New York, New York, USA), p. 145, ACM Press, 2004.

[145] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, "MaxProp: Routing for vehicle-based disruption-tolerant networks," *Proceedings - IEEE INFOCOM*, vol. 00, no. c, 2006.

[146] Z.-S. Yuan, Y.-A. Chen, B. Zhao, S. Chen, J. Schmiedmayer, and J.-W. Pan, "Experimental demonstration of a BDCZ quantum repeater node," *Nature*, vol. 454, pp. 1098–1101, aug 2008.

[147] S. Olmschenk, D. N. Matsukevich, P. Maunz, D. Hayes, L.-M. Duan, and C. Monroe, "Quantum Teleportation Between Distant Matter Qubits," *Science*, vol. 323, no. 5913, pp. 486–489, 2009.

[148] J.-W. Pan, D. Bouwmeester, H. Weinfurter, and A. Zeilinger, "Experimental Entanglement Swapping: Entangling Photons That Never Interacted," *Physical Review Letters*, vol. 80, pp. 3891–3894, may 1998.

[149] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Entanglement of single-atom quantum bits at a distance," *Nature*, vol. 449, pp. 68–71, sep 2007.

[150] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, "Quantum Privacy Amplification and the Security of Quantum Cryptography over Noisy Channels," *Physical Review Letters*, vol. 77, pp. 2818–2821, sep 1996.

[151] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels," *Physical Review Letters*, vol. 76, pp. 722–725, jan 1996.

[152] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, "Quantum repeater with encoding," *Physical Review A*, vol. 79, p. 032325, mar 2009.

[153] M. M. Wilde, M. Tomamichel, and M. Berta, "Converse bounds for private communication over quantum channels," *IEEE Transactions on Information Theory*, vol. 63, pp. 1792–1817, feb 2017.

[154] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, "Optimal architectures for long distance quantum communication," *Scientific Reports*, vol. 6, no. 1, p. 20463, 2016.

[155] E. Schoute, L. Mancinska, T. Islam, I. Kerenidis, and S. Wehner, "Shortcuts to quantum network routing," *arXiv preprint*, pp. 1–45, oct 2016.

[156] A. Acín, J. I. Cirac, and M. Lewenstein, "Entanglement percolation in quantum networks," *Nature Physics*, vol. 3, pp. 256–259, apr 2007.

[157] R. Van Meter, *Quantum Networking*. Wiley, 2014.

[158] M. Heuck, M. Pant, and D. Englund, "Temporally and frequency multiplexed single photon source using quantum feedback control for scalable photonic quantum technologies," aug 2017.

[159] N. C. Harris, D. Grassani, A. Simbula, M. Pant, M. Galli, T. Baehr-Jones, M. Hochberg, D. Englund, D. Bajoni, and C. Galland, "Integrated Source of Spectrally Filtered Correlated Photons for Large-Scale Quantum Photonic Systems," *Physical Review X*, vol. 4, p. 041047, dec 2014.

[160] C. K. Madsen and J. H. Zhao, *Optical Filter Design and Analysis: A Signal Processing Approach.* Wiley Online Library, 1999.

[161] J. M. Auger, H. Anwar, M. Gimeno-Segovia, T. M. Stace, and D. E. Browne, "Fault-tolerant quantum computation with non-deterministic entangling gates," aug 2017.

[162] D. Herr, A. Paler, S. J. Devitt, and F. Nori, "A local and scalable lattice renormalization method for ballistic quantum computation," oct 2017.