

CONNECTING TWO QUANTUM CITIES

*Understanding the Requirements
for Long-Distance Quantum Links*

MASTER'S THESIS



Tobias Ploeckinger

Connecting two quantum cities: understanding the requirements for long-distance quantum links

THESIS

submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE

in

APPLIED PHYSICS

by

Tobias Ploeckinger
born in Vienna, Austria



Wehner Group
QuTech, Delft University of Technology
Delft, the Netherlands

Connecting two quantum cities: understanding the requirements for long-distance quantum links

Author: Tobias Ploeckinger

Abstract

Long-distance quantum networks are expected to enable applications that rely on entanglement distributed between users in different metropolitan areas. A central question is what performance a long-distance entanglement-distribution link connecting these metropolitan areas must provide for such applications to become feasible on near-term hardware. This thesis studies this question for an intercity architecture in which users connect to local border nodes through metropolitan links, while the long-distance backbone connecting these border nodes is modeled abstractly by its entanglement delivery rate and fidelity. Using a NetSquid-based simulation framework, three benchmark protocols are analyzed: quantum key distribution, the CHSH game, and verifiable blind quantum computation. The main result is a protocol-dependent characterization of feasibility in backbone parameter space. For each protocol, thresholds are identified for the minimum metropolitan hardware improvement, backbone fidelity, and backbone rate required for successful implementation, and combined into feasibility maps. These maps reveal distinct tradeoffs between backbone quality and delivery rate, providing a comparative benchmark for assessing candidate backbone technologies based on hardware maturity.

Thesis Committee:

| | |
|--------------------|---------------------------------------|
| Supervisor: | Prof. Dr. S. Wehner, QuTech, TU Delft |
| Daily supervisor: | J. van Dam, QuTech, TU Delft |
| Committee Members: | Dr. S. Feld, QuTech, TU Delft |
| | Dr. F. L. Machado, QuTech, TU Delft |

Preface

Before presenting the work in this thesis, I would like to take a moment to thank the people who supported me throughout this project.

I am especially grateful to my daily supervisor, Janice van Dam, for her continued support, guidance, and feedback throughout the entirety of this thesis. Her insight and encouragement were invaluable at every stage of the project. I would also like to sincerely thank my supervisor, Stephanie Wehner, for the opportunity to carry out this research within the Wehner group and for her feedback and guidance during the project. I am also grateful to the members of my thesis committee for taking the time to evaluate this work and participate in the defense.

I also thank the members of the Wehner group for the helpful discussions, feedback, and pleasant working environment they provided throughout this thesis. In addition, I am grateful to my friends for their support and for the many interesting discussions along the way.

Finally, I would like to thank my parents and family for their continuous support and encouragement throughout my studies. I am also deeply grateful to my girlfriend for her patience, understanding, and constant support during the demanding periods of this thesis.

Tobias Ploeckinger
Delft, the Netherlands
June 16, 2026

Contents

| | |
|--|------------|
| Preface | iii |
| Contents | v |
| 1 Introduction | 1 |
| 1.1 Research Questions | 3 |
| 2 Background | 5 |
| 2.1 Foundation | 5 |
| 2.2 The network architecture | 14 |
| 3 Methods | 19 |
| 3.1 Network setup | 19 |
| 3.2 Simulation framework | 21 |
| 3.3 Protocol implementation | 26 |
| 3.4 Summary | 39 |
| 4 Results and discussion | 41 |
| 4.1 Perfect backbone | 41 |
| 4.2 Fast backbone | 44 |
| 4.3 Quality backbone | 47 |
| 4.4 Full backbone characterization | 52 |
| 4.5 Further discussion | 55 |
| 5 Conclusions and future work | 61 |
| 5.1 Contributions | 61 |
| 5.2 Conclusions | 61 |
| 5.3 Future work | 63 |
| Bibliography | 65 |

CONTENTS

| | | |
|----------|---|-----------|
| A | Additional material | 71 |
| A.1 | Declaration of AI usage | 71 |
| A.2 | Analytical model for the 5-qubit linear graph | 71 |
| A.3 | Extra Figures | 74 |

Chapter 1

Introduction

The vision of a quantum internet is to complement the existing classical internet with the ability to transmit and process quantum information between distant locations [41]. By distributing entanglement across a network, users would gain access to applications that have no classical counterpart: quantum key distribution offers cryptographic security grounded in the laws of physics rather than computational assumptions, non-local coordination games exploit correlations that no classical strategy can reproduce, and blind quantum computation enables clients to delegate computations to a remote server without revealing their input, algorithm, or output. These capabilities have motivated substantial research efforts aimed at realizing the hardware and network architectures required to bring a quantum internet closer to reality.

Considerable effort has recently been devoted to determining, primarily through simulation-based studies, what level of hardware performance is required for quantum networks to support useful applications. Rather than asking only whether a protocol is possible in principle, these studies aim to quantify how far existing hardware must improve, or what characteristics specific network components must attain, in order to make practical implementations feasible. Avis et al. studied the requirements for long-distance entanglement distribution on a real-world fiber grid between Delft and Eindhoven, using a repeater-based architecture and hardware-specific simulations to determine the improvements needed beyond current technology for verifiable blind quantum computation (VBQC) over a total distance of 227 km [4]. Ferreira da Silva et al. extended this type of requirements analysis to a 917 km real-world fiber route between Bonn and Berlin, using QKD performance targets to investigate how the required hardware improvements depend on both the number of processing-node repeaters and the target application when upgrading trusted-node infrastructure to a repeater chain [17]. In a complementary direction, Van Dam et al. analyzed the hardware improvements required to realize VBQC between a trapped-ion server and a measurement-only client separated by 50 km [39]. Most recently, Maiti et al. investigated teleportation requirements in an intercity architecture consisting of two metropolitan networks connected by a 450 km backbone, identifying the hardware improvements needed to reach a fidelity of $2/3$ using a more analytical approach [30].

These works have substantially advanced our understanding of near-term quantum-network requirements, but they remain focused on specific protocols and concrete archi-

tectural realizations. In particular, even in the intercity setting considered by Maiti et al., the analysis is centered on teleportation as a single benchmark. What is still missing is a comparative study that isolates the long-distance backbone itself as the object of interest and asks what entanglement quality and delivery rate such a backbone must provide in order to support different quantum-network applications. The present work adopts the intercity entanglement-distribution architecture of Maiti et al. and uses a simulation-based approach to investigate this gap. We now introduce this architecture.

The entanglement-distribution architecture

In the network architecture of this work, users are located within metropolitan areas and connect to their respective central metropolitan nodes, here referred to as border nodes, via short-range optical-fiber links of 20km. These border-nodes, in turn, are connected by a long-distance backbone that distributes entanglement between cities.

A central obstacle of the backbone is the distribution of entanglement over long distances. Unlike classical signals, quantum states cannot be amplified by copying, as this is forbidden by the no-cloning theorem [31]. Photon loss in optical fiber, which grows exponentially with distance, therefore severely limits the rate of direct-transmission entanglement distribution over long distances between remote parties. While several approaches have been proposed to overcome this challenge, including quantum repeater chains [38], satellite-based free-space links [43], and vacuum beam guides [22], each introduces distinct engineering trade-offs and none has yet established itself as the definitive solution for long-distance quantum communication.

Crucially, in this work we treat the backbone as an abstract resource, characterized entirely by two parameters. First, the rate R at which it delivers entangled states and second, the fidelity F of those states. This abstraction decouples the analysis from any particular backbone implementation and allows the investigation to focus on a question of broad practical relevance: what must the backbone deliver in order for useful quantum protocols to be executed between two parties in different metropolitan areas?

On the metropolitan side, the network is modeled using concrete physical hardware. The four nodes of the network can be divided into two types, namely processing nodes and simple clients. Processing nodes are equipped with quantum memories, qubit control, and readout capabilities, all capabilities required by the border-nodes. Clients, by contrast, are simpler end-user devices intended to provide broader access to quantum-network resources without requiring full quantum-processing capabilities. The border-nodes are realized as trapped-ion processors based on the system developed by Krutyanskiy et al. in Innsbruck [24, 25], while end-user devices are either processing nodes of the same type or lightweight measurement-only clients that interact with the network by rotating and measuring single photons. Accordingly, the simulated baseline hardware parameter set is based on values reported in current experiments, including parameters such as gate fidelities, emission durations, and coherence times. This combination of an abstract backbone with physically detailed metropolitan components provides a flexible framework for studying the interplay between local hardware quality and long-distance network requirements.

Three benchmark protocols

To evaluate how well our nodes and backbone perform, three quantum communication protocols are used as benchmarks: quantum key distribution [1, 7], the CHSH game [20, 11], and verifiable blind quantum computation [28, 10]. These protocols were selected not only because they represent distinct and practically relevant applications of a quantum network, but also because they span all possible combinations of end-node types available in the model. QKD can be implemented with two measurement-only clients, the CHSH game considered here requires two quantum processing nodes, and VBQC involves one measurement-only client and one server which is also a quantum processing nodes. Together, these three protocols test all end-node configurations considered in the model and provide a representative set of practically relevant applications.

Hardware parameter improvement framework

At their current performance levels, the baseline trapped-ion hardware parameters from the Innsbruck experiment are not yet sufficient to execute any of the three protocols, even assuming a noiseless and high-rate backbone. To study the conditions under which successful protocol performance becomes achievable, the simulated hardware parameters are systematically improved beyond their baseline using a single global improvement factor k , applied through probabilities of no imperfection as introduced by Avis et al. [4]. This factor acts uniformly on all relevant parameters through a mapping that accounts for the different physical domains of each quantity, allowing coherence times, gate fidelities, and photonic efficiencies to be improved on a common footing. While this uniform scaling does not reflect the unpredictable pace of real experimental progress, it provides a tractable way to explore how the backbone requirements evolve as the metropolitan hardware matures.

1.1 Research Questions

With this setup in place, the thesis addresses four research questions that progressively relax the assumptions on the backbone.

The investigation begins by assuming a perfect backbone, one that delivers entanglement with unit fidelity at an essentially infinite rate¹, in order to isolate the limitations imposed by the metropolitan hardware alone. This motivates the first question:

RQ1: *What is the minimum hardware improvement factor required to implement each protocol assuming a perfect backbone?*

This establishes the regime of improvement factors in which a meaningful investigation of backbone requirements is possible at all. Any hardware configuration that fails under a perfect backbone will necessarily also fail under a realistic one.

Next, the backbone fidelity is allowed to be imperfect while the rate remains infinite, isolating the effect of backbone quality:

¹That is, the backbone generates entanglement at such rates that its contribution to the end-to-end entanglement time is completely negligible and will therefore have no effect on the performance of the three protocols

RQ2: *For a given improved parameter set, what is the minimum required fidelity of the backbone to implement each protocol?*

The complementary question then restores perfect fidelity while reducing the backbone rate to finite values:

RQ3: *For a given improved parameter set, what is the minimum required rate of the backbone to implement each protocol?*

Like the first question, these two questions further restrict the relevant parameter space. If a protocol already fails with perfect backbone fidelity, then lowering the fidelity while keeping the same rate cannot restore feasibility. Likewise, if it fails when the backbone generates entanglement effectively instantaneously, it will also fail at any finite generation rate if we keep the fidelity constant. Together, these two limiting cases define necessary limits for the full two-parameter analysis considered in the final question.

Finally, both parameters are relaxed simultaneously to map out the full feasible region:

RQ4: *For a given improved parameter set, in what region of the backbone (R, F) space can each of the three protocols be successfully implemented?*

Contributions and outline

This thesis makes the following contributions. First, it presents a simulation-based analysis of three qualitatively distinct quantum communication protocols on a four-node trapped-ion network, implemented using the NetSquid discrete-event simulation platform [15]. The corresponding simulation code is made publicly available online [32]. Second, it identifies the minimum hardware improvement factors required for each protocol to become feasible, establishing the boundary below which no backbone can compensate for metropolitan hardware limitations. Third, it maps out the minimum backbone fidelity and rate requirements as functions of the hardware improvement factor, revealing how the demands on the backbone evolve as the local hardware improves. Finally, the combined analysis produces a two-dimensional characterization of the feasible backbone parameter space for each protocol, providing a practical reference for assessing whether a given backbone technology can support specific quantum network applications. We will further use this characterization to compare the backbone requirements across the different applications.

The remainder of this thesis is structured as follows. Chapter 2 introduces the theoretical background, covering the relevant quantum information/computing primitives and the physical systems that make up the network. Chapter 3 describes the network setup, simulation framework, hardware models, and protocol implementations in detail. Chapter 4 presents and discusses the results, organized according to the four research questions and discusses the limitations of the framework. Chapter 5 concludes with a summary of the findings and directions for future work.

Chapter 2

Background

This chapter introduces the background required for the remainder of this thesis. Section 2.1 first presents the theoretical foundations and application context, including the no-cloning theorem, teleportation, remote state preparation, quantum key distribution, the CHSH game, and blind quantum computing. Section 2.2 then turns to the physical network setting by introducing the architecture of a quantum internet considered here, the roles of different node types and introduces the hardware platforms considered in this work, namely trapped-ion quantum nodes and measurement-only clients. Together, these sections establish both the conceptual and physical building blocks of the network model studied in the following chapters.

2.1 Foundation

2.1.1 The no-cloning theorem

A fundamental property of quantum mechanics is that an unknown quantum state cannot be copied perfectly. This statement is formalized by the no-cloning theorem, which states that there exists no physical operation capable of producing identical copies of an arbitrary quantum state [31]. This has profound consequences for quantum communication: classical amplification methods, which rely on copying signals, simply do not work, making the transmission of quantum information over lossy channels a fundamental challenge. At the same time, the impossibility of copying unknown quantum states is also the foundation of many quantum communication protocols.

2.1.2 Teleportation

This section introduces quantum teleportation, a protocol that allows the transfer of any qubit state $|\psi\rangle$ from one party to another spatially separated party, requiring a shared entangled pair and classical communication. Teleportation is one of the most important tools for transmitting quantum information, as it directly addresses the challenge posed by the no-cloning theorem: sending specific qubits over conventional optical fiber carries a high chance of losing them, and lost qubits cannot be recovered through copying. It can therefore

2. BACKGROUND

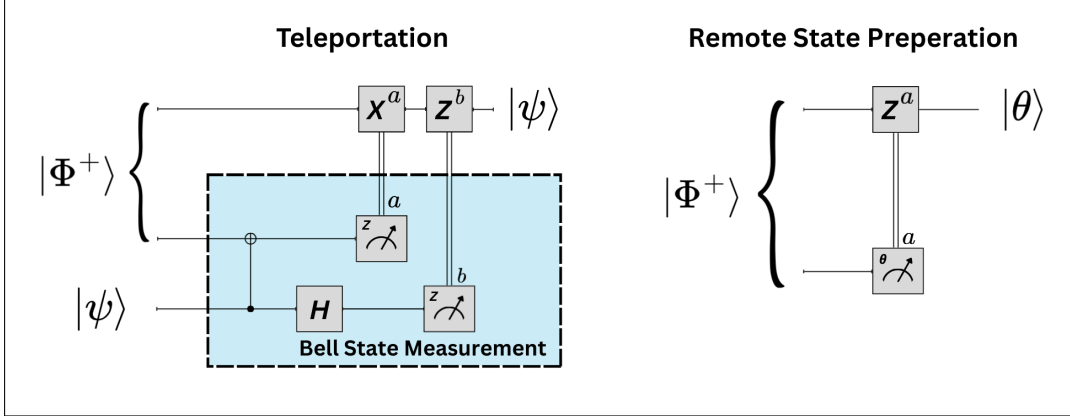


Figure 2.1: Circuit diagrams for quantum teleportation (left) and remote state preparation (RSP) (right). In both protocols, two parties exploit shared entanglement and classical communication to transfer quantum information without physically transmitting the quantum system. Additionally, the teleportation protocol highlights a deterministic Bell-state measurement, implemented by a CNOT gate followed by a Hadamard gate on the control qubit.

be preferable to first generate entanglement between two parties and then teleport the qubit [8].

The protocol works as follows. Two spatially separated parties, Alice and Bob, share any one of the four maximally entangled Bell states:

$$|\Phi^\pm\rangle = \frac{1}{\sqrt{2}} (|00\rangle \pm |11\rangle) \quad (2.1)$$

$$|\Psi^\pm\rangle = \frac{1}{\sqrt{2}} (|01\rangle \pm |10\rangle). \quad (2.2)$$

In addition, Alice also possesses a possibly unknown quantum state $|\psi\rangle$ that she wants to share with Bob without sending it directly. She performs a Bell-state measurement, i.e., a measurement in the Bell basis (see Equation 2.1 and 2.2), on the to-be-shared qubit and her half of the entangled pair. Immediately after the measurement, Bob's qubit collapses into one of four states: $|\psi\rangle$, $X|\psi\rangle$, $Z|\psi\rangle$, or $XZ|\psi\rangle$, depending on Alice's measurement outcome. Once Alice communicates her outcome, Bob can apply the corresponding correction gates to recover the original state. Classical communication is essential for this step because, without it, Bob can infer nothing about the teleported state. This ensures that the protocol does not violate causality. These steps are visualized as a circuit diagram in Figure 2.1.

Quantum teleportation is also central to quantum repeaters, where entanglement is first generated over shorter elementary links and then extended across the full network. If the to-be teleported qubit is already entangled with a third party, teleporting it creates an entangled link between the third party and the receiver of the teleportation. This process is known as entanglement swapping, and allows two clients who each share entanglement only with a central party to become entangled with each other.

2.1.3 Remote State Preparation

While quantum teleportation enables the transfer of an unknown quantum state using shared entanglement and classical communication, it requires the sender to physically prepare and manipulate the state to be transmitted. In many practical scenarios, this requirement can be experimentally demanding, as it involves state preparation and Bell-state measurements.

Remote State Preparation (RSP) provides a situational alternative, in which the sender needs to possess complete classical knowledge of the target state. Exploiting this knowledge, the sender can secretly prepare the desired state at a remote location using shared entanglement, a measurement and classical communication, without locally preparing the quantum state to be transmitted [9]. This allows, for example, remote preparation whereby the sender simply measures an incoming entangled photon in a specific basis.

To explain this phenomenon, we restrict our attention to equatorial states, which lie on the equator of the Bloch sphere and take the general form:

$$|\theta\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle). \quad (2.3)$$

The key insight behind RSP is the rotational invariance of the singlet Bell state, which allows it to be expressed in an antisymmetric form with respect to any orthonormal single-qubit basis. In particular, for any basis including the equatorial basis $\{|\theta\rangle, |\theta^\perp\rangle\}$, we can decompose the singlet state as:

$$|\Phi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle) = \frac{1}{\sqrt{2}}(|\theta\theta^\perp\rangle - |\theta^\perp\theta\rangle) \quad (2.4)$$

A measurement of one qubit in this basis then collapses the remote qubit into one of the two basis states: an outcome of $|\theta^\perp\rangle$ yields the desired state $|\theta\rangle$, while the opposite outcome produces $|\theta\rangle$, which can be corrected to $|\theta\rangle$ by applying a Z operation. Combining this principle with the general ability to convert any Bell state to another via local X or Z operations means that we can deterministically and remotely prepare any equatorial state from any known Bell state, by correcting according to the measurement outcome and the identity of the shared Bell state. It should be noted, however, that we cannot securely and deterministically remotely prepare an arbitrary state $|\psi\rangle$, as there is no unitary mapping from $|\psi^\perp\rangle$ to $|\psi\rangle$ in general. We note that within a protocol that nominally prepares equatorial states, the client may also choose to measure in the computational basis, thereby secretly preparing either $|0\rangle$ or $|1\rangle$ at the remote site. While the client cannot deterministically select which of the two states is prepared, this is sufficient for protocols in which either outcome is acceptable.

This protocol is important to this thesis as it allows for a cheaper and more technologically feasible method of secretly preparing remote states. With it, even a simple setup using a single-photon measurement device has access to some of the more advanced schemes of quantum cryptography requiring secret states.

2. BACKGROUND

Table 2.1: Examples of BB84 rounds with ideal correlations (no errors)

| Round | Basis _A | Bit _A | Basis _B | Bit _B |
|-------|--------------------|------------------|--------------------|------------------|
| 1 | Z | 0 | X | 1 |
| 2 | X | 1 | X | 1 |
| 3 | Z | 0 | Z | 0 |
| 4 | X | 0 | Z | 1 |

Table 2.2: Examples of BB84 rounds with a bit error

| Round | Basis _A | Bit _A | Basis _B | Bit _B |
|-------|--------------------|------------------|--------------------|------------------|
| 1 | Z | 0 | X | 1 |
| 2 | X | 1 | X | 0 |
| 3 | Z | 0 | Z | 0 |
| 4 | X | 0 | Z | 1 |

2.1.4 Quantum key distribution

Quantum key distribution (QKD) is a method for establishing secret cryptographic keys between distant parties whose security relies solely on the laws of quantum mechanics rather than assumptions about computational hardness. This stands in contrast to most classical cryptographic schemes, whose security is based on the presumed difficulty of certain mathematical problems.

Interest in quantum-secure communication increased significantly after the discovery of Shor’s quantum algorithm for efficient integer factorization [37]. This result demonstrated that widely used public-key cryptography systems, such as RSA, could in principle be broken by sufficiently powerful quantum computers. More broadly, it highlighted the inherent vulnerability of cryptographic schemes whose security relies on assumptions about computational hardness. Consequently, cryptographic protocols whose security does not depend on computational assumptions became an important area of research.

Quantum key distribution addresses this challenge by enabling two parties to generate a shared secret key with security guaranteed by the laws of physics. The security of QKD protocols is rooted deeply in fundamental principles of quantum mechanics, such as the no-cloning theorem and the disturbance caused by measurement. For a more comprehensive review of QKD protocols, security assumptions, and practical implementations, see Ref. [35].

The first QKD protocol was introduced by Bennett and Brassard in 1984 and is commonly referred to as the BB84 protocol [7]. In this scheme, one party (Alice) randomly chooses a bit and a preparation basis, here either the Z or X basis, encodes the bit into the corresponding basis state, and sends the resulting qubit to another party (Bob). Bob then measures each qubit in a randomly chosen basis, again either Z or X . After the transmission, Alice and Bob publicly compare their chosen bases while keeping the measurement outcomes private. Events in which both parties used the same basis are retained, while the others are discarded. In the absence of noise or eavesdropping, the remaining measurement outcomes form perfectly correlated classical symbols, as illustrated in Table 2.1.

In realistic implementations, however, errors inevitably occur due to channel noise, imperfect devices, or potential adversarial interference. These errors lead to only partially correlated measurement outcomes, as illustrated in Table 2.2. To quantify the level of disagree-

ment, Alice and Bob publicly compare a subset of their raw key and use it to estimate the quantum bit error rate. Depending on the protocol, this error rate may be basis-dependent, yielding separate estimates such as E_Z and E_X . Based on these estimates, they then apply classical post-processing steps, including error correction and privacy amplification, to distill a shorter but secure final key from the imperfect raw data [42]. During this process, part of the raw key is sacrificed in order to correct mismatches and remove any information that may have leaked to an eavesdropper.

In general, the fraction of the raw key that survives post-processing is referred to as the secret fraction s which depends on the quantum bit error rate (E), which quantifies the fraction of disagreeing bits when Alice and Bob compare their measurement outcomes. Together with the raw key generation rate R , which concerns the rate of obtaining correlated or semi-correlated symbols, this defines the practical secret key rate (SKR):

$$SKR = sR. \quad (2.5)$$

Its precise form, however, depends on the protocol, implementation, and security assumptions. One important distinction is between the asymptotic and finite-key regimes. In the asymptotic case, which assumes arbitrarily long keys, statistical fluctuations can be neglected and the achievable secret key rate can be expressed using simple analytical formulas. Since the simulations in this thesis rely on a specific set of assumptions, the explicit expression used for the secret key rate is introduced in Section 3.3.2.

2.1.5 The CHSH game

We now turn to another setting in which quantum mechanics can provide an advantage over classical systems, namely the generation of instantaneous correlations between distant systems that cannot be reproduced classically. The CHSH game provides an operational way to quantify such non-classical correlations and forms the basis for the second benchmark protocol considered in this thesis.

To introduce the game, it is useful to first recall the Clauser-Horne-Shimony-Holt (CHSH) inequality, which compares the correlations achievable with classical shared information to those predicted by quantum mechanics [14]. Consider two spatially separated parties, Alice and Bob, who each choose between two measurement settings, labeled 0 and 1, and each obtain a binary outcome. From these outcomes one defines the CHSH correlator

$$S = \langle A_0 B_0 \rangle + \langle A_0 B_1 \rangle + \langle A_1 B_0 \rangle - \langle A_1 B_1 \rangle. \quad (2.6)$$

If the correlations can be explained classically, using only shared randomness and without communication between Alice and Bob after the measurement settings are chosen, then the CHSH parameter is bounded by [31]

$$|S| \leq 2. \quad (2.7)$$

2. BACKGROUND

Quantum mechanics predicts stronger correlations. For a suitable entangled state and an appropriate choice of measurements, the CHSH parameter can reach

$$|S| = 2\sqrt{2}, \quad (2.8)$$

which exceeds the classical limit [31]. Any observed value above 2 therefore certifies correlations that cannot be reproduced by a classical strategy of this type. In practice, this makes the CHSH inequality a useful tool for quantifying whether the correlations shared by Alice and Bob are genuinely non-classical and whether, in suitable coordination tasks, these correlations can give rise to a quantum advantage. For the purposes of this thesis, it is convenient to express these correlations through the CHSH game, which provides a more intuitive framework for the advantage gained from non-classical correlations.

The CHSH game is a coordination game involving a referee, Charlie, and two players, Alice and Bob. Charlie sends a random bit $x \in \{0, 1\}$ to Alice and a random bit $y \in \{0, 1\}$ to Bob. Without communicating with each other, Alice and Bob must return bits a and b . They win the game if

$$a \oplus b = xy. \quad (2.9)$$

In order to see how the best classical and quantum strategies perform, it is useful to first show the equivalence between this game and the CHSH correlator. Assuming a uniform distribution of the inputs (x, y) , the probability of winning the game can be written as

$$p_{\text{win}} = \frac{1}{4} \sum_{x,y,a,b \in \{0,1\}} \delta_{a \oplus b, xy} p(a, b | x, y), \quad (2.10)$$

where the Kronecker delta enforces the winning condition. Introducing the parity variable $c = a \oplus b$, the joint probabilities can be expressed in terms of the correlators $\langle A_x B_y \rangle$ as [11]

$$p(c | x, y) = \frac{1}{2} (1 + (-1)^c \langle A_x B_y \rangle). \quad (2.11)$$

Substituting this expression into the winning probability yields

$$p_{\text{win}} = \frac{1}{8} \sum_{x,y,c} \delta_{c,xy} [1 + (-1)^c \langle A_x B_y \rangle]. \quad (2.12)$$

Evaluating the sum gives

$$p_{\text{win}} = \frac{1}{8} (4 + \langle A_0 B_0 \rangle + \langle A_0 B_1 \rangle + \langle A_1 B_0 \rangle - \langle A_1 B_1 \rangle). \quad (2.13)$$

Recognizing the CHSH correlator, the winning probability can be written compactly as

$$p_{\text{win}} = \frac{1}{2} \left(1 + \frac{S}{4} \right). \quad (2.14)$$

It follows that any classical strategy must satisfy the classical CHSH bound $S \leq 2$, corresponding to a maximum winning probability of $p_{\text{win}} = 0.75$. This value can be achieved,

for instance, by the simple strategy in which Alice and Bob always output 0, regardless of the inputs they receive from the referee. Quantum strategies, on the other hand, can obtain $S = 2\sqrt{2}$ by using suitably chosen measurement settings on a maximally entangled Bell state, resulting in

$$p_{\text{win}} = \frac{1}{2} \left(1 + \frac{\sqrt{2}}{2} \right) \approx 0.8536. \quad (2.15)$$

The CHSH game therefore turns the presence of non-classical correlations into a directly measurable protocol-performance metric. The concrete implementation of the game on the simulated network, including the chosen measurement settings and network-specific details, is reserved for Section 3.3.3.

The CHSH game also illustrates that entanglement can provide an advantage in concrete coordination tasks by enabling correlations that cannot be reproduced classically. More generally, CHSH-type correlations arise in a broader class of coordination problems, sometimes referred to as quantum telepathy [19]. In such scenarios, spatially separated players must coordinate their responses to randomly generated inputs without exchanging classical information after receiving those inputs. The presence of shared entanglement allows the players to achieve correlations that cannot be reproduced by any classical strategy. For example, recent work has studied entanglement-assisted coordination in distributed routing and scheduling problems, where separated routers must make correlated decisions without real-time communication. In such latency-constrained settings, shared entanglement can improve operational performance compared to optimal communication-free classical strategies [16].

2.1.6 Blind quantum computation

While the protocols discussed so far use entanglement as a resource for communication and correlation tasks, entanglement can also serve as the foundation for computation itself. This section introduces blind quantum computation (BQC), the basis of the third and final benchmark protocol considered in this thesis. BQC enables a client with limited quantum capabilities to delegate a quantum computation to a remote quantum server while keeping the computation private. In such a setting, the server performs the required quantum operations, while the client ensures that its input, algorithm, and output remain hidden [10]. This paradigm is expected to become increasingly relevant as large-scale quantum processors may one day be accessible through remote services.

To understand the BQC protocol considered in this thesis, it is useful to first introduce measurement-based quantum computing (MBQC), the computational model on which the protocol is built. In this model, computation is performed not by applying a sequence of unitary gates, but by measuring individual qubits of a fixed entangled resource state. The resource state is prepared once, and the entire computation is then driven by a pattern of adaptive single-qubit measurements together with classical feedforward of earlier outcomes. Because the entanglement in the resource state is consumed irreversibly by each measurement, this approach is sometimes referred to as a one-way quantum computer [34].

2. BACKGROUND

The entangled resource states used in MBQC are constructed from an underlying undirected graph $G = (V, E)$, where each vertex $v \in V$ corresponds to a qubit and each edge $(u, v) \in E$ prescribes an entangling operation. The preparation procedure is as follows: every qubit is initialized in the $|+\rangle$ state, and a controlled-Z (CZ) gate is applied between every pair of qubits connected by an edge. The resulting state is called a graph state [21], and the structure of the underlying graph determines the class of computations that can be implemented on it. For instance, the five-qubit linear graph state shown in Figure 2.2b supports universal single-qubit rotations, while the brickwork state of Figure 2.2a is a universal resource for arbitrary quantum computations, provided the graph is sufficiently large.

Computation on a graph state proceeds by performing single-qubit measurements in the equatorial basis $\{|\theta\rangle, |\theta^\perp\rangle\}$, introduced in Equation 2.3, where the angle θ is chosen by the party performing the computation. To build intuition, consider the five-qubit linear graph state, in which qubits are arranged in a chain connected by CZ gates. Measuring the first qubit at angle θ_1 implements a rotation on the logical information encoded in the chain, which is then passed to the remaining unmeasured qubits. Successive measurements at angles θ_2 , θ_3 , and θ_4 further transform this information, and the final qubit holds the output of the computation. By choosing the measurement angles appropriately, an arbitrary single-qubit unitary can be decomposed into this sequence of measurements on the linear graph.

Because measurement is inherently probabilistic, each measurement produces one of two possible outcomes. Depending on the outcome obtained, the state of the remaining qubits may differ from the intended one by a Pauli operator, referred to as a byproduct. Concretely, after measuring qubit i , the logical state carried by the unmeasured qubits is either the desired state or one related to it by an X or Z correction. A key insight is that these byproduct operators do not need to be corrected by applying physical gates. Instead, they can be absorbed into the measurement angle of a subsequent qubit using the relations [18]

$$M_i^\theta X_i = M_i^{-\theta}, \quad (2.16)$$

$$M_i^\theta Z_i = M_i^{\theta-\pi}. \quad (2.17)$$

with M_i^θ denoting a measurement of qubit i with angle θ in the equatorial plane of the Bloch sphere. An X byproduct is compensated by negating the measurement angle, while a Z byproduct is compensated by shifting it by π . This means that the measurement angle for each qubit depends on the outcomes of all previously measured qubits. The computation is therefore an inherently sequential and adaptive process in which one measures a qubit, records the outcome, updates the next measurement angle accordingly, and proceeds. A complete and formal treatment of how byproduct operators propagate through arbitrary graph states and how measurement angles must be adapted is provided by the measurement calculus [18].

The computational power of MBQC depends on the structure of the underlying graph. For general-purpose quantum computation, a universal resource state is required. The brickwork state shown in Figure 2.2a is one example of such a resource state, on which arbitrary unitary transformations can be implemented using only single-qubit measurements and classical feedforward, provided that the graph state is sufficiently large. In this thesis, however,




Figure 2.2: Graph states for measurement-based quantum computing. Edges symbolize Controlled-Z operations and circles are qubits. (a) A brickwork state with $n = 6$ rows and $m = 13$ columns, illustrating how larger graph states enable the implementation of more complex unitary operations. (b) A five-qubit linear graph state, which supports universal single-qubit rotations. This is the graph state investigated for blind quantum computation in this thesis.

the focus is not on universal computation, but on a five-qubit linear graph that is sufficient to benchmark the verified BQC protocol considered later.

The relevance of MBQC for BQC is that it naturally separates the party choosing the measurement instructions from the party physically preparing and measuring the graph state. A foundational protocol for blind quantum computation was introduced in the universal blind quantum computation (UBQC) protocol [10]. In this scheme, a client prepares single qubits with randomly chosen phases and sends them to a server, which constructs the entangled resource state and performs the measurements required for the computation. The client then instructs the server which measurement bases to use, while adapting these instructions based on previously reported measurement outcomes. Randomization of the prepared states and measurement angles ensures that the server cannot infer information about the underlying computation beyond trivial properties such as the size of the resource state.

Crucially, this framework can also be extended to verification. In verified blind quantum computation, the client does not only hide the computation, but also aims to detect whether the server deviates from the prescribed protocol. One way to achieve this is through the insertion of trap qubits, whose outcomes are known to the client and can therefore be used as tests. The original UBQC-based verification schemes, however, are highly sensitive to noise, limiting their practical applicability on near-term hardware. We therefore use a noise-robust verified blind quantum computation scheme that will be presented in Section 3.3.4 [28].

For the protocol used later, an important structural property is the coloring of the graph state. A k -coloring of a graph $G = (V, E)$ is a partition of the vertex set into k disjoint subsets

V_1, V_2, \dots, V_k satisfying

$$\bigcup_{i=1}^k V_i = V, \quad \text{and} \quad \forall i \in [k], \forall v \in V_i: N_G(v) \cap V_i = \emptyset, \quad (2.18)$$

where $N_G(v)$ denotes the set of neighbors of v in G [28]. In other words, no two adjacent vertices share the same color. Suppose that all qubits belonging to every color except one are prepared in computational basis states, $|0\rangle$ or $|1\rangle$, rather than in the $|+\rangle$ state. Since the CZ gate acts trivially when either of its input qubits is in the $|0\rangle$ state, and applies only a phase when the input is $|1\rangle$, the entangling operations between a computational-basis qubit and its neighbors reduce to at most local phase gates. As a result, the qubits of the remaining color become effectively disconnected from the rest of the graph: they remain in known pure states, up to Z corrections determined by the computational-basis states of their neighbors.

This mechanism is central to the verified BQC protocol considered in this thesis. By preparing certain qubits in computational-basis states, the client can isolate trap qubits from the rest of the graph and predict their measurement outcomes. These traps can then be used to test whether the server behaved honestly. The concrete noise-robust test-round protocol, specialized to the five-qubit linear graph, is described in Section 3.3.4.

2.2 The network architecture

This section connects the protocol-level background introduced above to the physical architecture considered in this thesis by introducing the relevant node types, their physical implementations, and the heralded entanglement-generation mechanism used between processing nodes in the same metropolitan area.

2.2.1 Quantum internet architecture

The vision of a quantum internet is to extend the capabilities of today’s classical internet by enabling the transmission and processing of quantum information between distant locations. Such a network would allow quantum devices to use and distribute entanglement, enabling applications such as quantum key distribution, non-local coordination strategies, and secure access to remote quantum computers.

At an abstract level, a quantum internet is built out of quantum nodes connected through quantum channels, allowing entanglement to be generated and distributed across the network. Nodes may serve different roles depending on their capabilities, ranging from simple end-user devices to more advanced processing nodes that facilitate long-distance communication [41].

In the remainder of this chapter we move from this abstract network description to the physical systems considered in this work. We first introduce the relevant types of quantum nodes and their respective platforms before introducing an entanglement generation scheme.

2.2.2 Network nodes

Within a quantum network, nodes can fulfill different roles depending on their position and physical capabilities. A common distinction is between end nodes, which act as users of the network, and intermediary nodes, which assist in the distribution of entanglement across the network. While intermediary nodes typically participate in entanglement generation and entanglement swapping, end nodes consume the resulting quantum resources in order to implement specific applications.

We focus on the quantum-network nodes relevant to this work, which can be divided into processing nodes and measurement-only (MO) nodes. Processing nodes represent the more capable elements of the network. They are responsible for generating, storing, and processing entanglement. These nodes support a limited but sufficient set of operations, including single-qubit rotations, two-qubit entangling gates, qubit readout, and the emission of photons entangled with the internal qubit state. MO nodes are simpler architectures capable of rotating and measuring individual incoming photonic qubits. In the following, we discuss two concrete physical realizations of these nodes: trapped-ion processing nodes and photonic measurement-only clients.

Ion traps

As a concrete physical realization of a processing node, we consider trapped ions. Trapped-ion systems are among the leading physical platforms for quantum technologies and are regarded as strong candidates for quantum repeaters and quantum information processing due to their long coherence times and high-fidelity control operations [36].

Here, we focus on radio-frequency (RF) linear Paul traps, in which ions are confined by an oscillating RF electric field in two spatial dimensions and by a static electric field in the remaining dimension, forming a one-dimensional ion chain. The dynamics of such systems can be described by the Mathieu equations, although a full derivation is beyond the scope of this thesis and a comprehensive treatment can be found in [27]. After loading, the ions are typically prepared using Doppler and sideband cooling. These are essential experimental steps, but they will not be discussed in detail here. For a broader review, see [12].

In this work, we consider the type of trapped-ion system developed by Krutyanskiy et al. [24], based on trapped $^{40}\text{Ca}^+$ ions in a linear Paul trap [12]. Only the capabilities




Figure 2.3: Recompilation of the CZ gate in terms of the trapped-ion native gate set, consisting of single-qubit rotations and the Mølmer-Sørensen gate. Traced from [39].

2. BACKGROUND

most relevant to this thesis are summarized here and a more detailed technical description is given in the appendix of [24]. The qubit is encoded in the internal electronic level structure of the ion, specifically in the electronic ground state and a metastable D -state connected via a narrow optical transition. This long-lived encoding enables coherent qubit manipulation on experimentally relevant timescales.

Lasers are used to manipulate these internal states, allowing individual addressing of ions and the implementation of single-qubit operations. In addition to their internal levels, the ions possess shared quantized motional modes arising from the trapping potential. These collective modes act as a quantum bus and enable entangling operations between ions within the same trap. In particular, the Mølmer–Sørensen gate uses a bichromatic laser interaction to couple the internal states to a shared motional mode, thereby generating an effective two-qubit interaction [12]. Together with single-qubit rotations, this provides a universal gate set for local quantum processing. This universality is illustrated by the recompilation of the CZ gate shown in Figure 2.3, since the CZ gate together with arbitrary single-qubit rotations forms a standard universal gate set.

Trapped ions additionally support high-fidelity qubit readout and can generate entanglement between an internal ion state and an emitted photon by correlating the ion state with the polarization of the emitted photon. Since these photons are not naturally emitted at telecom wavelengths, they can be converted to the telecom regime before being coupled into optical fibers, thereby significantly reducing transmission losses over long distances [24]. Entanglement between distant ions can then be established using the heralded entanglement generation scheme that will be introduced in Section 2.2.3.

High-fidelity coherent control, long-lived qubit memories, deterministic readout, and photon-mediated remote entanglement generation make them strong candidates for quantum network architectures.

Measurement-only clients

While trapped-ion systems provide powerful processing nodes, they are also complex devices that are challenging to build, operate, and scale to large numbers of end users. This motivates the development of lightweight client devices with only minimal quantum capabilities, while still allowing users to benefit from quantum network resources.

In this work, we therefore also consider measurement-only (MO) clients. In this model, the client does not generate, store, or process qubits internally, but instead measures and rotates photonic qubits received from a remote quantum node. Despite this restricted functionality, such a client can still access non-classical network resources by measuring photons that are entangled with qubits held by the server, thereby influencing the remote quantum state through remote state preparation.

A possible implementation of such a device measures the polarization state of incoming photons using a combination of a half-wave plate (H), quarter-wave plates (Q), a polarizing beam splitter (PBS), and single-photon detectors (SPD). The waveplates implement unitary rotations on the polarization qubit, while the PBS followed by the detectors performs a measurement in the computational basis. By adjusting the waveplate settings, the client can

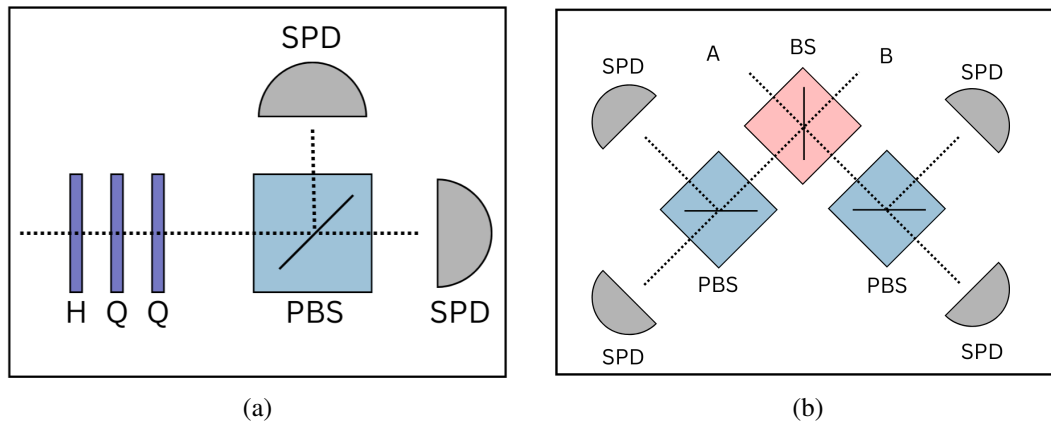


Figure 2.4: Optical components used in the network. Dotted lines indicate photon paths. (a) Measurement-only client capable of measuring an incoming photon in an arbitrary basis using a half- and quarter-wave plate. (b) Optical Bell-state measurement for polarization-encoded photons. A beam splitter interferes the photons from Alice and Bob, after which two polarizing beam splitters direct the outputs to four single-photon detectors.

therefore measure incoming photons in arbitrary single-photon bases. The client architecture considered here was proposed in [39] and is shown schematically in Figure 2.4a.

MO nodes thus represent simplified client devices that interact with the network solely through the manipulation and measurement of single photons emitted by a quantum node. Although their capabilities are limited compared to processing nodes, they are sufficient for several relevant protocols and provide a technologically lightweight interface to the network.

2.2.3 Heralded entanglement generation

Entanglement between distant quantum nodes can be generated probabilistically by interfering photons emitted from remote quantum memories. In heralded entanglement generation protocols, each node emits a photon whose photonic degree of freedom is entangled with the internal state of the emitter. The photons are transmitted through optical fibers to a midpoint station, where they interfere on a beam splitter and are subsequently measured, thereby probabilistically swapping the entanglement. Certain detection events therefore project the remote quantum memories onto an entangled state. Crucially, successful entanglement generation is heralded by a classical signal from the measurement station indicating that the required detection pattern occurred. Assuming the nodes wait for the classical notification of each attempt before initiating the next one, the repetition rate of the protocol is limited by the classical communication time between the nodes and the midpoint station, which is on the order of L/c for a link of length L .

When considering two trapped-ion processors connected through a midpoint station, a suitable approach is to use a heralded entanglement-generation scheme based on photon interference at that station [4], following the class of protocols originally proposed by Barrett

2. BACKGROUND

and Kok [6]. Each node emits a photon whose polarization is maximally entangled with the state of the emitter, resulting in the ion–photon state $\frac{1}{\sqrt{2}}(|0H\rangle + |1V\rangle)$. The photons from the two nodes are interfered on a non-polarizing beam splitter at the midpoint station and subsequently measured using polarization-resolving single-photon detectors. A successful heralding event corresponds to the detection of one horizontally polarized and one vertically polarized photon at different detectors, projecting the remote memories onto one of the Bell states

$$|\Psi^\pm\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle) \quad (2.19)$$

depending on the detection combination.

Chapter 3

Methods

This chapter presents the methodological framework used throughout this thesis. Section 3.1 first defines the physical and architectural setting of the problem by introducing the four-node network model and the assumptions made about the metropolitan links and the abstract backbone. Section 3.2 then describes how this setup is translated into a Net-Squid simulation, including the hardware and noise models, the baseline parameter set, and the global hardware-improvement framework used in later analyses. Finally, Section 3.3 details the implementation of the three benchmark protocols on this simulated network, starting from the shared features and then turning to the protocol-specific procedures and performance quantities used to assess QKD, the CHSH game, and VBQC. These sections provide the basis for the analyses in Chapter 4.

3.1 Network setup

Having introduced the fundamental building blocks of quantum networks in the previous chapter, we now describe the network architecture considered in this work. In this work, we consider an intercity quantum network in which users located in different metropolitan areas access quantum communication services through local infrastructure nodes. These so called border-nodes connect local clients to a long-distance backbone that distributes entanglement between cities. For instance, one may consider users located in cities such as Amsterdam and Paris or Amsterdam and Berlin. In such scenarios, entanglement must be distributed across distances that significantly exceed the scale of a single metropolitan area.

At the metropolitan scale, direct optical fiber connections provide a realistic means of linking clients to nearby infrastructure nodes. Fiber losses over these distances remain manageable. For example, a 20 km fiber link results in approximately 60% photon loss, which still allows practical heralded entanglement generation. Consequently, in the model considered here, clients within a metropolitan area are assumed to be connected to a border-node via direct fiber links and processing node clients generate entanglement via a midpoint heralding station as seen in Figure 3.1.

In contrast, the physical realization of the long-distance quantum backbone connecting metropolitan areas remains an open question. Several architectures have been proposed,

3. METHODS

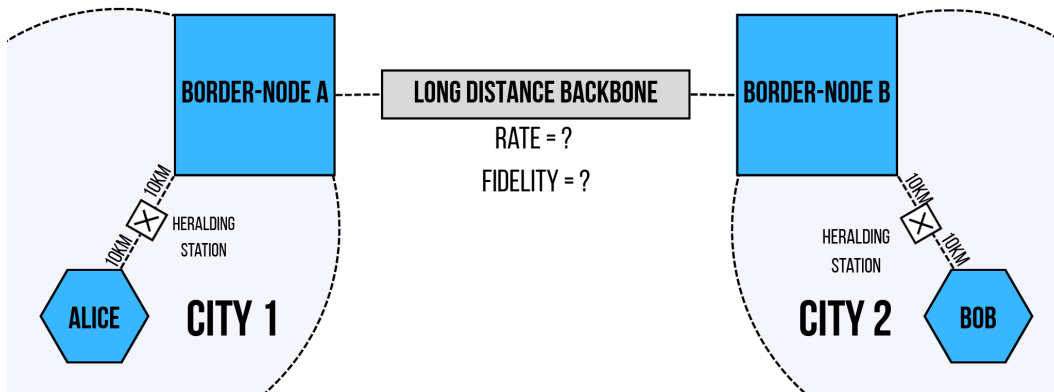


Figure 3.1: Overview of the network architecture simulated in this work. Two parties, Alice and Bob, are each located in a separate metropolitan area and connect to a local border-node. For quantum end-nodes, the metropolitan link uses a midpoint heralding station with 10 km fiber on each side; for measurement-only clients, it consists of a direct 20 km optical fiber connection. The two border-nodes are connected via a long-distance backbone that delivers entanglement at a rate R and with fidelity F .

including repeater chains, multiplexed photonic links, and satellite-based quantum communication systems. Each of these approaches introduces different implementation challenges and performance characteristics. However, from the perspective of applications running on the network, these architectures ultimately provide the same resource: entangled states distributed between distant nodes with some generation rate and fidelity.

For this reason, the intercity backbone is treated as an abstract resource in our model. Rather than specifying a particular physical implementation, it is characterized only by the fidelity F of the entangled states delivered between metropolitan network nodes and by their average delivery rate R . To capture the stochastic nature of entanglement generation, delivery times are modeled as a geometric process with per-attempt success probability $p_{succ} = 0.01$, where the attempt duration is chosen such that the corresponding mean delivery rate equals the desired rate R . This does not mean that any actual backbone must have an inherent success probability of 1%. It is a convenient reference choice, and moderate changes in the success probability can be compensated by adjusting the attempt duration accordingly while preserving the same average delivery rate.

This abstraction allows the analysis to remain agnostic to the specific technology used to realize the backbone, while focusing instead on the end-to-end performance available to network users. By contrast, the metropolitan part of the network, including the four nodes and the entanglement-generation processes between them, is modeled using concrete physical systems. In particular, we consider ion-trap processing nodes and measurement-only clients connected via optical fiber. Communication between a processing node and an MO client is modeled as a direct optical-fiber link, whereas entanglement between two processing nodes within the same metropolitan area is generated through a midpoint heralding station connected via optical fiber using the protocol introduced in Section 2.2.3.

The resulting network architecture considered in this work is illustrated in Fig. 3.1. The setup consists of two metropolitan networks connected through the abstract backbone described above. Each network hosts a quantum border-node that serves as the interface between the intercity connection and local users. Within each metropolitan area, clients connect to the local border node, which mediates the generation and distribution of entanglement between network participants.

Several assumptions are made regarding the network configuration considered in this work. First, the distance between clients and their respective border nodes is assumed to be 20 km. This distance is representative of metropolitan-scale connections and would, for example, allow hypothetical clients located in Delft to access a border node in Rotterdam. More generally, this range is sufficient to provide coverage across large European metropolises such as London or Paris.

Furthermore, the heralding station used for entanglement generation is assumed to be placed at the midpoint of the two communicating nodes. This location is, for two parties with identical parameters, optimal, as it minimizes the classical communication time. Although the midpoint placement is assumed throughout this work, the results obtained remain mostly valid for asymmetric configurations within the parameter regimes considered here. This will be discussed further in the analysis of the simulation results.

To evaluate the capabilities of the network architecture introduced above, we consider the three aforementioned quantum communication protocols as benchmarks: quantum key distribution (QKD), the CHSH game, and verified blind quantum computation (VBQC). These protocols represent distinct ways in which distributed entanglement can be used by network participants.

However, the protocols were not only chosen because of their respective uses, but they also place different requirements on the capabilities of the participating nodes. The CHSH protocol requires two quantum nodes, as both parties must perform measurements on locally stored qubits that share entanglement. In contrast, QKD can be implemented using two measurement-only clients. Blind quantum computation requires one measurement-only client and one quantum node acting as the computational server. These three protocols therefore span the possible combinations of client capabilities available in the model. Before we do so, we first need to take a look at our simulation framework.

3.2 Simulation framework

3.2.1 NetSquid

For the entirety of this simulation we use the Network Simulator for Quantum Information using Discrete events (NetSquid) [15]. NetSquid is a discrete-event simulation platform developed at QuTech for modeling scalable quantum networks and modular quantum computing systems. Its modular design allows individual physical components to be assembled into larger network simulations, while accurately tracking the effects of time-dependent noise such as qubit decoherence. This makes it the perfect choice for modeling our four-node setup. In addition, this project builds upon two NetSquid extension packages: `netsquid-netbuilder` [40], which provides a modular framework for constructing network configu-

rations, and netsquid-trappedions [2], which supplies the physical models for trapped-ion hardware.

3.2.2 Hardware modeling

This section outlines the noise models used to describe the physical hardware components in the simulation. Each imperfection is captured by a corresponding quantum channel, allowing the simulation to faithfully reproduce the effects of realistic noise on protocol performance. The noise models described below were already implemented in the aforementioned NetSquid packages and are used here as part of the underlying simulation framework.

Gate noise

Imperfect quantum gates are modeled using depolarizing channels. For a single-qubit gate, the noise is described by

$$\mathcal{D}_p^{(1)}(\rho) = (1-p)\rho + p \frac{I_2}{2}, \quad (3.1)$$

where $I_2/2$ is the single-qubit maximally mixed state and $p \in [0, 1]$ is the depolarizing probability. The fidelity of the output state with respect to the target pure state $|\psi\rangle$ is given by $F = 1 - p/2$, yielding the inverse relation $p = 2(1 - F)$.

For two-qubit gates, such as the Mølmer–Sørensen gate, the noise is modeled analogously by a two-qubit depolarizing channel

$$\mathcal{D}_p^{(2)}(\rho) = (1-p)\rho + p \frac{I_4}{4}, \quad (3.2)$$

where $I_4/4$ is the two-qubit maximally mixed state. In this case, the relationship between the depolarizing probability and the gate fidelity becomes $p = \frac{4}{3}(1 - F)$. Gate durations for single-qubit rotations, the Mølmer–Sørensen gate, and qubit readout are included in the simulation as fixed time costs that contribute to the total protocol duration and, consequently, to the amount of memory decoherence experienced by stored qubits.

Emission noise

The entanglement between an emitted photon and the trapped ion is also subject to noise. This takes the form of a Werner state by applying the two-qubit depolarizing channel 3.2 to the ideal ion-photon state $\rho = |\Phi^+\rangle\langle\Phi^+|$, where the depolarizing parameter is determined by the emission fidelity through the same relation as above.

Memory decoherence

The dominant source of memory noise in trapped-ion systems is dephasing caused by fluctuations in the ambient magnetic field. Since ions confined in the same trap are subject to the same field fluctuations, the resulting noise is correlated across qubits. This collective dephasing is modeled as a Gaussian decoherence channel of the form

$$\rho \rightarrow \int_{-\infty}^{\infty} K_r \rho K_r^\dagger p(r) dr, \quad (3.3)$$

with Kraus operators

$$K_r = \exp\left(-ir \frac{t}{T_{\text{coh}}} \sum_{j=1}^n \sigma_z^{(j)}\right) \quad (3.4)$$

and a Gaussian weight function

$$p(r) = \frac{1}{\sqrt{2\pi}} e^{-r^2/2}, \quad (3.5)$$

where t is the storage time, T_{coh} is the coherence time of the memory, n is the number of ions in the trap, and $\sigma_z^{(j)}$ denotes the Pauli-Z operator acting on ion j . A detailed derivation and discussion of this model can be found in [4].

Photonic parameters

In addition to the noise channels described above, several parameters govern the photonic components of the network. The emission duration determines the time required to generate an ion-photon entangled pair and therefore directly affects the rate at which entanglement generation can be attempted and therefore generated. The initial photon survival probability captures the combined effect of the photon collection efficiency, the frequency conversion efficiency from the ion emission wavelength to the telecom regime, and the detector efficiency. These parameters determine the overall probability that an emitted photon is successfully detected at the heralding station or MO-client per generation attempt, excluding attenuation losses.

Two further parameters characterize the quality of the photon interference at the mid-point beam splitter relevant for the entanglement generation scheme from Section 2.2.3. The visibility parameter accounts for imperfect photon indistinguishability in the heralding process. In the model implemented by Avis et al. [3] and used here, this imperfection enters the off-diagonal coherence terms through a factor proportional to \sqrt{V} , thereby reducing the fidelity of the generated entanglement. The parameter η_{penalty} is a multiplicative reduction of the heralding success probability that arises from restricting the accepted detection window in order to improve temporal overlap between the arriving photons. These two quantities are inherently linked: a narrower detection window improves the visibility at the cost of a larger penalty to the success probability, and vice versa. In this work, we do not optimize over this tradeoff but instead adopt a fixed operating point taken from [4]. Both parameters are discussed in greater detail therein. Finally, it should be noted that measurement-only clients are not modeled as a significant additional source of state infidelity. Since they only rotate and measure incoming photons, their imperfections mainly affect the probability of detecting a photon rather than the fidelity of the generated state. In addition, any realization of this measurement device needs to be carefully implemented, with basis independent detection efficiencies and without exploitable detector side channels. This assumption is important for the security of the protocols considered later, since imperfections in the detector module could otherwise be exploited by an adversary.

3. METHODS

Table 3.1: Baseline hardware parameters used throughout this work. Ion trap parameters are taken from the Innsbruck repeater node experiment [24], while photonic parameters are based on the entanglement distribution experiment reported in [25]. Derived quantities are indicated by footnotes.

| Parameter | Value |
|--|-----------------------|
| Ion Trap | |
| Coherence time (ns) [24] | 62,000,000 |
| Single-qubit gate fidelity [24] | 0.99 |
| MS gate fidelity [24] | 0.95 |
| Single-qubit rotation duration (ns) [24] | 12,000 |
| MS $\pi/2$ duration (ns) [24] | 107,000 |
| Measurement duration (ns) [10] | 100,000 |
| Photonic Parameters | |
| Initial photon survival probability ^a | 0.1245 |
| Emission fidelity [25] | 0.974 |
| Emission duration (ns) [25] | 666,666 |
| Dark count probability ^b | 8.75×10^{-6} |
| Visibility [25] | 0.89 |
| η_{penalty} [25] | 0.12 |

^a Computed as the product of the collection efficiency (0.53) [25], frequency conversion efficiency (0.25) [26], and detector efficiency (0.94) [13].

^b Estimated from a dark count rate of 500 Hz [13] and a detection window of $17.5 \mu\text{s}$ [25] via $p_{\text{dc}} = 1 - e^{-R\tau}$.

3.2.3 Baseline parameters

The simulation requires a consistent set of hardware parameters that together describe a complete network node. Rather than selecting individually optimized values for each component, we adopt a single parameter set derived from the system developed by Krutyanskiy et al. [24, 25]. This system was designed as an all-round quantum repeater node, integrating ion-photon entanglement generation, local two-qubit gates, and telecom-wavelength photon conversion within a single platform. While individual subsystems with superior performance exist, for instance qubits with significantly longer coherence times or higher gate fidelities, such systems are typically optimized for a specific task and do not necessarily offer the full set of capabilities required for the protocols considered here. The Innsbruck system therefore provides a realistic and self-consistent baseline from which improvements can be studied systematically. We also use the same platform for end users that require quantum processing capabilities, since it provides a well-rounded processor model with the same functional requirements as those user nodes.

The baseline parameters are summarized in Table 3.1. Not all values listed are directly measured quantities. The initial photon survival probability is a composite parameter obtained from the product of the emission probability, the frequency conversion efficiency, and the detector efficiency. Similarly, the dark count probability is estimated from the dark

count rate and the detection window duration, as detailed in the table footnotes. All remaining parameters are taken directly from the experimental references indicated.

In addition to the parameters listed in the table, two further quantities are fixed throughout this work. The fiber length between each client and its respective border node is set to $L = 20$ km, and the fiber attenuation is assumed to be 0.2 dB/km, consistent with standard telecom-wavelength optical fiber.

3.2.4 Probabilities of no-imperfection

The baseline parameters introduced in the previous section represent the current state of trapped-ion hardware. For the 4-node network configuration considered here, current hardware parameters are simply too noisy and slow to support successful implementation of any of the three protocols. To study the conditions under which positive protocol performance becomes achievable, it is therefore necessary to systematically improve the simulated hardware parameters beyond their current values.

A complication that arises when attempting to improve parameters uniformly is that they live on fundamentally different domains. Coherence time takes values in $[0, \infty)$, single qubit gate fidelities are physically meaningful within $[0.5, 1]$, and quantities such as the initial photon survival probability lie in $[0, 1]$. An improvement by a fixed percentage therefore has very different physical implications depending on the parameter in question, and may even be unphysical for parameters already close to their theoretical maximum.

To address this, we map each hardware baseline parameter b_i onto a common domain $[0, 1]$ by defining a probability of no-imperfection, denoted $p_{\text{NI}}(b_i)$ [4]. This quantity represents the probability that the corresponding hardware component introduces no error. A value of $p_{\text{NI}} = 0$ corresponds to the worst possible performance while $p_{\text{NI}} = 1$ corresponds to perfect hardware. The mappings for all relevant parameters are listed in Table 3.2.

For gate fidelities, these mappings follow directly from the depolarizing channels introduced in Section 3.2.2. The factors $1 - p$ in equations 3.1 and 3.2 already have the desired form and expressing them in terms of the fidelity F yields the relations listed in Table 3.2. For the remaining parameters, such as the photon survival probability, visibility, and η_{penalty} , the values already lie in $[0, 1]$ with the correct interpretation and are therefore used directly.

The mapping for the emission duration requires additional explanation as it is not present in the original work [4]. The emission duration determines the time required for a single entanglement generation attempt. In the network architecture considered here, backbone entanglement is always generated before the local client links. As a result, qubits that are already stored in memory continue to decohere during subsequent entanglement generation attempts on other links. The approximate decoherence accumulated from the emission time during a single attempt is then characterized by the probability of no-imperfection

$$p_{\text{NI}} = \exp\left(-\frac{t_e^2}{\tau^2}\right), \quad (3.6)$$

where t_e is the emission duration and τ is some coherence time. A detailed derivation of other mappings and their relation to the physical error mechanisms can be found in [4].

3. METHODS

Table 3.2: Probabilities of no-imperfection for all relevant parameters. The quantities T_{coh} and t_e denote the coherence time and emission duration respectively. We restrict attention to ion-trap and heralding-station parameters, as these are the only parameters that significantly affect performance.

| Parameter | p_{NI} |
|--|--|
| Coherence time | $\exp\left(-\frac{t^2}{T_{\text{coh}}^2}\right)$ |
| Emission fidelity F , MS gate fidelity F | $\frac{1}{3}(4F - 1)$ |
| Single-qubit gate fidelity F | $2F - 1$ |
| Emission time | $\exp\left(-\frac{t_e^2}{\tau^2}\right)$ |
| η_{penalty} , visibility V , photon survival probability p | $\eta_{\text{penalty}}, V, p$ |

Having defined the probability of no-imperfection for each parameter, we can now introduce the improvement procedure used throughout this work. A given parameter with baseline probability of no-imperfection $p_{\text{NI}}(b_i)$ is improved by a factor k by taking the k -th root:

$$p_{\text{NI}}(x_i) = (p_{\text{NI}}(b_i))^{1/k} \quad (3.7)$$

where b_i are the baseline parameters and x_i are then the improved parameters. For the coherence time mapping $p_{\text{NI}} = \exp(-t^2/T_{\text{coh}}^2)$, this is equivalent to multiplying the coherence time by \sqrt{k} . Analogously, for the emission duration the same operation corresponds to dividing the emission time by \sqrt{k} .

In this work, a selected subset of parameters is always improved by the same global factor k . Assigning independent improvement factors to individual parameters would in principle be possible, but would lead to a high-dimensional parameter space that is impractical to explore.

It should be noted that not all parameters listed in Table 3.1 are varied in this analysis and are therefore not mentioned in Table 3.2. Parameters with only a minor impact on protocol performance are kept fixed at their baseline values throughout. This allows the analysis to focus on those parameters whose improvement most strongly affects protocol performance. For gate durations in particular, this also reflects the practical reality that gains in gate fidelity and gate speed are typically not achieved simultaneously. In some cases, improving fidelity may even require slower operations.

3.3 Protocol implementation

3.3.1 General remarks

This section describes the features of the protocol execution pipeline that are shared across all three benchmark protocols. These shared elements arise from the network architecture

and the capabilities of the trapped-ion hardware, and are stated here once to avoid repetition in the protocol-specific sections that follow. For simplicity, we assume in the method descriptions that all successful entanglement generation events produce the state $|\Phi^+\rangle$. The same procedures can be adapted straightforwardly to any other Bell state by modifying the Pauli corrections applied in the subsequent protocol steps.

Entanglement generation order. In the network setup considered here, the backbone connecting the two metropolitan areas spans a significantly larger distance than the local links between clients and border nodes. As a consequence, the expected time to successfully generate entanglement over the backbone is substantially longer than over the metropolitan links. This motivates a fixed ordering of entanglement generation in which backbone entanglement is always established first. Only after a successful heralding event over the backbone do the two border nodes simultaneously initiate entanglement generation with their respective local clients. This ordering has a direct impact on protocol performance, as the qubits stored at the border nodes accumulate decoherence while waiting for the metropolitan entanglement attempts to succeed. Figure 3.2 showcases the entanglement generation order that is used throughout this work.

Ion trap symmetry. The trapped-ion nodes considered in this work contain two qubits that are physically symmetric. There is no distinction between dedicated memory and communication qubits, since either qubit can be used to generate entanglement with a remote node or participate in local gate operations. A practical consequence of this symmetry is that after a successful entanglement generation event, no additional swap operation is required to transfer the entangled state to a designated memory spot. The qubit that participated in the entanglement generation is simply used directly in subsequent protocol steps. While this simplifies the protocol execution, it is a feature specific to this hardware platform. Other physical implementations may require an explicit swap, which would introduce additional gate noise and processing time.

Sequential entanglement attempts. Each ion trap can only attempt entanglement generation on one link at a time, as it is difficult to simultaneously emit photons towards two different nodes. For the border nodes, which must establish entanglement both over the backbone and with a local client, this means the two links are handled sequentially. Combined with the ordering described above, the procedure at each border node is therefore: first participate in backbone entanglement generation, then, once the backbone link has been established, begin quantum communication with the end nodes. Although the two border nodes attempt their local links simultaneously, each individual node is engaged with only one link at any given time.

Native Bell state measurement. All three protocols require Bell state measurements at the border nodes, either for teleportation or for entanglement swapping. Luckily, we can perform them deterministically instead of relying on probabilistic optical Bell state measurements. In the trapped-ion platform, this operation is implemented using the native

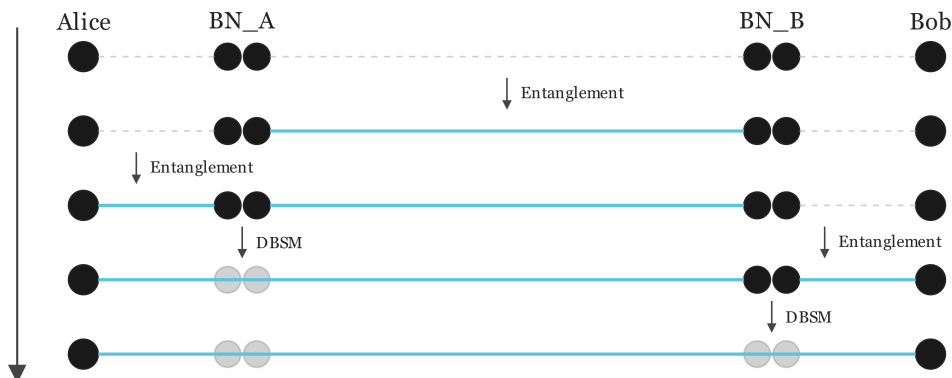


Figure 3.2: Order of entanglement generation for the four-node setup. In all protocols, entanglement along the backbone between the border nodes is established first. Only afterwards, will the metropolitan entanglement be generated between clients and border nodes. Note that Steps 3 and 4 therefore may occur in either order, depending on probabilistic outcomes. The corresponding Bordernodes then perform entanglement swaps For MO-clients, the generated entanglement is consumed immediately to remotely prepare the desired state.

gate set via a single-qubit Z rotation followed by one Mølmer-Sørensen gate, after which both qubits are measured. This decomposition requires only two noisy gate operations before readout, making it as efficient in terms of accumulated gate noise as the standard Bell state measurement. The outcome labeling of this native BSM differs from that of the standard CNOT+Hadamard circuit, but the four measurement outcomes can be mapped one-to-one onto the same four Bell state identifications. This remapping is applied consistently throughout all three protocols, ensuring that the teleportation corrections derived from each BSM outcome are interpreted correctly.

Generality. While the simulation framework accommodates other hardware platforms and parameter sets, the specific assumptions described above, namely the qubit symmetry within each trap, the one-sided emission constraint, and the native gate decomposition of the BSM, are particular to the trapped-ion system considered here. Extending the analysis to other physical implementations would require revisiting these assumptions and adapting the protocol execution pipeline accordingly.

3.3.2 Quantum key distribution

This section describes the implementation of quantum key distribution on the four-node network architecture. We first identify a suitable protocol and security analysis, then describe the protocol execution, and finally derive the secret key rate used to evaluate performance.

Protocol choice. Before investigating the implementation of QKD on the network architecture, a suitable protocol must be identified. The protocol must be compatible with two

MO-clients as end nodes, and we restrict the analysis to the asymptotic key regime in order to probe the theoretical maximum capacity of the setup.

The network architecture considered here closely resembles a quantum repeater chain, as the two border-nodes play the same functional role as repeater nodes. They both generate entanglement with adjacent parties and perform Bell-state measurements to extend the entangled link. This observation allows us to directly adopt the QKD analysis developed for repeater networks in [1]. In that work, Alice and Bob measure the state distributed over the network by applying local rotations and using single-photon detectors, while an adversary (Eve) is assumed to control the network that handles the entanglement generation. This is identical to the setup considered here, where the network architecture generates and distributes entanglement, and Alice and Bob as MO-clients measure incoming photons in arbitrary bases using waveplates and single-photon detectors. We restrict our attention to the BB84 scenario considered in [1], for which the asymptotic secret key fraction takes the standard form

$$s = 1 - h(E_Z) - h(E_X), \quad (3.8)$$

where $h(E)$ is the binary entropy function defined as

$$h(E) := -E \log_2 E - (1 - E) \log_2 (1 - E). \quad (3.9)$$

The complete expression for the secret key rate, incorporating the raw key rate and error correction, will be presented after the protocol execution has been described.

Protocol execution. The protocol proceeds as follows. First, entanglement is generated over the backbone between the two border nodes. Once successful, both border nodes simultaneously emit entangled photons towards their respective MO clients until arrival. Alice and Bob each measure the incoming photon in a uniformly randomly chosen basis, either X or Y , thereby remotely preparing a qubit at their respective border node in a state chosen to encode their raw key bit $b \in \{0, 1\}$. After both remote state preparations have succeeded, the two border nodes perform a deterministic Bell state measurements. This teleports the two RSP'd qubits to one border-node and then performs a parity measurement. The BSM outcomes, together with the known remote state preparation outcomes, allow Alice and Bob to determine whether their bits are correlated or anti-correlated (see Table 3.3). Rounds in which the bases do not match are discarded, and the remaining data is used as an estimate for the quantum bit error rate E .

Basis choice. In most formulations of BB84, the two measurement bases are typically taken to be Z and X . Here, we use the X and Y bases instead, as they are both equatorial states. Since X and Y also form a pair of mutually unbiased bases, this is equivalent to the standard BB84 protocol up to a relabeling of the qubit basis. All noise models employed in this work affect X and Y eigenstates identically. This means that in this work the error rates in both bases are identical ($E_X = E_Y = E$)¹. For simplicity, and to avoid additional

¹Choosing Z and X as the two bases would in principle lead to lower error rates on the Z -basis states, which are unaffected by dephasing. However, as can be seen in later results, memory decoherence plays a negligible role for the MO-client-based QKD setup considered here, making this distinction minor in practice.

3. METHODS

Table 3.3: Deterministic Bell-state measurement outcomes for X - and Y -basis input states. Alice and Bob's input states are mapped through a CNOT gate followed by a Hadamard on the first qubit.

| X basis | | | Y basis | | |
|--------------|---|----------|----------------|---|----------|
| Input | Before measurement | Outcomes | Input | Before measurement | Outcomes |
| $ ++\rangle$ | $\frac{1}{\sqrt{2}}(00\rangle + 01\rangle)$ | 00, 01 | $ +i+i\rangle$ | $\frac{1}{\sqrt{2}}(i 01\rangle + 10\rangle)$ | 01, 10 |
| $ --\rangle$ | $\frac{1}{\sqrt{2}}(00\rangle - 01\rangle)$ | 00, 01 | $ -i-i\rangle$ | $\frac{1}{\sqrt{2}}(-i 01\rangle + 10\rangle)$ | 01, 10 |
| $ +-\rangle$ | $\frac{1}{\sqrt{2}}(10\rangle - 11\rangle)$ | 10, 11 | $ +i-i\rangle$ | $\frac{1}{\sqrt{2}}(00\rangle - i 11\rangle)$ | 00, 11 |
| $ -+\rangle$ | $\frac{1}{\sqrt{2}}(10\rangle + 11\rangle)$ | 10, 11 | $ -i+i\rangle$ | $\frac{1}{\sqrt{2}}(00\rangle + i 11\rangle)$ | 00, 11 |

parameters that can be varied in our simulation, we implement symmetric basis choices for Alice and Bob, such that $p_X = p_Y = 1/2$. In many implementations of QKD, such as the aforementioned work in [1], the basis choice is typically biased towards the key-generating basis in order to maximize the raw key rate. In the strongly biased limit the fraction of useful rounds approaches unity, whereas under symmetric basis sampling only half of the rounds correspond to matching bases. Raw-key-rates with twice the magnitude are therefore theoretically possible.

Secret key rate. The secret key rate is obtained by combining a secret key fraction with a raw key rate. For the secret key fraction, we follow [5] and adapt it to the symmetric error rate, while introducing an error correction inefficiency factor f . The resulting expression is

$$s = 1 - [1 + f]h(E), \quad (3.10)$$

where $f = 1$ corresponds to the ideal case of perfect error correction. Throughout this work we use $f(E) = 1.16$, following the assumption made in [29]. As noted therein, this value is rather conservative and lower penalties are achievable with modern reconciliation protocols. The raw key rate is defined as

$$R_{\text{raw}} = \frac{N_{XX} + N_{YY}}{T_{\text{tot}}}, \quad (3.11)$$

where N_{XX} and N_{YY} denote the number of successful rounds for the XX and YY basis choices, respectively, and T_{tot} is the total simulated time. The secret key rate is then given by the product $SKR = s \cdot R_{\text{raw}}$.

Corrections. Each remotely prepared qubit is subject to Pauli corrections arising from the RSP measurement outcome and the BSM at the border nodes. A Z correction always flips the prepared bit, regardless of the chosen basis. An X correction flips the bit only for states prepared in one of the two bases and acts trivially on the other. Since Alice and Bob know both their own RSP outcomes and the announced BSM results, they can

independently determine whether their respective bits need to be flipped and apply these corrections classically. No physical operations on quantum states are required.

Analytical model for the perfect metropolitan hardware limit. To complement the simulation results, it is useful to derive an analytical expression for the secret key rate in the idealized limit where all metropolitan hardware parameters are perfect, i.e., the improvement factor $k \rightarrow \infty$. In this limit, the only remaining source of error is a possible imperfect fidelity F of the backbone entanglement used to teleport the RSP'd qubits. The QBER is then determined exclusively by the depolarizing noise on the teleported state, giving

$$E = \frac{p}{2} \quad (3.12)$$

where p is the depolarizing probability from Equation 3.1. Since teleporting a pure state through a depolarized Bell pair with parameter p applies the single-qubit depolarizing channel with the same parameter p to the teleported state, the backbone fidelity F directly determines the depolarizing parameter of the teleported qubit. Substituting this relation into the secret key fraction in Equation 3.10 then allows the minimum backbone fidelity required for a positive secret key rate to be determined analytically.

3.3.3 CHSH Game

The CHSH game implemented in this work follows the formulation of the game introduced in Section 2.1.5. In contrast to QKD, both Alice and Bob are quantum end-nodes equipped with ion-trap processors. This section describes how the game is executed on the network architecture, the measurement strategies employed, and the fixed protocol parameters.

Game execution. To coordinate the game, a fifth node is introduced: a central classical referee (Charlie) that communicates with both Alice and Bob over classical channels. Charlie sends requests at regular intervals, each containing a random input bit $x \in \{0, 1\}$ for Alice and $y \in \{0, 1\}$ for Bob. Upon receiving a request, each party checks whether an entangled pair is currently available in their quantum memory. If both parties have access to an entangled pair, they employ the quantum strategy by performing a measurement on their respective qubit. If either party does not have an entangled pair available, both fall back to the classical strategy for that round. After obtaining their outputs, Alice and Bob each return a bit to Charlie, who evaluates the winning condition $a \oplus b = x \cdot y$.

The reported metric is the overall winning probability p_{win} , which is a weighted average of the quantum and classical strategies depending on the fraction of rounds in which an entangled pair was available. A winning probability meaningfully exceeding the classical limit of 0.75 therefore requires both that the shared entangled states are of sufficient quality and that they are available sufficiently often.

Classical and quantum strategies. The optimal classical strategy for the CHSH game achieves a winning probability of 0.75. This is realized by both parties always outputting

3. METHODS

the same fixed bit regardless of their input: $a = b = 0$. This strategy is used as the fallback whenever an entangled pair is not available.

For the quantum strategy, Alice and Bob perform measurements on their respective halves of the shared entangled pair and return the outcome. Assuming the state is $|\Phi^+\rangle$, the optimal measurement settings are as follows [20]. Alice measures in the Z basis for input $x = 0$ and in the X basis for input $x = 1$. Bob measures in the eigenbasis of $(Z + X)/\sqrt{2}$ for input $y = 0$ and in the eigenbasis of $(Z - X)/\sqrt{2}$ for input $y = 1$. These are the standard CHSH-optimal angles and achieve the maximum quantum winning probability of approximately 0.8536 on a perfect Bell pair. It should be noted that in the presence of dephasing noise, alternative measurement settings may in principle lead to a larger CHSH correlator and therefore higher win probability. However, no such optimization was performed in this work.

Teleportation corrections. The entangled pair shared between Alice and Bob is not generated directly but results from two entanglement swaps at the border nodes. Depending on the BSM outcomes at each border node, the resulting state may be any of the four Bell states rather than $|\Phi^+\rangle$. Since any Bell state is related to $|\Phi^+\rangle$ by a local Pauli operation on one of the two qubits, we can simply adapt the measurement settings accordingly.

Cutoff. We impose a cutoff at half the coherence time of the simulated hardware to prevent excessively decohered qubits from degrading the protocol performance. To see how this effects the win-rate, let

$$p_Z(t) = \frac{1 - e^{-2t^2/T^2}}{2} \quad (3.13)$$

denote the single-qubit phase-flip probability under Gaussian dephasing[4]. Since the two qubits are in two separate nodes, the Bell state is mapped to its Z -flipped counterpart whenever an odd number of phase flips occurs, which happens with probability

$$p_{\text{odd}} = 2p_Z(1 - p_Z). \quad (3.14)$$

For the fixed CHSH measurement settings that are optimal for $|\Phi^+\rangle$, the state $|\Phi^+\rangle$ yields $S = 2\sqrt{2}$, while its Z -flipped counterpart $|\Phi^-\rangle$ yields $S = 0$. The average correlator therefore becomes

$$S(t) = 2\sqrt{2}(1 - 2p_Z(1 - p_Z)) = \sqrt{2}(1 + e^{-4t^2/T^2}). \quad (3.15)$$

At $t = T/2$, this evaluates to

$$S\left(\frac{T}{2}\right) = \sqrt{2}(1 + e^{-1}) \approx 1.93, \quad (3.16)$$

which lies slightly below the classical threshold $S = 2$ and therefore leads to a win rate of 0.741 via Equation 2.14. This shows that, for the fixed CHSH measurement settings considered here, a cutoff at half the coherence time is close to, but not fully optimal, since states stored this long no longer outperform a classical strategy. Nevertheless, the choice of $T_{\text{cutoff}} = T_{\text{coh}}/2$ is retained in this work in order to use a uniform cutoff criterion across the different protocols, particularly for consistency with the VBQC analysis that follows. In addition, we discuss this further in the conclusion of this work.

Request rate. The baseline request rate is fixed at 16 Hz, corresponding to one request per baseline coherence time ($T_{\text{coh}} = 62$ ms). This rate is held constant across all values of the improvement factor k , so that different hardware configurations are compared on equal footing. The choice is motivated by both physical and practical considerations. From a physical perspective, if the backbone is unable to deliver entanglement within one coherence time, the majority of stored qubits will be discarded before a request arrives, and the quantum strategy is rarely used regardless of the request rate. From a simulation perspective, a request rate much lower than the entanglement generation rate leads to the simulation spending most of its time producing entangled pairs that are never consumed, drastically increasing the computation time per useful data point.

As an additional safeguard, the request rate is increased whenever the expected end-to-end entanglement generation rate exceeds ten times the request rate. In this regime the quantum strategy is always used, so the overall winning probability is dominated by the quantum strategy which can be seen in Figure A.2. This allows sampling of fast entanglement rates without massive simulation times.

Analytical model for the perfect metropolitan hardware limit. As for QKD, it is useful to consider the limit of perfect hardware ($k \rightarrow \infty$), where the backbone fidelity is the only remaining source of imperfection. For a maximally entangled state subject to a two-qubit depolarizing channel with parameter p , the CHSH correlator with optimal measurement settings evaluates to $S = (1 - p)2\sqrt{2}$ [11]. The corresponding winning probability is then

$$p_{\text{win}} = \frac{1}{2} \left(1 + \frac{(1-p)\sqrt{2}}{2} \right). \quad (3.17)$$

Using the relation $p = \frac{4}{3}(1 - F)$ between the depolarizing probability and the fidelity of the backbone Bell pair, this provides a direct mapping from backbone fidelity to winning probability.

3.3.4 Blind quantum computation

The VBQC protocol considered in this work is based on the noise-robust verified blind quantum computation scheme introduced in [28]. In this protocol, computation rounds and test rounds are interleaved. The computation rounds are used to produce the desired classical output, while the test rounds serve to estimate the error rate of the system and to detect deviations from the prescribed protocol. In the present work, only the test rounds are considered, as the aim is to investigate the protocol failure probability under realistic noise. In addition, we focus our investigation on a 5-qubit linear graph.

Test-Rounds The idea of the test round is based on embedding trap and dummy qubits into a graph. The former are prepared in random states from the XY -plane of the Bloch sphere: $|\theta\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle)$, where θ is chosen uniformly at random from the set $\{k\pi/4\}_{k=0}^7$. The latter are each randomly initialized in one of the computational basis states $\{|0\rangle, |1\rangle\}$.

3. METHODS

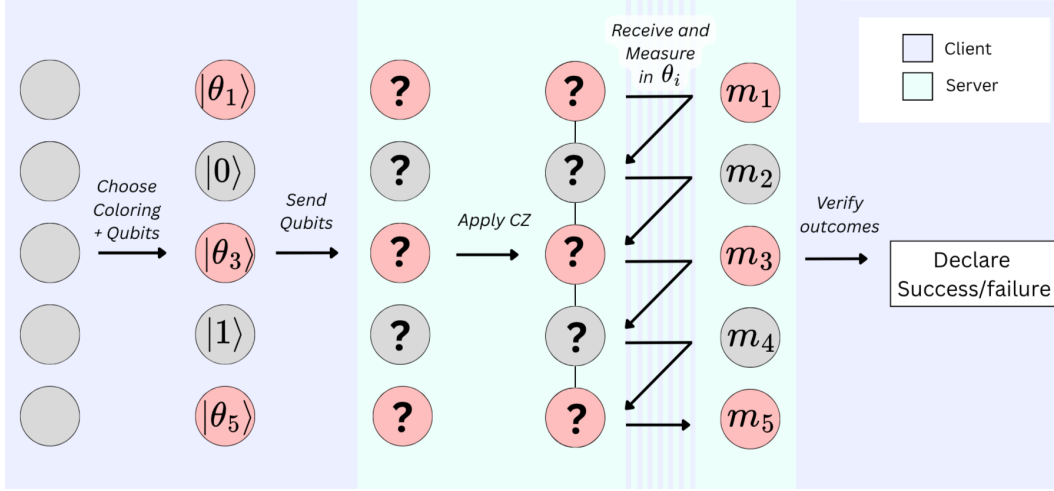


Figure 3.3: Visualization of a VBQC test round on a five-qubit linear graph. The client selects a coloring, prepares trap qubits in random equatorial states and dummy qubits in random computational basis states, and sends them to the server. The server applies the CZ gates and measures each qubit as instructed. The client then verifies whether the trap outcomes match the expected results and declares the round a success or failure.

The trap and dummy assignments are determined by a graph coloring: all qubits of the chosen color become traps, and all others become dummies. This ensures that no two traps are adjacent in the graph, so every neighbor of a trap is a dummy. For the 5-qubit linear graph, this leaves only two possible assignments: T,D,T,D,T and D,T,D,T,D, where T denotes a trap qubit and D denotes a dummy qubit. Since the dummy qubits are in the computational basis, the controlled-Z gates connecting them to neighboring traps act in a simple manner. A dummy in $|0\rangle$ leaves the trap unchanged, while a dummy in $|1\rangle$ applies a Z gate to it. The traps are therefore disconnected from the rest of the graph, and each trap remains in a known pure state up to Z corrections determined by its neighboring dummies. If the client instructs the server to measure a trap qubit in the basis defined by the angle $\delta_v = \theta_v + r_v\pi$, where $r_v \in \{0, 1\}$ is a uniformly random bit serving as a one-time pad and θ_v absorbs any corrections, the expected measurement outcome in the absence of noise or tampering is

$$b_v = r_v \oplus d_v, \quad (3.18)$$

where $d_v = \bigoplus_{i \in N(v)} s_i$ is the parity of the computational basis states $d_i \in \{0, 1\}$ of the neighboring dummy qubits. Since the client knows both the dummy states and the random pad r_v , it can predict b_v and verify the server's reported outcome. A test round is declared successful if all trap outcomes match these predictions, and failed otherwise.

Protocol metric We now use the results found in [28], restricted to the five-qubit linear graph state shown in Figure 2.2b. This graph admits a 2-coloring ($k = 2$), and we assume an inherent error probability of $p = 0$. The security threshold from [28] then requires the

test round failure probability to satisfy

$$p_{\text{fail}} < \frac{1}{k} \cdot \frac{2p-1}{2p-2} = \frac{1}{4}. \quad (3.19)$$

Qubit delivery. Each test round requires the client (Alice) to deliver five qubits to the server (Bob), where each qubit is either a trap prepared in a random equatorial state or a dummy prepared in a random computational basis state. In the network architecture considered here, Alice does not send these qubits directly. Instead, each qubit is delivered through the following sequence of operations. First, entanglement is generated over the backbone between the two border nodes. Subsequently, both metropolitan links start entanglement generation simultaneously: Alice performs remote state preparation on her side by measuring the photon received from her border node in the appropriate basis, while Bob's border node generates entanglement with Bob. As soon as either succeeds, the corresponding border node performs a DBSM and announces the corrections. Depending on the order in which the quantum communication steps succeed, either entanglement is swapped at a border node, or Alice's remotely prepared state is teleported further through the network. In both cases, once both border nodes have successfully established quantum communication with their respective end nodes and performed their DBSM, Bob receives Alice's state up to the announced correction operations.

Corrections and angle adaptation. The delivery process introduces a maximum total of five correction operations per qubit: one Z correction from the remote state preparation and four corrections (two X and two Z) from the two teleportation steps. Denoting the original state that Alice intends to prepare as $|\theta_{1,v}\rangle$ and the state that arrives at Bob as $|\theta_{2,v}\rangle$, the relationship between the two is

$$|\theta_{2,v}\rangle = Z^{a_2} X^{b_2} Z^{a_1} X^{b_1} Z^{m_{\text{RSP}}} |\theta_{1,v}\rangle = Z^{a_1+a_2+m_{\text{RSP}}} X^{b_1+b_2} |\theta_{1,v}\rangle, \quad (3.20)$$

up to a global phase, where $m_{\text{RSP}} \in \{0, 1\}$ is the outcome of Alice's remote state preparation measurement and $a_i, b_i \in \{0, 1\}$ are the outcomes of the i -th Bell state measurement. These corrections do not need to be physically applied at the server. Instead, they can be absorbed into the measurement instructions using the measurement calculus relations given in Equation 2.16 and 2.17. Applying these relations to Equation 3.20 yields the adapted measurement angle

$$\theta_{2,v} = (-1)^{b_1+b_2} \theta_{1,v} + (a_1 + a_2 + m_{\text{RSP}}) \pi \quad (3.21)$$

Alice then instructs the server to measure qubit v at the angle

$$\delta_v = \theta_{2,v} + r_v \pi, \quad (3.22)$$

where $r_v \in \{0, 1\}$ is a uniformly random bit that serves as a one-time pad on the measurement outcome. After accounting for the effect of neighboring dummy qubits, the expected measurement outcome for a trap qubit is

$$b_v = r_v \oplus d_v, \quad (3.23)$$

3. METHODS

where d_v is the parity of the computational basis states of the neighboring dummy qubits, including the teleportation corrections on those qubits. Since the client knows both the dummy states and all correction terms, they can predict d_v and verify whether the server's reported outcome matches the expectation.

For dummy qubits, the situation is simpler. Since dummies are prepared in the computational basis, only the X corrections from the teleportation steps are relevant: Z corrections act trivially on $|0\rangle$ and $|1\rangle$. The client tracks these X corrections to determine the effective computational basis state of each dummy at the server, which in turn determines the expected phase contribution to neighboring trap qubits via the controlled- Z gates.

Measure-as-you-go. In the architecture considered here, only one qubit can be delivered to the server at a time. Receiving all five qubits sequentially while maintaining coherence on the earliest arrivals poses a significant challenge, as each subsequent delivery attempt adds decoherence to the qubits already in memory. Fortunately, the protocol does not require all five qubits to be present at the server simultaneously. In a linear graph, each qubit is connected to at most two neighbors. As soon as a qubit and all of its neighbors are available, the required controlled- Z gates can be applied, after which the qubit may be measured immediately, provided that this measurement is consistent with the flow. This technique, referred to as Measure-as-you-go, means that the server needs to hold at most two qubits in memory at any point during the protocol in the case of the linear graph. Consequently, the test round shown in Figure 3.3 is not exactly how it is implemented in the simulation. Instead, for the five-qubit linear graph, the procedure is as follows: after the first two qubits have been received, they are connected via a controlled- Z gate and the first qubit is measured; when the third qubit arrives, it is connected to the second and the second is measured and so on. This substantially reduces the memory time per qubit and thereby the accumulated decoherence. However, received qubits may, by chance, still remain in memory long enough to decohere beyond the point of useful contribution.

Cutoff. To prevent such excessively decohered qubits from corrupting the protocol, a cutoff time is imposed: any qubit that has been stored in memory for longer than half the coherence time of the simulated hardware is discarded. If a qubit is discarded during a test round, the entire round is restarted. The choice of half the coherence time as the cutoff is adopted from [39] and is not optimized over in this work. In certain parameter regimes there is therefore room for improvement via an optimized cutoff time, but finding such an optimum is tricky due to the sequential and interdependent nature of the five-qubit deliveries, as well as the significant rate trade-off introduced by shorter cutoff times.

Analytical model for the perfect metropolitan hardware limit. To complement the simulation results, an analytical expression for the test round failure probability can be derived, again under the assumption of perfect metropolitan hardware, modeling only the noise introduced by the backbone as a depolarizing channel where teleportation with the produced pair introduces a single-qubit error probability p which can be mapped to the fidelity of the backbone. For the five-qubit linear graph with two possible colorings, averaging over the

two scenarios yields

$$p_{\text{fail}} = 1 - \left(1 - \frac{p}{2}\right)^5 + \left(1 - \frac{p}{2}\right)^3 \left(\frac{p}{2}\right)^2 + 3 \left(1 - \frac{p}{2}\right)^2 \left(\frac{p}{2}\right)^3 + \left(1 - \frac{p}{2}\right) \left(\frac{p}{2}\right)^4. \quad (3.24)$$

where p is the depolarizing probability per qubit assumed uniform over the graph (and not to be confused with the failure probability p_{fail}). This expression comes from the fact that the depolarizing channel maps a pure state to its orthogonal state with probability $p/2$ following Equation 3.1. The equation above is then obtained by identifying which combinations of such flips can be compensated by corresponding flips on neighboring qubits, such that the effective measurement outcome remains correct. The full derivation is provided in Appendix A.2. This expression allows the backbone fidelity threshold to be estimated analytically as setting $p_{\text{fail}} = 0.25$ and solving for p yields the maximum tolerable depolarizing probability. This serves as a consistency check for the simulation results.

3.3.5 Simulation procedure

With the network model, hardware parameters, and protocol implementations specified, the simulation framework allows each protocol to be evaluated for arbitrary choices of the backbone rate R , backbone fidelity F , and hardware improvement factor k . We now describe how these tunable parameters are varied in order to answer the research questions. Each question is translated into a parameter sweep in which all parameters are kept fixed except for the quantity whose threshold is being investigated. For each point in such a sweep, the corresponding protocol is simulated repeatedly, the relevant protocol-specific performance metric is extracted, and the threshold crossing is estimated from the resulting data.

For RQ1, the backbone is assumed to be perfect. It delivers Bell pairs with unit fidelity and at an effectively infinite rate. In this regime, only the hardware improvement factor k is varied. The resulting performance curve is then used to determine the smallest value of k for which the protocol reaches the required performance threshold via linear interpolation.

For RQ2, the backbone rate is kept effectively infinite, while the backbone fidelity F is varied for fixed values of k . This sweep isolates the effect of backbone fidelity and gives the minimum backbone fidelity required at each chosen hardware-improvement factor. For RQ3, the complementary sweep is performed: the backbone fidelity is set to unity, while the delivery rate R is varied. This determines the minimum backbone rate required when the delivered states themselves are assumed to be perfect.

Finally, RQ4 combines these two backbone parameters. For fixed values of k , simulations are performed over a two-dimensional grid in (R, F) -space. The same thresholding procedure is then applied at each point to identify the region in which the protocol can be implemented successfully. The statistical treatment of the simulated data and the extraction of uncertainties are described in the following section.

3.3.6 Data analysis

The simulation outputs for all three benchmark protocols can be reduced to binary outcomes. For each protocol, we define a single trial as the smallest unit of data that contributes to the metric of interest. For the CHSH game, one trial corresponds to one request from the referee. For VBQC, one trial is a completed test round that was not restarted due to a cutoff event. For QKD, one trial corresponds to one raw key bit, that is, a round in which Alice and Bob chose matching bases. In each case, a trial either succeeds or fails: Alice and Bob either win or lose the CHSH round, a VBQC test round either passes or fails trap verification, and a QKD raw key bit is either correct or erroneous. The metrics of interest, namely the winning probability, the failure probability, and the quantum bit error rate, are therefore all proportions estimated from a finite number of trials n .

This structure makes the error analysis straightforward. Throughout this work, statistical uncertainties are quantified using the standard error of a proportion,

$$\text{SE} = \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}}, \quad (3.25)$$

where \hat{p} is the estimated proportion and n is the number of trials. When the metric of interest is derived from the raw proportion, such as the secret key rate which depends on the QBER through Equation 3.10, the uncertainty is propagated using standard first-order error propagation. Unless stated otherwise, all uncertainties reported in this thesis correspond to ± 2 standard errors, i.e. approximately 95% confidence intervals.

To ensure sufficient statistical precision, we require a minimum of 60 000 trials per simulation run for each setting. At this sample size, the standard error on the estimated proportions remains below one percent of the metric value in the parameter regimes of interest.

Crossing point estimation. A recurring task in the analysis is to determine the critical backbone fidelity at which a protocol metric crosses a given threshold, for example the fidelity at which the CHSH winning probability exceeds 0.75 or the VBQC failure probability drops below 0.25. Since the simulation provides metric estimates at discrete fidelity values, the crossing point must be interpolated. We estimate it by linear interpolation between the two simulated points that bracket the threshold. Given two fidelity values F_1 and F_2 with corresponding metric estimates $\hat{m}_1 \pm \sigma_1$ and $\hat{m}_2 \pm \sigma_2$, the critical fidelity is

$$F_c = F_1 + \frac{(m^* - \hat{m}_1)(F_2 - F_1)}{\hat{m}_2 - \hat{m}_1}, \quad (3.26)$$

where m^* is the threshold value. The uncertainty on F_c is obtained via first-order error propagation, treating \hat{m}_1 and \hat{m}_2 as independent:

$$\sigma_{F_c} = \sqrt{\left(\frac{\partial F_c}{\partial \hat{m}_1} \sigma_1\right)^2 + \left(\frac{\partial F_c}{\partial \hat{m}_2} \sigma_2\right)^2}, \quad (3.27)$$

with partial derivatives

$$\frac{\partial F_c}{\partial \hat{m}_1} = \frac{(F_2 - F_1)(m^* - \hat{m}_2)}{(\hat{m}_2 - \hat{m}_1)^2}, \quad \frac{\partial F_c}{\partial \hat{m}_2} = \frac{-(F_2 - F_1)(m^* - \hat{m}_1)}{(\hat{m}_2 - \hat{m}_1)^2}. \quad (3.28)$$

This estimate assumes that the metric varies approximately linearly between the two bracketing points, which is reasonable when the fidelity spacing is sufficiently small. The accuracy of this approximation and its implications for the reported crossing points will be revisited in the discussion of the results. Unless stated otherwise, uncertainties on reported crossing points are again given as $\pm 2\sigma_{F_c}$.

3.4 Summary

This chapter has described the simulation framework and how we simulate the three quantum communication protocols on the four-node network architecture introduced. Each protocol defines a performance metric with a clear threshold that separates feasible from infeasible operation: a positive secret key fraction for QKD, a winning probability exceeding the classical limit of 0.75 for the CHSH game, and a test round failure probability below 0.25 for VBQC. These thresholds serve as the targets of the investigation, as they mark the boundary at which the protocol provides a quantum advantage or satisfies its security guarantees.

In addition to differing in their metrics and noise sensitivity, the three protocols require different end-node configurations. QKD uses two measurement-only clients, the CHSH game requires two full quantum nodes, and VBQC combines one measurement-only client with one quantum server. Together, these three protocols exhaust all possible combinations of end-node types available in the model, ensuring that the benchmark suite probes the network architecture comprehensively. In the simulations, we systematically vary the backbone fidelity F , the backbone rate R , and the hardware improvement factor k in order to answer the research questions formulated in Section 1.1.

The key properties of each protocol are summarized in Table 3.4. The backbone fidelity bounds listed in the table correspond to the analytical limits derived under the assumption of perfect metropolitan hardware ($k \rightarrow \infty$), and represent the strictest possible requirement imposed by the backbone alone. We already observe that VBQC is the most demanding protocol in terms of fidelity. This is further amplified by its sequential structure, since five qubits must be delivered one after another, increasing the exposure to memory decoherence and causing errors to accumulate over the course of the protocol. As a result, VBQC will consistently impose the strongest requirements on the network throughout the results presented in this chapter.

Table 3.4: Summary of the three benchmark protocols, their feasibility thresholds, the minimum backbone fidelity required under perfect metropolitan hardware, and the end-node configuration used in each case.

| Protocol | Threshold metric | Min. backbone fidelity | End-node types |
|-----------|--------------------------|------------------------|---------------------------|
| QKD | $s > 0$ | $F \gtrsim 0.853$ | 2 × MO client |
| CHSH game | $p_{\text{win}} > 0.75$ | $F \gtrsim 0.780$ | 2 × quantum node |
| VBQC | $p_{\text{fail}} < 0.25$ | $F \gtrsim 0.915$ | 1 × MO client, 1 × server |

Chapter 4

Results and discussion

This chapter is organized around the four research questions introduced in Section 1.1. We begin in Section 4.1 with the perfect-backbone regime, in which the backbone is assumed to deliver entanglement with unit fidelity and effectively infinite rate. This isolates the limitations imposed by the metropolitan hardware alone and allows us to determine the minimum hardware improvement factor required for each protocol, thereby answering RQ1. We then turn to Section 4.2, which considers a fast backbone with effectively infinite entanglement generation rate but imperfect fidelity, in order to determine the minimum backbone fidelity required once the local hardware has been improved, which addresses RQ2. In Section 4.3, we study the complementary quality-backbone regime, in which the backbone fidelity is fixed to unity while the entanglement generation rate is finite, allowing us to identify the minimum required backbone rate, corresponding to RQ3. Finally, Section 4.4 relaxes both assumptions simultaneously and maps out the feasible operating region in the full backbone parameter space spanned by rate and fidelity, in line with RQ4. Section 4.5 then discusses the broader trends, limitations, and robustness of these results.

4.1 Perfect backbone

Even though the backbone plays the central role in this investigation, we begin by essentially removing it from the analysis. This is done by assuming a perfect backbone, that is, one that delivers entanglement with unit fidelity $F = 1$ at an infinite rate. While this may seem counterintuitive for a study centered on backbone requirements, it serves a clear purpose: by eliminating the backbone as a source of noise, we isolate the contributions of the metropolitan hardware and identify the regime of improvement factors in which the protocols are feasible at all. Any hardware configuration that fails to meet the protocol thresholds under a perfect backbone will necessarily also fail under any realistic one, regardless of its quality. This effectively reduces the parameter space and establishes a lower bound on the improvement factor k for each protocol. It is worth noting that the assumption of backbone-first entanglement generation, introduced in Section 3.3.1, is also somewhat artificial in this limit, since with an infinite-rate backbone, performance would certainly improve if the ordering were reversed. Nevertheless, we retain this ordering throughout, as the goal of this

4. RESULTS AND DISCUSSION

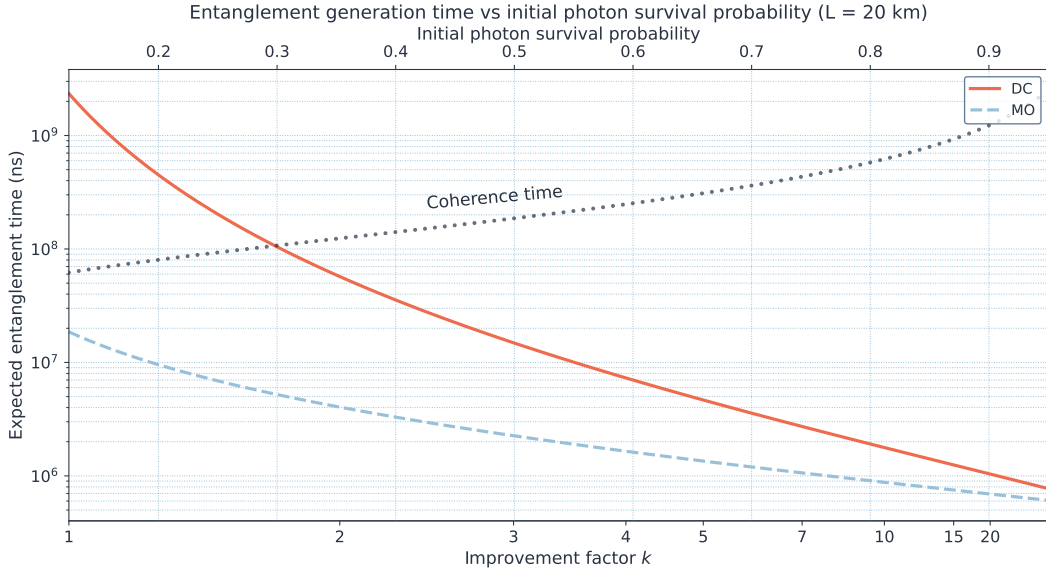


Figure 4.1: Average metropolitan entanglement-generation times as a function of the hardware improvement factor k . The plot compares the double-click entanglement-generation time between two quantum processors with the time required for a measurement-only client to successfully detect a photon. As k increases, the metropolitan entanglement times become progressively less relevant to protocol performance, since they become small compared with the simultaneously increasing coherence time.

analysis is to quantify the strain that the metropolitan hardware places on the setup under the same assumptions used in all subsequent sections.

Metropolitan entanglement times.

Before addressing the minimum improvement factors directly, it is instructive to examine how long metropolitan entanglement generation is expected to take and how this compares to the coherence time of the quantum memories. Figure 4.1 shows the average successful entanglement time as a function of the improvement factor for both double-click (DC) heralded entanglement generation for processing node clients and measurement-only (MO) remote state preparation over a 20 km metropolitan link.

At the baseline parameter set ($k = 1$), the average DC entanglement time exceeds the coherence time by more than an order of magnitude. Under the backbone-first ordering assumed throughout this work, this means that by the time a DC entanglement attempt succeeds, the qubit stored at the border node from the backbone link will have almost certainly completely decohered. Only at an improvement factor of roughly $k = 2$ does the average DC entanglement time drop below half the coherence time, which is the cutoff threshold adopted in this work. This marks the onset of the regime in which DC-based protocols can frequently produce qubits that retain useful coherence.

MO-based remote state preparation, by contrast, is substantially faster even at current




Figure 4.2: Minimum hardware improvement factor required for each protocol to cross its respective threshold metric under a perfect backbone. Error bars indicate the uncertainty from linear interpolation between simulated data points.

hardware parameters. This difference is explained by the scaling of the two schemes with the initial photon survival probability η . The DC protocol requires two photons to arrive at the midpoint station and be detected and therefore scales as η^2 , whereas MO entanglement generation requires only a single photon to reach the client and scales as η . Attenuation alone accounts for the same loss as the total length is 20km each. As the improvement factor increases and η grows, the two curves converge as observed in Figure 4.1. They eventually meet because the DC protocol benefits from requiring only half the classical communication time compared to MO, while its inherent success probability of 0.5, due to the probabilistic nature of the Bell state measurement, cancels this advantage.

It should be noted that, since the two metropolitan links are generated simultaneously, the relevant timescale for the protocol is not the individual entanglement time but the time until both links have succeeded. For symmetric setups where both sides use the same scheme (MO-MO or DC-DC), the expected time until both links are ready is roughly 50% longer than the expected time for a single link, due to the statistics of waiting for two independent geometric processes to complete. For the asymmetric setup used in VBQC (MO-DC), the slower DC link dominates and the combined expected time is well approximated by the DC entanglement time alone.

Minimum Improvement Parameters

Having established the entanglement time landscape, we now turn to the central question of this section, corresponding to Research Question 1: what is the minimum hardware improvement factor required to implement each protocol assuming a perfect backbone? The answer is obtained by finding, for each protocol, the smallest value of k at which the metric from Table 3.4 is satisfied.

The results are shown in Figure 4.2. The minimum improvement factors are $k = 1.85 \pm 0.03$ for QKD, $k = 2.86 \pm 0.06$ for the CHSH game, and $k = 6.42 \pm 0.09$ for VBQC. Given the entanglement-time analysis presented above, this ordering is as expected. QKD relies on two MO clients, which benefit from the more favorable η scaling and do not require quantum memories at the end nodes. In addition, MO clients do not introduce further fidelity

degradation through local ion-photon emission or noise from the double-click entanglement generation scheme. As a result, QKD can be implemented at the lowest improvement factor. By contrast, the CHSH game requires two quantum end-nodes with memory and depends on DC entanglement generation on both sides, yet its very modest fidelity threshold, seen in Table 3.4, allows it to be performed well before VBQC becomes feasible. VBQC, in turn, demands not only heralded entanglement generation between the server and its border node but also the sequential delivery of five qubits, each requiring additional gate operations. Combined with the strictest fidelity requirement of the three protocols, this places it firmly as the most demanding.

The uncertainties on the three values differ noticeably, growing from QKD to VBQC. This is a consequence of the nonlinear nature of the improvement procedure defined in Equation 3.7. At low values of k , a unit increase in the improvement factor produces a large change in the underlying hardware parameters and therefore a large change in the protocol metric. At higher values of k , the same unit increase produces progressively smaller relative improvements. Since the crossing points are estimated via linear interpolation between simulated data points, a steeper metric-versus- k curve yields a tighter bracket and thus a smaller uncertainty, while a shallower curve leads to a wider one.

These three improvement factors define the regime in which a meaningful investigation of backbone requirements is possible. For any k below these thresholds, no backbone, however ideal, can compensate for the limitations of the metropolitan hardware. In the sections that follow, we take k above these thresholds and progressively relax the backbone assumptions to determine the minimum fidelity and rate that the backbone must provide.

4.2 Fast backbone

Having established the minimum improvement factors under a perfect backbone, we now begin to relax the assumptions on the backbone itself. In this section, the backbone rate is kept practically infinite while the fidelity is allowed to take values below unity. This corresponds to a backbone that can deliver entanglement quickly but with imperfect quality, and allows us to answer RQ2: for a given improved parameter set, what is the minimum backbone fidelity required to implement each protocol?

Metric vs. backbone fidelity

Figure 4.3(a)-(c) shows the performance metric of each protocol as a function of the backbone fidelity for several values of the improvement factor k . In each panel, the dashed line indicates the analytical model corresponding to perfect metropolitan hardware ($k \rightarrow \infty$), which isolates the contribution of the backbone fidelity alone.

For QKD (Figure 4.3a), the metric of interest is the secret key fraction obtained from Equation 3.10. Even at moderate improvement factors the protocol performs reasonably well, consistent with the low minimum k found in the previous section. As the hardware improves, the curves shift to the left, meaning that progressively lower backbone fidelities can still sustain a positive secret key rate. The gains, however, are strongly nonlinear. Compared with $k = 2$, increasing the improvement factor to $k = 5$ allows the backbone-fidelity

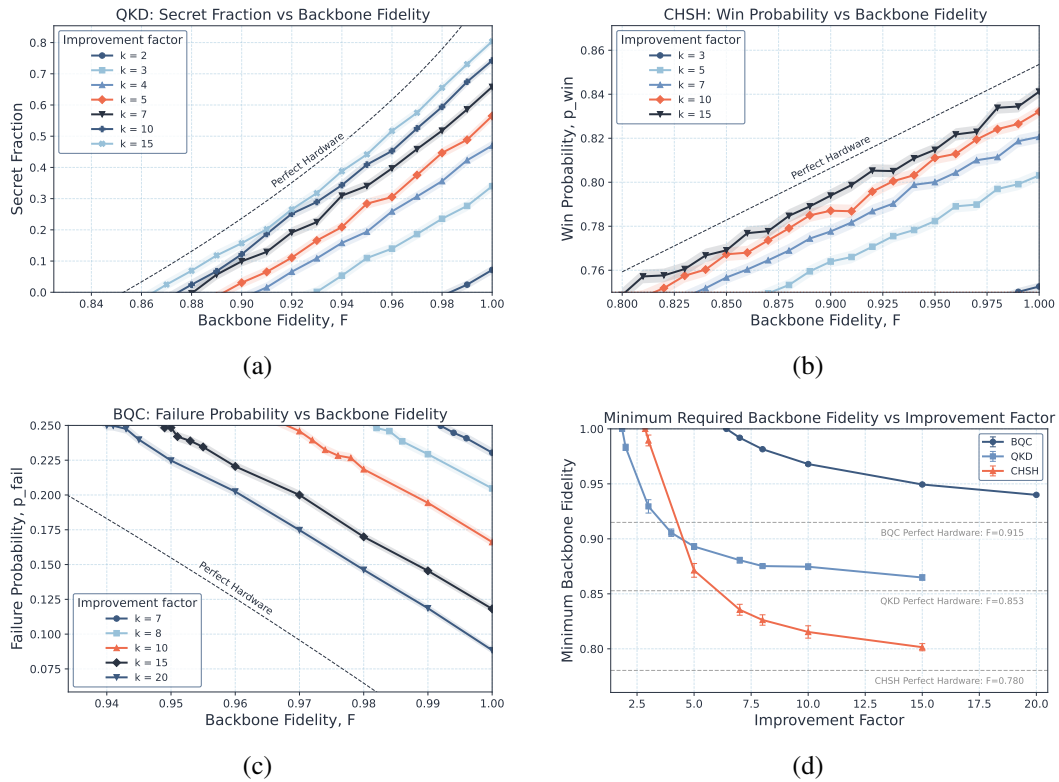


Figure 4.3: Fast-backbone regime. (a)-(c) Performance metric versus backbone fidelity for QKD, CHSH, and VBQC, respectively, for several metropolitan hardware improvement factors k , assuming an infinite backbone rate. Dashed curves show the analytical perfect-metropolitan-hardware limit ($k \rightarrow \infty$). The protocol thresholds are a positive secret key fraction for QKD, a winning probability above 0.75 for CHSH, and a test-round failure probability below 0.25 for VBQC. (d) Minimum backbone fidelity obtained from the threshold crossings in (a)-(c), shown as a function of k . Horizontal dashed lines denote the analytical limits, and the points at $F=1$ coincide with the minimum improvement factors from the perfect-backbone analysis. Shaded regions in (a)-(c) show two-standard-error uncertainties, while error bars in (d) indicate $2\sigma_{F_c}$ uncertainties on the interpolated threshold fidelities.

requirement to be relaxed by approximately 0.09 in absolute value while still obtaining positive protocol performance. Increasing the improvement factor further to $k = 15$ provides a much smaller additional relaxation of only about 0.03. This diminishing return reflects the way the improvement factor k acts on probabilities of no imperfection. The effective reduction in noise is strongest at small k , while further increases yield progressively smaller gains.

The CHSH game (Figure 4.3b) shows the winning probability at the fixed request rate. At $k = 3$, the protocol operates at the very edge of feasibility and requires a backbone of near-unit fidelity. The jump from $k = 3$ to $k = 5$ is substantial: at an improvement fac-

tor of 5, a quantum advantage can be achieved for backbone fidelities as low as roughly 0.875. The simulated curves exhibit an approximately linear dependence on the backbone fidelity, which is expected in this regime. With entanglement generation happening effectively instantaneously, memory decoherence plays a negligible role and the dominant noise mechanisms are depolarizing channels, which contribute linearly to the winning probability through Equation 3.17. For $k = 3$, the approximately linear behavior should be interpreted with some caution. At this low improvement factor, the quantum strategy is employed in only about 60% of the rounds, so the resulting winning probability contains a substantial non-quantum contribution which decreases the overall slope of the line. Moreover, memory decoherence is no longer negligible, such that the effective noise is not expected to be captured fully by the depolarizing model underlying Equation 3.17. However, since the viable fidelity range at $k = 3$ is so narrow, neither effect is clearly visible in the figure. Another interesting feature visible in this figure is that, as the hardware improves, the fidelity requirement for CHSH decreases more rapidly than for QKD. This is a consequence of the lower fidelity bound inherent to the CHSH-game (Table 3.4), which leaves more room for the backbone quality to be relaxed once the metropolitan noise is sufficiently suppressed.

For VBQC (Figure 4.3c), the metric is the test round failure probability, which must remain below 0.25. The protocol threshold therefore corresponds to the top of the plot rather than the bottom, and the relevant question is at what backbone fidelity the curves cross upward through this bound. It is worth noting that the fidelity axis in this panel spans only the range 0.93-1.0, which visually amplifies the gap between the simulated curves and the analytical limit. Despite this compressed scale, even at $k = 20$ a noticeable offset from the perfect-hardware curve remains, reflecting the higher sensitivity of the VBQC protocol to residual gate and emission noise, whose effects compound over the full protocol.

Minimum fidelity requirements

Figure 4.3d summarizes the crossing points extracted from Figures 4.3a-4.3c, where each crossing point corresponds to the backbone fidelity at which the protocol metric reaches the relevant performance threshold, thereby addressing RQ2. The resulting curve shows the minimum backbone fidelity required to implement each protocol as a function of the improvement factor. As the metropolitan hardware improves, the backbone fidelity can be progressively relaxed for all three protocols, with each curve converging toward the analytical limit indicated by the horizontal dashed lines.

At the upper-left edge of the figure, the three curves meet a minimum backbone fidelity of $F = 1.0$ at precisely the improvement factors found in the previous section ($k = 1.85$, 2.86, and 6.42 for QKD, CHSH, and VBQC respectively), as those data points can be taken as starting points of this analysis. As k increases, the curves separate and their relative ordering evolves. For near-term hardware improvements, QKD requires the lowest backbone fidelity of the three protocols, owing to its simpler metropolitan hardware requirements. However, at higher improvement factors the CHSH curve crosses below the QKD curve. This is a direct consequence of the lower fidelity bound inherent to the CHSH protocol (Table 3.4), which leaves more room for the backbone quality to be relaxed once metropolitan noise has been sufficiently suppressed.

Even at high improvement factors of $k \geq 15$, a noticeable gap remains between the simulated crossing points and the analytical limits for both CHSH and VBQC. This gap reflects the fact that the dominant fidelity-related noise sources in these protocols, particularly the emission fidelity, visibility and gate fidelities, are more prevalent due to the additional operations required by these protocols. Under the improvement mapping defined in Equation 3.7, these parameters approach unity only slowly after the initial quick improvement, leaving residual noise that continues to tighten the backbone requirement relative to the analytical prediction. QKD, which avoids some of these operations by relying on MO clients, converges to its analytical limit considerably faster.

4.3 Quality backbone

We now consider the complementary regime to the previous section. The backbone fidelity is restored to $F = 1$ while the entanglement generation rate is reduced to finite values. This allows us to investigate RQ3, concerning the minimum backbone rate required to implement each protocol for a given improved parameter set. However, as we will discuss, the original threshold metrics from Table 3.4 do not all translate cleanly to this setting, and adapted targets are needed for two of the three protocols.

Adapted metrics

The threshold metrics used in the previous sections were designed to identify the boundary between feasible and infeasible operation. For the fidelity analysis this worked well, as backbone fidelity directly affects the quality of the distributed entangled states and therefore the protocol metrics themselves. Reducing the backbone rate, however, introduces a fundamentally different kind of degradation, and the original thresholds are no longer sufficient for two of the three protocols.

For both QKD and the CHSH game, the issue stems from the backbone-first entanglement generation ordering. Under this assumption, memory decoherence only begins after the backbone link has been successfully established. A slower backbone therefore does not affect the fidelity of the states used in the protocol. It only determines how frequently those states become available. For QKD, this means that the quantum bit error rate, and consequently the secret key fraction, remains unchanged regardless of the backbone rate. A slower backbone reduces the secret key rate but cannot, on its own, push the secret fraction to zero. Similarly, for the CHSH game, a slower backbone simply increases the fraction of rounds handled by the classical strategy. Even if only a small fraction of requests coincide with an available entangled pair, the overall winning probability will still exceed the classical bound of 0.75 given sufficient sampling, as long as the entangled states themselves are of adequate quality. In both cases, the original threshold cannot be crossed by reducing the rate alone, making it impossible to define a meaningful minimum backbone rate using the criteria from Table 3.4.

To proceed, we therefore adopt adapted targets that capture the practical utility of the protocols rather than their fundamental feasibility. For QKD, we replace the condition $s > 0$ with minimum secret key rate thresholds of $R_s \geq 5$ Hz and $R_s \geq 10$ Hz as they were

identified as reasonable targets within the regime of investigation. For the CHSH game, we replace the bound $p_{\text{win}} > 0.75$ with targets of $p_{\text{win}} \geq 0.76$ and $p_{\text{win}} \geq 0.80$. The lower target of 0.76 corresponds roughly to the winning probability obtained when a perfect-fidelity Bell pair is available in one out of every ten rounds, while the higher target of 0.80 corresponds approximately to availability in half of the rounds. It could of course also correspond to imperfect Bell pairs with compensated higher availability. These values are chosen to provide two reference points that span a practically relevant range, rather than to represent specific application requirements.

BQC does not suffer from this issue. Because the five qubits are delivered sequentially, each subsequent qubit requires a new round of backbone entanglement generation. Qubits that have already been delivered continue to decohere while waiting for the next backbone entanglement attempt to succeed. A slower backbone therefore directly increases the memory decoherence accumulated per qubit, which raises the test round failure probability and can push it above the 0.25 threshold. The original metric thus remains well-defined in this setting. The practical difficulty lies elsewhere. When the backbone rate is low, the cut-off condition triggers frequently, causing rounds to restart and dramatically increasing the time required to complete a single successful test round. This makes the low-rate regime increasingly expensive to simulate, placing a practical limit on the backbone rates that can be explored.

Adopting these adapted targets does not undermine the relevance of the analysis performed in the previous sections. The simulation data underlying the fast backbone analysis remains unchanged. Only the criterion used to extract crossing points is different. In particular, the protocol performance curves as a function of backbone fidelity still hold and will be reused when both fidelity and rate are varied simultaneously.

Metric vs. backbone rate

Figure 4.4 (a)-(c) shows the performance metric of each protocol as a function of the backbone entanglement rate at unit backbone fidelity, for several values of the improvement factor.

For QKD (Figure 4.4a), three distinct regimes can be identified. At low backbone rates, the time per protocol round is dominated by the backbone entanglement generation and the secret key rate increases approximately linearly with the backbone rate. This is expected, as the secret fraction is determined entirely by the QBER, which, as discussed in the previous section, does not depend on the backbone rate. In this regime the secret key rate is therefore directly proportional to how frequently backbone entangled pairs become available. As the backbone rate increases, the metropolitan entanglement generation, gate and measurement times begin to constitute a more significant fraction of the total round duration, causing the secret key rate to grow increasingly slowly. At sufficiently high rates, the backbone is no longer the bottleneck and further increases yield only marginal improvements, as the protocol duration is now limited by the metropolitan hardware. The boundaries between these three regimes naturally shift with the improvement factor, since better metropolitan hardware also reduces the entanglement generation times.

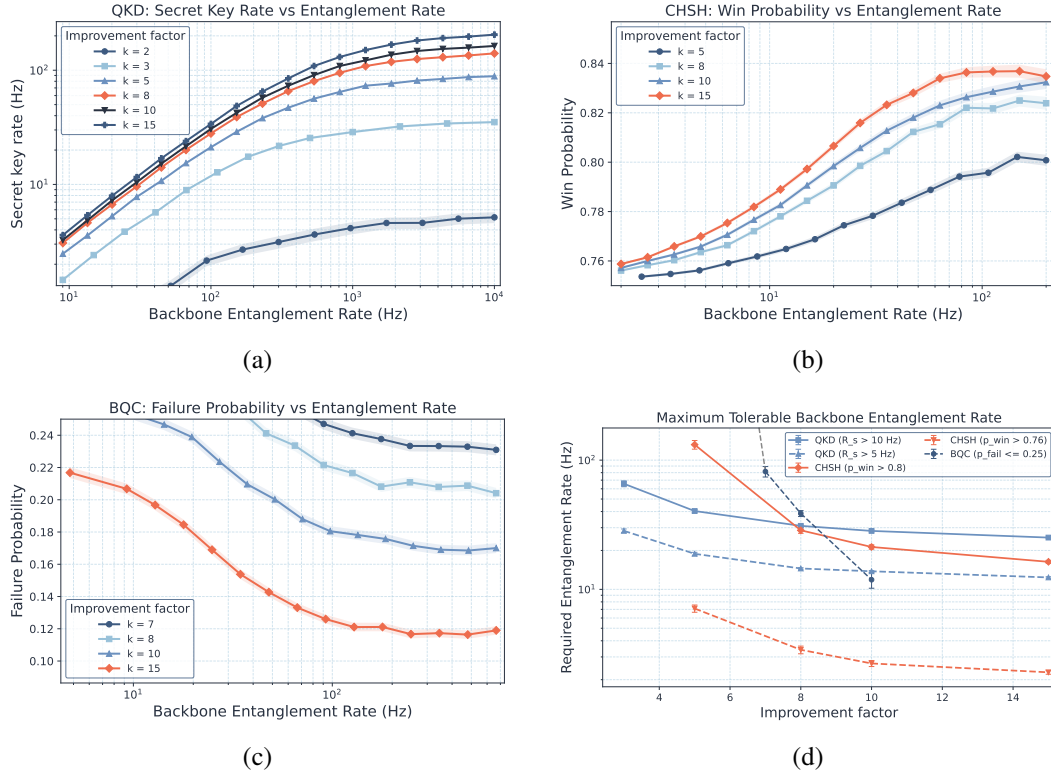


Figure 4.4: Finite-rate backbone regime at unit backbone fidelity. (a)-(c) show the performance metric of QKD, CHSH, and VBQC, respectively, as a function of the backbone entanglement rate for several values of the metropolitan hardware improvement factor k . In all panels, the backbone fidelity is fixed at $F = 1$, such that only the effect of a finite entanglement generation rate is probed. (d) summarizes the crossing points extracted from panels (a)-(c), showing the minimum backbone entanglement rate required to satisfy the adapted target metrics for each protocol as a function of k .

The log-log scaling of Figure 4.4a also makes the need for adapted metrics visually apparent. The secret key rate decreases continuously with the backbone rate but will never reach zero, confirming that the original threshold of $s > 0$ cannot be crossed by reducing the rate alone. Additionally, the error bars at a low improvement factor ($k=2$) are noticeably larger than at higher values. This is a consequence of error propagation through the secret key fraction. Since the secret fraction depends on the QBER through Equation 3.10, its absolute uncertainty remains roughly constant, but its relative uncertainty grows as the secret fraction itself approaches zero. This effect is then directly seen in the secret key rate, producing the visible spread at low k .

For the CHSH game (Figure 4.4b), the same fundamental issue is visible. As the backbone rate decreases, the winning probability approaches the classical bound of 0.75 asymptotically but never drops below it, since a slower backbone simply means that fewer rounds employ the quantum strategy while the quality of the entangled pairs will provide a bet-

ter win-rate thanks to the cutoff. At the opposite end, as the backbone rate increases, the curves saturate at the values found in the $F = 1$ regime of the fast backbone analysis, which is expected since a fast backbone with unit fidelity recovers exactly that regime. At the highest simulated rates, the quantum strategy is used in essentially all rounds. However, even among configurations where the quantum strategy is employed in more than 99% of the rounds¹, the winning probability continues to improve slightly with increasing backbone rate. This residual improvement arises because a faster backbone means that less time elapses between entanglement generation over the backbone and the subsequent metropolitan operations, resulting in fresher entangled states with less accumulated memory decoherence at the moment of measurement.

For VBQC (Figure 4.4c), as the backbone rate decreases, the failure probability increases as expected and can cross the 0.25 threshold for lower improvement factors. However, a notable feature of the figure is that at $k = 15$, the failure probability appears to remain below the threshold regardless of the backbone rate. The curve flattens rather than continuing to rise, suggesting that some mechanism prevents the failure probability from growing indefinitely as the backbone slows down.

This behavior can be explained by the role of the cutoff. To see this, it is useful to analyze the conditional memory time of a qubit that is not discarded. If we approximate the discrete geometric entanglement success distribution as a continuous exponential $p_{succ}(t) = Re^{-Rt}$ with rate R , the probability that an entanglement generation attempt succeeds before the cutoff T_c is

$$P(t < T_c) = 1 - e^{-RT_c}. \quad (4.1)$$

The conditional expected waiting time, given arrival before the cutoff, is then

$$E[t \mid t < T_c] = \frac{1}{P(t < T_c)} \int_0^{T_c} t R e^{-Rt} dt, \quad (4.2)$$

which evaluates to

$$E[t \mid t < T_c] = \frac{1}{R} - \frac{T_c}{e^{RT_c} - 1}. \quad (4.3)$$

In the limit $R \rightarrow 0$ this reduces to $T_c/2$. With the cutoff set at $T_c = T_{coh}/2$, the conditional expected memory time therefore converges to $T_{coh}/4$. The continuous approximation is justified in the rate regimes investigated here, as the per-attempt success probability of heralded entanglement generation is on the order of 1%, making the geometric distribution well approximated by an exponential. Naturally, the approximation breaks down completely as the rate approaches zero, where not even a single attempt would fall within the cutoff window. In addition, this average memory time cannot directly be used to estimate the average fidelity, as the Gaussian dephasing channel depends nonlinearly on the storage time. Nevertheless, it captures the essential mechanism behind the stagnation observed in Figure 4.4c.

Therefore, while a slower backbone causes the vast majority of rounds to be discarded, the rounds that do survive have a bounded conditional memory time that converges to

¹At a rate of 200 Hz, the quantum strategy is used in more than 99% of rounds for all improvement factors. This threshold is reached at progressively lower rates for higher k : the last 3 simulated rate values for $k=5$, 4 for $k=8$, and 5 for both $k=10$ and $k=15$

$T_{\text{coh}}/4$. In addition, dummy qubits in the Z basis are completely unaffected by the dephasing noise of the memory. At a sufficiently high improvement factor such as $k = 15$, the residual decoherence and gate noise accumulated over this bounded time are not enough to push the failure probability above the threshold. The practical cost of this tradeoff, however, is severe. At low backbone rates, the fraction of rounds that survive the cutoff becomes vanishingly small, meaning that the time required to complete a single successful test round grows rapidly. This places a practical limit on the backbone rates that can be explored in simulation and would similarly constrain any experimental implementation of the protocol in this regime. The exact extent of this effect is visualized in Figure A.1. In practice, this limitation could be mitigated by relaxing the cutoff to a value larger than $T_{\text{coh}}/2$. A less strict cutoff would allow more rounds to complete at the cost of increased decoherence per round, potentially shifting the balance in regimes where the current cutoff is overly conservative. However, optimizing the cutoff introduces an additional parameter that would need to be tuned carefully for each improvement factor and backbone rate, significantly expanding the parameter space while also needing to consider the resulting rate-fidelity trade-off. Such an optimization is beyond the scope of this work and is left for future investigation and is explicitly discussed in Chapter 5.

Minimum rate requirements

Figure 4.4d summarizes the crossing points extracted from the previous figures, corresponding to the backbone rates at which the protocol metrics first reach the adapted threshold values, thereby addressing RQ3. It therefore shows the minimum backbone entanglement rate required to achieve the adapted target for each protocol as a function of the improvement factor.

For QKD, the two curves corresponding to the $R_s \geq 5$ Hz and $R_s \geq 10$ Hz targets are separated by a nearly constant factor of two across all simulated improvement factors. This confirms that, in the regime where the backbone rate is the dominant bottleneck, the required backbone rate scales linearly with the target secret key rate. The implication is practical: a single curve can be used to extrapolate the backbone requirement for any desired secret key rate, as long as the required backbone rate remains much slower than the metropolitan entanglement generation and BSM times. Beyond this regime the proportionality breaks down, as the metropolitan operations begin to limit the achievable throughput regardless of the backbone rate.

For CHSH, the same logic applies in principle, and the two target curves ($p_{\text{win}} \geq 0.76$ and $p_{\text{win}} \geq 0.80$) would be expected to differ by a factor of roughly five based on the quantum fraction argument outlined in the adapted metrics discussion. In practice, however, the observed ratio is somewhat larger, and the discrepancy grows at lower improvement factors. This can be understood by noting that at lower k , the metropolitan entanglement generation times are no longer negligible compared to the backbone waiting time. Looking at Figure 4.1, at an improvement factor of 5 the DC entanglement protocol operates at a rate of roughly 250 Hz, which is comparable to the backbone rates required for the $p_{\text{win}} \geq 0.80$ target. The end-to-end entanglement time is therefore no longer dominated by the backbone alone, and the simple linear scaling between backbone rate and quantum fraction breaks

down. As the improvement factor increases and metropolitan times become faster, this effect diminishes and the ratio between the two curves moves closer to the expected factor of five.

For VBQC, the data is more limited than for the other two protocols. At the low end of the improvement-factor range, the curve connects to the perfect-backbone analysis. As k decreases toward the minimum value of $k \approx 6.42$, even minimal decoherence becomes sufficient to cross the threshold, and the required backbone rate grows rapidly. This trend is visible in the figure as the dashed extrapolation toward diverging rates. In the opposite direction, however, no crossing points can be extracted at $k = 15$, while they still exist at $k = 10$. This means that somewhere between these two values, the required backbone rate drops to zero and the curve terminates. As argued previously, at sufficiently high improvement factors the decoherence accumulated over the bounded conditional memory time of $T_{\text{coh}}/4$ is simply not enough to make the test rounds fail, regardless of the backbone rate. While the exact location of this asymptote cannot be determined from the available data, it can be constrained to lie in the range $10 < k < 15$.

4.4 Full backbone characterization

In the previous two sections, the backbone fidelity and rate were varied independently while the other was kept ideal. We now allow both to be imperfect simultaneously, mapping out the region in backbone parameter space where each protocol can be successfully implemented. Since the parameter space has now grown by one, the protocol performance itself can no longer be easily displayed directly. Instead, we show threshold contours that separate the feasible region from the infeasible one. Any point above a contour satisfies the corresponding target metric, while points below it do not. The adapted metrics introduced in the quality backbone analysis carry over to this setting. Points from both previous analyses serve as references. The crossing points from the quality-backbone analysis appear as the $F = 1$ intercepts at the top of each figure, while the fast-backbone results correspond to the high-rate limits on the right. This allows us to answer the fourth and final research question: For a given improved parameter set, in what region of the backbone (R, F) space can each of the three protocols be successfully implemented?

Protocol-specific backbone characterization

Figure 4.5 (a)-(c) shows the threshold contours for each protocol individually, for several values of the improvement factor and target metric, thereby addressing RQ4.

For QKD (Figure 4.5a), the contours reveal a clear tradeoff between backbone fidelity and rate. At high backbone fidelities, a modest decrease in fidelity can be offset by increasing the backbone rate, as the secret fraction remains comfortably above zero and the reduced quality is compensated by higher throughput. This is reflected in the moderate slope of the contours in this regime, which shows that losses in fidelity can be compensated by increases in rate while maintaining the same secret-key-rate target. As the backbone fidelity decreases further, however, the secret fraction approaches zero and progressively larger rate increases are needed to maintain the same secret key rate which becomes apparent by the slopes flat-

4.4. Full backbone characterization

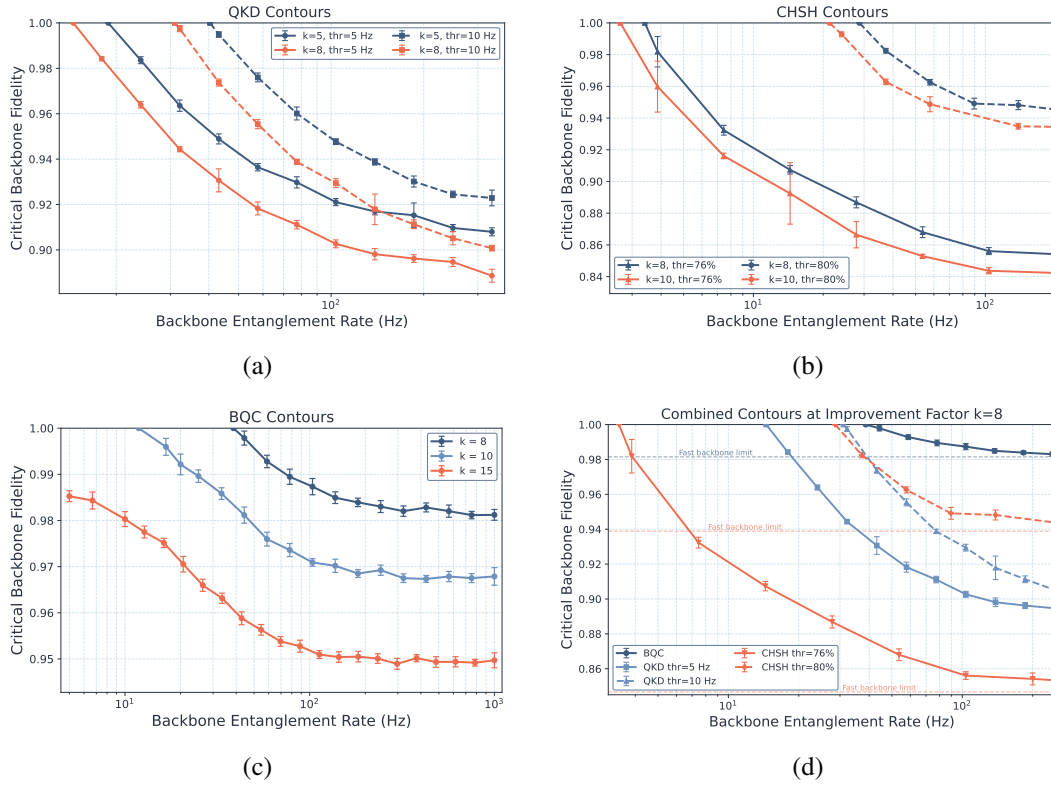


Figure 4.5: Full backbone characterization. (a)-(c) show the threshold contours in backbone parameter space for QKD, CHSH, and VBQC, respectively, for several values of the metropolitan hardware improvement factor k and, where applicable, several target metrics. Each contour separates the feasible and infeasible regions for the corresponding protocol. Points above a contour satisfy the target metric, while points below do not. The $F = 1$ intercepts at the top reproduce the minimum backbone-rate requirements obtained in the quality-backbone analysis, while the high-rate limits on the right recover the minimum backbone-fidelity requirements from the fast-backbone analysis. (d) combines the contours of all three protocols at a fixed improvement factor of $k = 8$, allowing a direct comparison of their backbone requirements. Dashed lines indicate the fast-backbone limiting fidelities where applicable.

tening. In the idealized limit where the backbone rate directly sets the raw key rate, all thresholds for the same k would therefore approach the same critical backbone fidelity at which the secret fraction vanishes, explaining why the contours flatten and begin to group together as the required rate increases.

As the backbone rate increases, however, the raw key rate becomes less dominated by the backbone entanglement rate and increasingly limited by the metropolitan entanglement generation and BSM times. The contours therefore flatten not only because the secret fraction approaches zero, but also because further increases in the backbone rate yield diminishing returns in the raw key rate. For the same reason, contours with the same k do

not converge exactly to a single critical backbone fidelity. Nevertheless, the trend remains clearly visible, as increasing the backbone rate causes contours with the same k , but different target metrics, to group together. This also explains the crossing point observed between the $k = 5, R_s \geq 10, \text{Hz}$ contour and the $k = 8, R_s \geq 5, \text{Hz}$ contour.

For CHSH (Figure 4.5b), the contours illustrate that demonstrating any quantum advantage places relatively mild requirements on the backbone, while achieving a substantially higher winning probability demands considerably more. The contours for $p_{\text{win}} \geq 0.76$ lie well below those for $p_{\text{win}} \geq 0.80$, reflecting the steep increase in backbone quality needed to move from a marginal to a convincing violation of the classical bound. One that also cannot be offset by rate, as opposed to QKD. It should be noted that the uncertainty bands in this figure are less reliable than in the other two protocols. The threshold contours are obtained by bisecting and linearly interpolating between simulated points, and the tolerance for this bisection was not adapted per protocol. For CHSH, the metric varies slowly with backbone fidelity, meaning that the bracketing points can end up very close in metric value. This was not an issue for VBQC due to the vast difference in sensitivity as can be seen in Figure A.3. In such cases, small statistical fluctuations can produce unreliable slopes in the linear interpolation of Equation 3.26, leading to inflated uncertainty estimates. The contour positions themselves remain trustworthy, as they were still interpolated, but the error bands should be interpreted with caution.

For VBQC (Figure 4.5c), the cutoff stagnation observed in the quality backbone analysis reappears. At $k = 15$, nearly the entire region above $F \approx 0.99$ satisfies the threshold regardless of the backbone rate, although as discussed previously, operating at very low rates would not be experimentally feasible. The overall shape of the contours bears a strong resemblance to the failure probability curves of Figure 4.4c. This is not coincidental, since a lower failure probability at unit backbone fidelity corresponds to a larger margin before the threshold is crossed, which translates directly into a higher tolerance for backbone noise and therefore a lower required backbone fidelity. The ordering of the contours thus mirrors the ordering of the quality backbone curves.

Combined backbone characterization

Figure 4.5d combines the threshold contours of all three protocols at a fixed improvement factor of $k = 8$. This value was chosen because it provides sufficient margin above the minimum improvement factor for VBQC ($k \approx 6.42$) to allow a meaningful investigation, while still representing a relatively modest hardware improvement. The fast backbone limits obtained from the analysis in Section 4.2 are, if applicable, indicated by dashed lines on the right side of the figure, while the quality backbone crossing points from Section 4.3 appear as the $F = 1$ intercepts at the top.

At this improvement factor, demonstrating a win rate $p_{\text{win}} \geq 0.76$ places the mildest requirements on the backbone by a comfortable margin. The next easiest target is QKD at $R_s \geq 5 \text{ Hz}$, followed by a group of two: QKD at $R_s \geq 10 \text{ Hz}$ and CHSH at $p_{\text{win}} \geq 0.80$ start at similar minimum backbone rates in the high-fidelity regime. As the backbone fidelity decreases, however, the two diverge. The QKD contour drops more strongly, reflecting the fact that a reduced secret fraction can be effectively compensated by a higher backbone

rate. The CHSH contour, by contrast, steepens less prominently, as the winning probability is less easily recovered through increased throughput alone. Achieving a strong quantum advantage in the CHSH game therefore requires a higher quality backbone than generating secret keys at a comparable rate, making $p_{\text{win}} \geq 0.80$ the second most demanding target after VBQC.

VBQC is by far the most demanding of the three protocols. Its contour lies above those of all other targets across the entire parameter space, requiring both higher backbone fidelity and higher backbone rate. While one might argue that this is simply a consequence of the adapted metrics for QKD and CHSH being set above their respective minimum thresholds, it is worth emphasizing that the VBQC contour still corresponds to the bare minimum metric of $p_{\text{fail}} < 0.25$. Even with stricter targets imposed on the other two protocols, VBQC remains the hardest to implement. Achieving secure blind quantum computation on this architecture therefore places substantially stronger demands on the backbone than either key distribution or performing quantum based communication strategies.

4.5 Further discussion

The preceding sections addressed the four research questions by determining when each benchmark protocol becomes feasible, how the required backbone fidelity and rate depend on metropolitan hardware quality, and how these requirements combine into feasible operating regions in backbone parameter space. Taken together, they provide a broad picture of the backbone demands of QKD, CHSH, and VBQC on the considered architecture.

Before drawing final conclusions, however, it is important to examine how strongly these results depend on the specific modeling assumptions used throughout this work. In particular, three points merit further discussion: the robustness of the conclusions under nearby setup variations, the extent to which the considered improvement factors are experimentally realistic, and the interpretation and limitations of the global improvement-factor framework itself.

On the robustness of the results to setup variations

The results in this thesis were obtained for one specific metropolitan setup where processor clients connect to the border-nodes through midpoint-heralded entanglement generation, while measurement-only clients connect through direct optical links with high-efficiency detectors. These assumptions provide a concrete and physically motivated baseline, but they are not unique. It is therefore useful to ask whether nearby setup variations would change the main conclusions. Here, we consider two examples, namely moving the heralding station for processing-node clients to the border node and reducing the detector efficiency of MO clients.

Throughout this thesis, the heralding station for entanglement generation between a processing node client and a border-node was assumed to be placed at the midpoint of the link. A more practical metropolitan deployment may instead place this station at the border-node, allowing multiple clients to share the same infrastructure more naturally [30]. The main effect of this asymmetric placement is an increase in the classical communication

4. RESULTS AND DISCUSSION

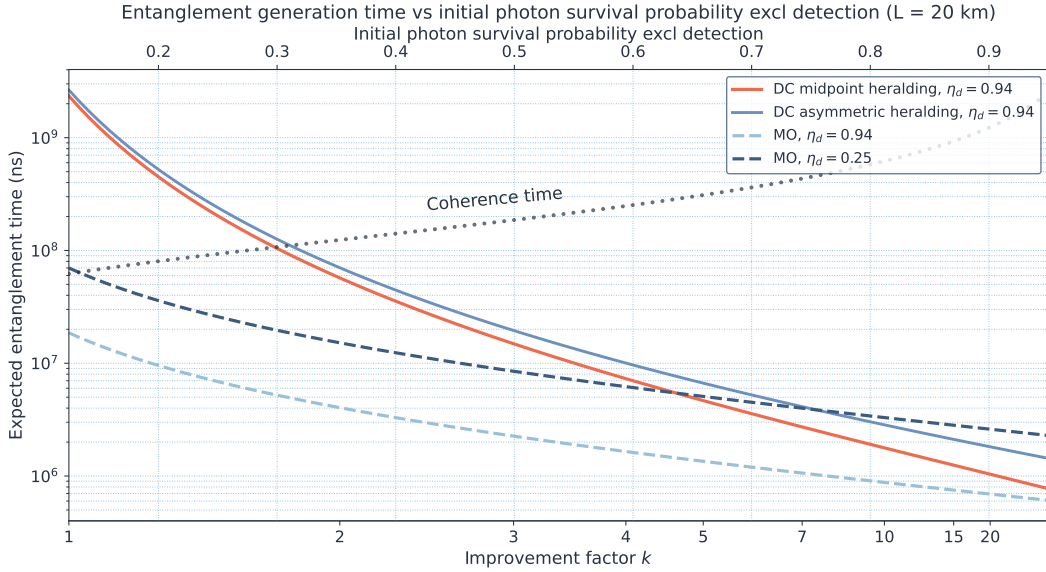


Figure 4.6: Average metropolitan entanglement-generation times as a function of the hardware improvement factor k . The plot compares double-click (DC) entanglement generation between two quantum processors for midpoint and asymmetric heralding-station placement with the photon-detection time of a measurement-only (MO) client for different detector efficiencies.

time per entanglement attempt. For a 20 km client-to-border-node separation, the relevant communication distance increases from 20 km to 40 km, corresponding to an increase from $100\mu\text{s}$ to $200\mu\text{s}$ in the simulated fiber.

As shown in Figure 4.6, this makes the asymmetric double-click link slower than the midpoint configuration, especially at larger k , where the emission duration has been improved and the fixed communication time becomes more important. However, the protocol-level impact is expected to remain limited. In the regimes where CHSH and VBQC become feasible, the metropolitan entanglement-generation time is already small compared with the coherence time. Consequently, the additional waiting time contributes only weakly to memory decoherence. The midpoint assumption should therefore be understood as optimistic, but not overly restrictive. It may affect exact threshold values, particularly close to the onset of feasibility, but is not expected to change the qualitative conclusions for CHSH or VBQC.

A second relevant variation is the detector efficiency of the MO-clients. The baseline simulations assume a high detector efficiency of $\eta_d = 0.94$, whereas cheaper telecom-range detectors may have substantially lower efficiencies. Figure 4.6 therefore also shows the effect of reducing and not improving the detector efficiency to $\eta_d = 0.25$ corresponding to commercially available architectures also specializing in single photon telecom photon detection [23].

This change directly lowers the success probability of the MO-link and therefore increases the expected photon-detection time by roughly the inverse ratio of the detector efficiencies. Its importance is protocol dependent. For QKD, which uses two MO-clients, the

reduced detector efficiency can noticeably affect the achievable raw key rate whenever the metropolitan links become comparable to, or slower than, the backbone. This is especially relevant in the quality-backbone analysis, where the backbone rate is varied explicitly and can enter the same range as the metropolitan link rate.

For VBQC, the effect is expected to be much smaller. In the regime where VBQC becomes feasible, the relevant metropolitan timescales remain far below the coherence time even for the reduced detector efficiency shown in Figure 4.6. The dominant limitation therefore remains the quality and timing of the entanglement delivered through the backbone, rather than the MO-client detection time. The same reasoning also suggests that the main CHSH conclusions are insensitive to this particular variation, since CHSH does not rely on MO-clients.

Overall, these examples indicate that the conclusions of this thesis are reasonably robust to nearby changes in the metropolitan setup, but not in a protocol-independent way. Variations that only weakly affect memory decoherence are unlikely to change the CHSH and VBQC conclusions substantially, whereas changes that directly reduce the MO-link rate can matter for QKD whenever the metropolitan links become a relevant bottleneck.

On the interpretation of the improvement factor

Throughout this thesis, hardware progress has been described through a single global improvement factor k . This was primarily a modeling choice because assigning an independent improvement factor to every hardware parameter would make the parameter space too large to explore systematically. The global- k framework instead reduces the problem to a one-dimensional scan, allowing the backbone requirements to be studied as the metropolitan hardware is improved in a controlled way.

However, k should not be interpreted as a literal forecast of experimental progress. In practice, different hardware parameters will improve at very different rates. To illustrate this, Table 4.1 compares the baseline values used in this thesis with optimistic values taken either from demonstrated results on the same general trapped-ion platform or from targets (5-10 years) in consultation with the relevant experiment groups. For each parameter, the table also lists the effective improvement factor obtained by inverting the probability-of-no-imperfection mapping,

$$k_i = \frac{\ln(p_{\text{NI}}(b_i))}{\ln(p_{\text{NI}}(x_i))}, \quad (4.4)$$

where b_i is the baseline value and x_i is the optimistic value.

The table shows that the effective values of k differ strongly between parameters. Some improvements fall within the range studied in this thesis. For example, the optimistic single-qubit gate fidelity, emission fidelity, and initial photon survival probability correspond to effective improvement factors of order $k \sim 2-5$. These values are directly comparable to the regimes explored in the main results and therefore suggest that parts of the required parameter improvement may be experimentally plausible.

Other parameters, however, map to much larger effective improvement factors. A coherence-time target of 4 s corresponds to $k \approx 4160$, while a tenfold reduction in emis-

4. RESULTS AND DISCUSSION

Table 4.1: Comparison between the baseline hardware parameters used in this thesis and optimistic values taken from demonstrated improvements or near-term target values. The final column gives the effective improvement factor required to map the baseline value to the optimistic value within the probability-of-no-imperfection framework.

| Parameter | Baseline value | Optimistic value | Effective k |
|-------------------------------------|----------------|------------------|---------------|
| Coherence time (ns) | 62,000,000 | 4×10^9 | 4160 |
| Single-qubit gate fidelity | 0.99 | 0.998 [33] | 5.04 |
| MS gate fidelity | 0.95 | 0.997 [33] | 17.2 |
| Initial photon survival probability | 0.1245 | 0.375 | 2.12 |
| Emission fidelity | 0.974 | 0.995 | 5.27 |
| Emission duration (ns) | 666,666 | 66,666 | 100 |
| Visibility | 0.89 | – | – |
| η_{penalty} | 0.12 | 0.2 | 1.32 |

sion duration corresponds to $k = 100$, placing both improvements well beyond the k -range explored in this thesis. These values should not be interpreted as meaning that the corresponding hardware has improved by the same amount in a physically comparable sense. Rather, they show that the probability of no-imperfection mapping is much more sensitive for some parameter classes than for others.

This highlights the main limitation of the global improvement factor. It is useful for organizing the parameter scan and identifying broad performance trends, but it does not capture the way experimental progress is actually expected to occur. Hardware development is parameter-specific, where gate fidelities, photonic efficiencies, coherence times, and emission durations may improve on very different timescales. A single value of k therefore represents an effective hardware maturity level, not a concrete future device.

This limitation is partly softened by the fact that improvements in different parameters can compensate for one another at the protocol level. For example, the initial photon survival probability is itself a product of several underlying efficiencies, so unequal improvements in collection efficiency, frequency conversion, and detector efficiency can still produce the same effective success probability. More generally, weaker progress in one parameter may be offset by stronger progress in another, although the extent of this compensation is protocol-dependent. However, for drastic changes such as those expected for the coherence time, such a compensation argument becomes difficult to make, since the system may enter an entirely different operating regime than investigated here.

This limitation affects how the results should be interpreted, but it does not undermine the trends identified in the analysis. The purpose of a global- k was not to predict a unique future parameter set, but to identify how the backbone requirements change as the metropolitan hardware improves. In that sense, the framework remains useful. It shows when each protocol first becomes feasible, how the required backbone fidelity and rate relax with improved local hardware, and which regimes of backbone parameter space are relevant for each application.

The simulation framework also remains useful beyond the specific baseline values cho-

sen in this thesis. If updated parameter estimates become available, the analysis can be repeated by simply updating the corresponding entries in the hardware parameter set, for example by replacing the baseline coherence time with an optimistic projected value. In this way, the same pipeline can quickly reassess the protocol thresholds and backbone requirements for an updated hardware set. More targeted future work could then go one step further by introducing parameter-specific scans, targeted improvements of selected hardware components, or protocol-dependent optimization of the most relevant parameter groups.

Chapter 5

Conclusions and future work

5.1 Contributions

This thesis investigated the backbone requirements for implementing three quantum communication protocols on an intercity trapped-ion network architecture. The main contributions are as follows.

First, a simulation framework was developed using the NetSquid discrete-event simulation platform that models a four-node quantum network with physically detailed metropolitan components and an abstract long-distance backbone. The framework, whose simulation code is available online [32], is modular: the baseline hardware parameters, the backbone characteristics, and the choice of protocol can all be varied independently, allowing the entire analysis to be reproduced for different hardware platforms or updated parameter sets with minimal modification.

Second, three benchmark protocols, quantum key distribution, the CHSH game, and verified blind quantum computation, were implemented on this architecture, covering all three possible combinations of end-node types (two measurement-only clients, two quantum nodes, and one of each).

Third, a systematic analysis was carried out across the four research questions formulated in the introduction. The minimum hardware improvement factors were identified (RQ1), the minimum backbone fidelity was mapped as a function of the improvement factor (RQ2), the minimum backbone rate was characterized using adapted metrics where necessary (RQ3), and the full feasible region in the backbone (R, F) parameter space was determined for each protocol (RQ4). Together, these results provide a comprehensive picture of the interplay between metropolitan hardware quality and long-distance backbone requirements.

5.2 Conclusions

This thesis investigated what rate and fidelity a long-distance entanglement-distribution backbone must provide in order to support QKD, the CHSH game, and VBQC on an intercity trapped-ion network architecture. The results show that the answer depends strongly

on the application as the three protocols impose qualitatively different requirements on the backbone and respond differently to improvements in metropolitan hardware.

The first research question asked what minimum hardware improvement factor is required to implement each protocol assuming a perfect backbone. This question admits the most direct numerical answer. QKD becomes feasible at $k = 1.85 \pm 0.03$, the CHSH game at $k = 2.86 \pm 0.06$, and VBQC at $k = 6.42 \pm 0.09$. This ordering reflects the different end-node requirements of the protocols. QKD uses two measurement-only clients and is therefore least affected by slow two-photon heralded entanglement generation and fidelity losses. The CHSH game requires two quantum processing nodes, while VBQC combines a measurement-only client with a quantum server and requires the successful preparation and verification of a multi-qubit graph state, making it the most demanding of the three protocols.

The second research question asked what minimum backbone fidelity is required for a given improved hardware parameter set. Unlike RQ1, this question does not have a simple numerical answer. The answer is the set of threshold curves shown in Figure 4.3d. For all three protocols, improving the metropolitan hardware relaxes the required backbone fidelity, but with diminishing returns as k increases. QKD benefits strongly already at moderate improvement factors and approaches its analytical perfect-metropolitan-hardware limit relatively quickly. The CHSH game initially requires high backbone fidelity near its feasibility threshold, but at larger k its required fidelity can fall below that of QKD because the intrinsic perfect-hardware fidelity bound for CHSH is lower. VBQC remains the most fidelity-sensitive protocol due to its inherent nature.

The third research question asked what minimum backbone rate is required for a given improved hardware parameter set. Here, the original feasibility thresholds do not all define meaningful rate thresholds. Under the backbone-first ordering used in this thesis, reducing the backbone rate affects how often entangled states are available, but for QKD and CHSH it does not directly degrade the quality of the states once they are used. Therefore, practical target metrics were introduced, namely secret-key-rate targets for QKD and stronger winning-probability targets for CHSH. With these adapted targets, Figure 4.4d shows that the required backbone rate decreases as the hardware improves, but the dependence is again protocol- and target-specific. For VBQC, the rate dependence is strongly shaped by the memory cutoff. At sufficiently high k , the failure probability can remain below threshold even at low backbone rates, but only because most slow rounds are discarded. Thus, low-rate feasibility in this regime should not be interpreted as practical efficiency, since the time required to obtain a successful round can become very large.

The fourth research question combined both effects and asked which regions of the full backbone (R, F) parameter space allow successful protocol implementation. The answer is given by the threshold contours in Figure 4.5. These contours combine the limiting cases of the previous two research questions. Their high-rate limits reproduce the minimum-fidelity requirements from the fast-backbone analysis, while their $F = 1$ intercepts reproduce the minimum-rate requirements from the quality-backbone analysis. The resulting feasible regions show that rate and fidelity can compensate for each other only to a limited and protocol-dependent extent. For QKD, a lower secret fraction caused by reduced fidelity can partly be compensated by a higher backbone rate. For CHSH, increasing the rate cannot

compensate for insufficient entanglement quality in the same way, especially when aiming for a stronger violation of the classical bound. VBQC is the most restrictive case across the investigated parameter space, requiring both high fidelity and sufficiently high rate.

Overall, the central conclusion is that intercity quantum networks do not have a universal backbone requirement. The required rate and fidelity depend on both the target application and the maturity of the metropolitan hardware. QKD is the least demanding benchmark, CHSH becomes relatively tolerant to backbone infidelity once local hardware improves, and VBQC remains the most demanding protocol across the investigated parameter space.

Finally, these conclusions should be interpreted within the scope of the modeling choices made in this thesis. The global improvement factor provides a useful way to compare different hardware improvements on a common scale, but it should not be read as a literal forecast of future hardware progress, especially for time-like parameters such as coherence time and emission duration. Similarly, the protocol implementations were not optimized exhaustively. Cutoff strategies for CHSH and VBQC, as well as the basis choice in QKD, could be refined further.

5.3 Future work

This work opens several directions for further investigation.

Alternative baseline parameter sets and platforms. The simulation framework developed in this thesis can be extended to other repeater platforms with relatively modest effort. Most notably, nitrogen-vacancy centers in diamond represent a well-studied alternative to trapped ions and are already supported within the NetSquid ecosystem. Implementing an NV-center-based parameter set would allow a direct comparison between platforms, potentially revealing regimes in which one outperforms the other for specific protocols or backbone configurations. Beyond comparing platforms, the framework is equally suited to tracking the progress of a single platform over time. As the Innsbruck trapped-ion system continues to improve, updated baseline parameters can be substituted into the existing simulation to reassess how close the hardware is to the regime where each protocol becomes feasible. Each such update would immediately propagate through the entire analysis pipeline, from the minimum improvement factors to the full backbone characterization.

Parameter-specific improvement factors. The uniform improvement factor used in this thesis treats all hardware parameters equally, but experimental progress is rarely so balanced. A natural extension would be to assign independent improvement factors to individual parameters or groups of parameters. This would enable several types of investigation, including identifying which parameters have the greatest leverage on protocol performance, finding the optimal allocation of a fixed total improvement budget across parameters, and studying how an improvement in one quantity, such as coherence time, can compensate for stagnation in another, such as gate fidelity. Such an analysis would substantially expand the parameter space, but could yield practically valuable guidance for experimentalists prioritizing their development efforts and for assessing network requirements. Similar

requirements-driven analyses have been carried out for near-term quantum-network applications in Refs. [39, 17, 4].

Protocol investigation. An interesting aspect is in optimizing or investigating the protocols themselves. The fixed cutoff of $T_{\text{coh}}/2$ used throughout this work leaves considerable room for improvement. For the CHSH game, the data from the fast backbone analysis could be used to fit an effective depolarizing parameter that captures the total accumulated noise from all sources with negligible memory effects. Combined with a model of the memory decoherence, this would allow an estimation of the optimal moment at which a stored qubit no longer contributes a net benefit to the winning probability, providing a principled basis for setting the cutoff.

For VBQC, the cutoff problem is substantially richer. Here, the cutoff introduces a trade-off between fidelity and rate. A shorter cutoff generally reduces the test-round failure probability, but also lowers the probability that all required qubits are delivered successfully, thereby increasing the time per completed round. If the objective is instead to minimize the total time required to obtain a failure probability below 0.25 with some statistical certainty, the optimal cutoff will depend on the underlying hardware parameters and could be optimized for. This optimization is further complicated by the sequential delivery of the five qubits. The optimal cutoff for a given qubit may depend on the arrival times of the preceding qubits. A dynamic cutoff strategy could exploit this structure. For example, if the first four qubits arrived unusually quickly, a longer cutoff could be tolerated for the fifth. Such an investigation would address the stagnation effect identified in this thesis, which suggests that the current fixed cutoff is overly conservative at high improvement factors.

Bibliography

- [1] Silvestre Abruzzo, Sylvia Bratzik, Nadja K. Bernardes, Hermann Kampermann, Peter van Loock, and Dagmar Bruß. Quantum repeaters and quantum key distribution: Analysis of secret-key rates. *Phys. Rev. A*, 87:052315, May 2013. doi: 10.1103/PhysRevA.87.052315. URL <https://link.aps.org/doi/10.1103/PhysRevA.87.052315>.
- [2] Guus Avis. Netsquid trapped-ions snippet. <https://docs.netsquid.org/snippets/netsquid-trappedions/>, 2023.
- [3] Guus Avis, Tim Coopmans, Axel Dahlberg, Francisco Ferreira da Silva, Hana Jirovská, David Maier, Julian Rabbie, and Matthew Skrzypczyk. Netsquid magic snippet. <https://docs.netsquid.org/snippets/netsquid-magic/>, 2020.
- [4] Guus Avis, Francisco Ferreira da Silva, Tim Coopmans, Axel Dahlberg, Hana Jirovská, David Maier, Julian Rabbie, Ariana Torres-Knoop, and Stephanie Wehner. Requirements for a processing-node quantum repeater on a real-world fiber grid. *npj Quantum Information*, 9(1), October 2023. ISSN 2056-6387. doi: 10.1038/s41534-023-00765-x. URL <http://dx.doi.org/10.1038/s41534-023-00765-x>.
- [5] Koji Azuma, Sophia E. Economou, David Elkouss, Paul Hilaire, Liang Jiang, Hoi-Kwong Lo, and Ilan Tzitrin. Quantum repeaters: From quantum networks to the quantum internet. *Rev. Mod. Phys.*, 95:045006, Dec 2023. doi: 10.1103/RevModPhys.95.045006. URL <https://link.aps.org/doi/10.1103/RevModPhys.95.045006>.
- [6] Sean D. Barrett and Pieter Kok. Efficient high-fidelity quantum computation using matter qubits and linear optics. *Physical Review A*, 71(6), June 2005. ISSN 1094-1622. doi: 10.1103/physreva.71.060310. URL <http://dx.doi.org/10.1103/PhysRevA.71.060310>.
- [7] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, December 2014. ISSN 0304-3975. doi: 10.1016/j.tcs.2014.05.025. URL <http://dx.doi.org/10.1016/j.tcs.2014.05.025>.

- [8] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Phys. Rev. Lett.*, 70:1895–1899, Mar 1993. doi: 10.1103/PhysRevLett.70.1895. URL <https://link.aps.org/doi/10.1103/PhysRevLett.70.1895>.
- [9] Charles H. Bennett, David P. DiVincenzo, Peter W. Shor, John A. Smolin, Barbara M. Terhal, and William K. Wootters. Remote state preparation. *Physical Review Letters*, 87(7), July 2001. ISSN 1079-7114. doi: 10.1103/physrevlett.87.077902. URL <http://dx.doi.org/10.1103/PhysRevLett.87.077902>.
- [10] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 517–526. IEEE, October 2009. doi: 10.1109/focs.2009.36. URL <http://dx.doi.org/10.1109/FOCS.2009.36>.
- [11] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. Bell nonlocality. *Reviews of modern physics*, 86(2):419–478, 2014.
- [12] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2), May 2019. ISSN 1931-9401. doi: 10.1063/1.5088164. URL <http://dx.doi.org/10.1063/1.5088164>.
- [13] J. Chang, J. W. N. Los, J. O. Tenorio-Pearl, N. Noordzij, R. Gourgues, A. Guardiani, J. R. Zichi, S. F. Pereira, H. P. Urbach, V. Zwiller, S. N. Dorenbos, and I. Esmail Zadeh. Detecting telecom single photons with 99.5207±0.5% system detection efficiency and high time resolution. *APL Photonics*, 6(3): 036114, March 2021. ISSN 2378-0967. doi: 10.1063/5.0039772. URL <https://doi.org/10.1063/5.0039772>. eprint: https://pubs.aip.org/aip/app/article-pdf/doi/10.1063/5.0039772/14572142/036114.1_online.pdf.
- [14] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23:880–884, Oct 1969. doi: 10.1103/PhysRevLett.23.880. URL <https://link.aps.org/doi/10.1103/PhysRevLett.23.880>.
- [15] Tim Coopmans, Robert Knegjens, Axel Dahlberg, David Maier, Loek Nijsten, Julio de Oliveira Filho, Martijn Papendrecht, Julian Rabbie, Filip Rozpedek, Matthew Skrzypczyk, Leon Wubben, Walter de Jong, Damian Podareanu, Ariana Torres-Knoop, David Elkouss, and Stephanie Wehner. Netsquid, a network simulator for quantum information using discrete events. *Communications Physics*, 4(1), July 2021. ISSN 2399-3650. doi: 10.1038/s42005-021-00647-8. URL <http://dx.doi.org/10.1038/s42005-021-00647-8>.
- [16] Francisco Ferreira da Silva and Stephanie Wehner. Entanglement improves coordination in distributed systems, 2026. URL <https://arxiv.org/abs/2602.04588>.

-
- [17] Francisco Ferreira da Silva, Guus Avis, Joshua A. Slater, and Stephanie Wehner. Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber, 2023. URL <https://arxiv.org/abs/2303.03234>.
- [18] Vincent Danos, Elham Kashefi, and Prakash Panangaden. The measurement calculus, 2007. URL <https://arxiv.org/abs/0704.1263>.
- [19] Dawei Ding and Liang Jiang. Coordinating decisions via quantum telepathy, 2025. URL <https://arxiv.org/abs/2407.21723>.
- [20] David Elkouss. Key distribution and the chsh game, 2018.
- [21] M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest, and H. J. Briegel. Entanglement in graph states and its applications, 2006. URL <https://arxiv.org/abs/quant-ph/0602096>.
- [22] Yuexun Huang, Francisco Salces–Carcoba, Rana X. Adhikari, Amir H. Safavi-Naeini, and Liang Jiang. Vacuum beam guide for large scale quantum networks. *Physical Review Letters*, 133(2), July 2024. ISSN 1079-7114. doi: 10.1103/physrevlett.133.020801. URL <http://dx.doi.org/10.1103/PhysRevLett.133.020801>.
- [23] ID Quantique. Id qube series nir free-running, 2026. URL <https://www.idquantique.com/quantum-detection-systems/products/id-qube-nir-free-running/>. Accessed 2026-04-14.
- [24] V. Krutyanskiy, M. Canteri, M. Meraner, J. Bate, V. Krcmarsky, J. Schupp, N. Sangouard, and B. P. Lanyon. Telecom-wavelength quantum repeater node based on a trapped-ion processor. *Phys. Rev. Lett.*, 130:213601, May 2023. doi: 10.1103/PhysRevLett.130.213601. URL <https://link.aps.org/doi/10.1103/PhysRevLett.130.213601>.
- [25] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D.A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B.P. Lanyon, and T.E. Northup. Entanglement of trapped-ion qubits separated by 230 meters. *Physical Review Letters*, 130(5), February 2023. ISSN 1079-7114. doi: 10.1103/physrevlett.130.050803. URL <http://dx.doi.org/10.1103/PhysRevLett.130.050803>.
- [26] Viktor Krutyanskiy, Martin Meraner, J. Schupp, Vojtech Krcmarsky, Helene Hainzer, and Ben Lanyon. Light-matter entanglement over 50 km of optical fibre. *npj Quantum Information*, 5:1–5, 08 2019. doi: 10.1038/s41534-019-0186-3.
- [27] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, Mar 2003. doi: 10.1103/RevModPhys.75.281. URL <https://link.aps.org/doi/10.1103/RevModPhys.75.281>.

BIBLIOGRAPHY

- [28] Dominik Leichtle, Luka Music, Elham Kashefi, and Harold Ollivier. Verifying bqp computations on noisy devices with minimal overhead. *PRX Quantum*, 2:040302, Oct 2021. doi: 10.1103/PRXQuantum.2.040302. URL <https://link.aps.org/doi/10.1103/PRXQuantum.2.040302>.
- [29] Hoi-Kwong Lo, Marcos Curty, and Bing Qi. Measurement-device-independent quantum key distribution. *Physical Review Letters*, 108(13), March 2012. ISSN 1079-7114. doi: 10.1103/physrevlett.108.130503. URL <http://dx.doi.org/10.1103/PhysRevLett.108.130503>.
- [30] Soubhadra Maiti, Guus Avis, Sounak Kar, and Stephanie Wehner. Requirements for teleportation in an intercity quantum network, 2026. URL <https://arxiv.org/abs/2602.04869>.
- [31] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [32] Tobias Ploeckinger. Code underlying the thesis: Connecting Two Quantum Cities: Understanding the Requirements for Long-Distance Quantum Links. https://gitlab.tudelft.nl/wehner-research/backbone_threeprotocols, 2026.
- [33] I. Pogorelov, T. Feldker, Ch. D. Marciniak, L. Postler, G. Jacob, O. Kriegelsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskii, R. Blatt, P. Schindler, and T. Monz. Compact ion-trap quantum computing demonstrator. *PRX Quantum*, 2:020343, Jun 2021. doi: 10.1103/PRXQuantum.2.020343. URL <https://link.aps.org/doi/10.1103/PRXQuantum.2.020343>.
- [34] Robert Raussendorf and Hans J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86:5188–5191, May 2001. doi: 10.1103/PhysRevLett.86.5188. URL <https://link.aps.org/doi/10.1103/PhysRevLett.86.5188>.
- [35] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J. Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81:1301–1350, Sep 2009. doi: 10.1103/RevModPhys.81.1301. URL <https://link.aps.org/doi/10.1103/RevModPhys.81.1301>.
- [36] David Schwerdt, Lee Peleg, Yotam Shapira, Nadav Priel, Yanay Florshaim, Avram Gross, Ayelet Zalic, Gadi Afek, Nitzan Akerman, Ady Stern, Amit Ben Kish, and Roei Ozeri. Scalable architecture for trapped-ion quantum computing using rf traps and dynamic optical potentials. *Phys. Rev. X*, 14:041017, Oct 2024. doi: 10.1103/PhysRevX.14.041017. URL <https://link.aps.org/doi/10.1103/PhysRevX.14.041017>.
- [37] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997. ISSN 1095-7111. doi: 10.1137/s0097539795293172. URL <http://dx.doi.org/10.1137/s0097539795293172>.

-
- [38] TU Delft. **Quantum repeaters and teleportation.** <https://www.tudelft.nl/over-tu-delft/strategie/vision-teams/quantum-internet/basics-of-quantum-mechanics/quantum-repeaters-and-teleportation>. Accessed: 2025-07-09.
- [39] J van Dam, G Avis, Tz B Propp, F Ferreira da Silva, J A Slater, T E Northup, and S Wehner. Hardware requirements for trapped-ion-based verifiable blind quantum computing with a measurement-only client. *Quantum Science and Technology*, 9(4): 045031, August 2024. ISSN 2058-9565. doi: 10.1088/2058-9565/ad6eb2. URL <http://dx.doi.org/10.1088/2058-9565/ad6eb2>.
- [40] Michał van Hooft. **Netsquid netbuilder snippet.** <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-netbuilder>, 2024.
- [41] Stephanie Wehner, David Elkouss, and Ronald Hanson. Quantum internet: A vision for the road ahead. *Science*, 362(6412):eaam9288, 2018. doi: 10.1126/science.aam9288. URL <https://www.science.org/doi/abs/10.1126/science.aam9288>.
- [42] Feihu Xu, Xiongfeng Ma, Qiang Zhang, Hoi-Kwong Lo, and Jian-Wei Pan. Secure quantum key distribution with realistic devices. *Reviews of Modern Physics*, 92(2), May 2020. ISSN 1539-0756. doi: 10.1103/revmodphys.92.025002. URL <http://dx.doi.org/10.1103/RevModPhys.92.025002>.
- [43] Juan Yin, Yuan Cao, Yu-Huai Li, Sheng-Kai Liao, Liang Zhang, Ji-Gang Ren, Wen-Qi Cai, Wei-Yue Liu, Bo Li, Hui Dai, Guang-Bing Li, Qi-Ming Lu, Yun-Hong Gong, Yu Xu, Shuang-Lin Li, Feng-Zhi Li, Ya-Yun Yin, Zi-Qing Jiang, Ming Li, Jian-Jun Jia, Ge Ren, Dong He, Yi-Lin Zhou, Xiao-Xiang Zhang, Na Wang, Xiang Chang, Zhen-Cai Zhu, Nai-Le Liu, Yu-Ao Chen, Chao-Yang Lu, Rong Shu, Cheng-Zhi Peng, Jian-Yu Wang, and Jian-Wei Pan. Satellite-based entanglement distribution over 1200 kilometers. *Science*, 356(6343):1140–1144, 2017. doi: 10.1126/science.aan3211. URL <https://www.science.org/doi/abs/10.1126/science.aan3211>.

Appendix A

Additional material

A.1 Declaration of AI usage

Given the increasing prevalence of AI tools, as well as the ongoing discussion surrounding their appropriate use in scientific work, I briefly outline the extent to which AI was used throughout this project. AI tools were used in a supportive capacity during this project. Their use included assistance with writing, coding-related tasks, and occasional brainstorming about the formulation and presentation of ideas. However, the scientific choices, implementation of the core protocols, and final interpretation of the results were derived without the use of AI.

- For coding, AI was not used to write the protocols themselves. Its use was limited to tasks such as debugging and generating or refining code for plot styling and visualization.
- For writing, AI was used more extensively as a language and presentation aid. This included spell-checking, improving the flow and coherence of text, and assisting with LaTeX formatting for tables and figures. While the underlying content, arguments, and scientific ideas remain my own, AI contributed to the exact way these were formulated.

Naturally, any AI-generated output was critically assessed before use and handled in accordance with the policy of the research group. The AI systems used during this project were GPT-5.4 and 5.5 for writing and Opus 4.5 and 4.6 for both writing assistance and coding.

A.2 Analytical model for the 5-qubit linear graph

In this appendix, we derive an analytical expression for the test-round success probability of a 5-qubit linear graph state in the presence of single-qubit depolarizing noise. The goal is to obtain a simple model for how the depolarizing channel affects the probability that a test round is accepted.

Noise model and basic assumptions

We consider the single-qubit depolarizing channel introduced in Section 3.2.2,

$$\mathcal{D}_p^{(1)}(\rho) = (1-p)\rho + p\frac{I_2}{2}, \quad (\text{A.1})$$

so that with probability p the qubit is replaced by the maximally mixed state. Since

$$I = |0\rangle\langle 0| + |1\rangle\langle 1| = |\theta\rangle\langle\theta| + |\theta^\perp\rangle\langle\theta^\perp|, \quad (\text{A.2})$$

the maximally mixed contribution can be interpreted, in the relevant preparation basis, as a flip to the orthogonal state with probability $p/2$.

Throughout this derivation, we assume that the depolarizing channel acts on each qubit before the controlled- Z gates are applied to generate the linear graph state. Under this assumption, each qubit independently flips to its orthogonal state with probability $p/2$, and remains unchanged with probability $1 - p/2$.

Effect of a flip on dummy and trap qubits

A flip on a trap qubit directly changes its measurement outcome and therefore causes the test round to fail. A flip on a dummy qubit has an indirect effect: after the controlled- Z gates are applied, it induces a Z error on each neighboring trap qubit. Since

$$|\theta^\perp\rangle = Z|\theta\rangle, \quad (\text{A.3})$$

such an induced Z error flips the expected measurement outcome of a neighboring trap and therefore also leads to failure.

The only nontrivial cases are those in which multiple flips occur and cancel each other. In particular, if both a dummy qubit and a neighboring trap qubit are flipped, the induced Z error from the dummy can compensate the trap flip, since $ZZ = I$. The problem therefore reduces to a combinatorial counting problem: for each possible dummy–trap coloring of the 5-qubit chain, we count the flip patterns that still lead to a successful test round.

Successful flip patterns

Let the 5-qubit linear graph be represented schematically as

— — — — —

and denote a flipped qubit by x . For example, the pattern $xx---$ means that the first two qubits are flipped.

There are two different coloring assignments of the 5-qubit chain:

1. DTDTD, containing three dummy qubits and two trap qubits,
2. TDTDT, containing three trap qubits and two dummy qubits.

We now count the flip patterns that still produce a successful test round. Those patterns can be found in Table A.1

Table A.1: Flip patterns that still yield a successful test round for the two possible dummy-trap colorings of the 5-qubit linear graph. A flipped qubit is denoted by x .

| Number of flips | DTDTD | TDTDT |
|-----------------|---------------------------------------|--------------------|
| 0 | ----- | ----- |
| 1 | None | None |
| 2 | $xx---$, $---xx$ | None |
| 3 | $x-x-x$, $-xx-x$, $x-xx-$, $-xxx-$ | $xxx--$, $---xxx$ |
| 4 | $xx-xx$ | $xx-xx$ |
| 5 | None | None |

Coloring DTDTD

Looking at the central column of Table A.1 we can combine the outcomes with their respective probabilities and obtain the success probability:

$$\begin{aligned}
p_{\text{succ},1} &= \left(1 - \frac{p}{2}\right)^5 + 2 \left(1 - \frac{p}{2}\right)^3 \left(\frac{p}{2}\right)^2 \\
&\quad + 4 \left(1 - \frac{p}{2}\right)^2 \left(\frac{p}{2}\right)^3 + \left(1 - \frac{p}{2}\right) \left(\frac{p}{2}\right)^4.
\end{aligned} \tag{A.4}$$

Coloring TDTDT

Looking at the last column of Table A.1 we can combine the outcomes with their respective probabilities and obtain the success probability for this pattern

$$\begin{aligned}
p_{\text{succ},2} &= \left(1 - \frac{p}{2}\right)^5 + 2 \left(1 - \frac{p}{2}\right)^2 \left(\frac{p}{2}\right)^3 \\
&\quad + \left(1 - \frac{p}{2}\right) \left(\frac{p}{2}\right)^4.
\end{aligned} \tag{A.5}$$

Average success probability

Assuming that the two colorings are chosen with equal probability, the overall test-round success probability is

$$p_{\text{succ}} = \frac{1}{2} (p_{\text{succ},1} + p_{\text{succ},2}). \tag{A.6}$$

Substituting the two expressions above yields

$$\begin{aligned}
p_{\text{succ}} &= \left(1 - \frac{p}{2}\right)^5 + \left(1 - \frac{p}{2}\right)^3 \left(\frac{p}{2}\right)^2 \\
&\quad + 3 \left(1 - \frac{p}{2}\right)^2 \left(\frac{p}{2}\right)^3 + \left(1 - \frac{p}{2}\right) \left(\frac{p}{2}\right)^4.
\end{aligned} \tag{A.7}$$

This expression gives the analytical success probability of a test round for the 5-qubit linear graph under the simplified depolarizing-noise model considered here. The probability of failure can easily be obtained by taking the conjugate probability

A.3 Extra Figures

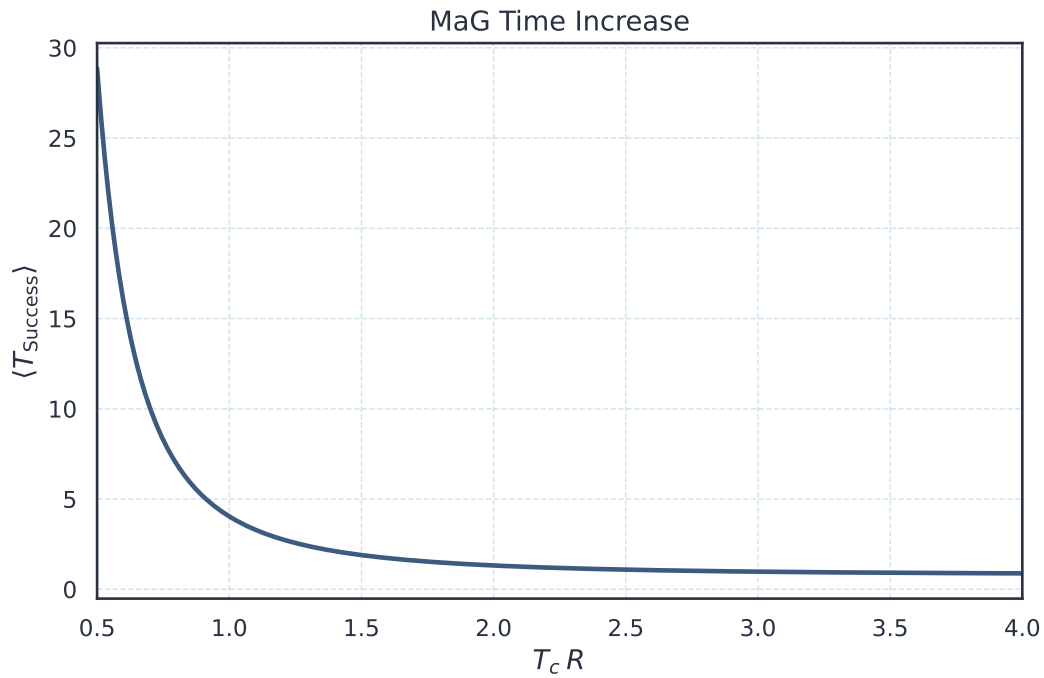


Figure A.1: Normalized time required to complete a successful BQC test round on a five-qubit linear graph as a function of the backbone entanglement rate, for several values of the cutoff time. As the expected backbone waiting time $1/R$ exceeds the cutoff T_c , the required time per successful round grows rapidly, reflecting the increasing fraction of rounds that are discarded and restarted.

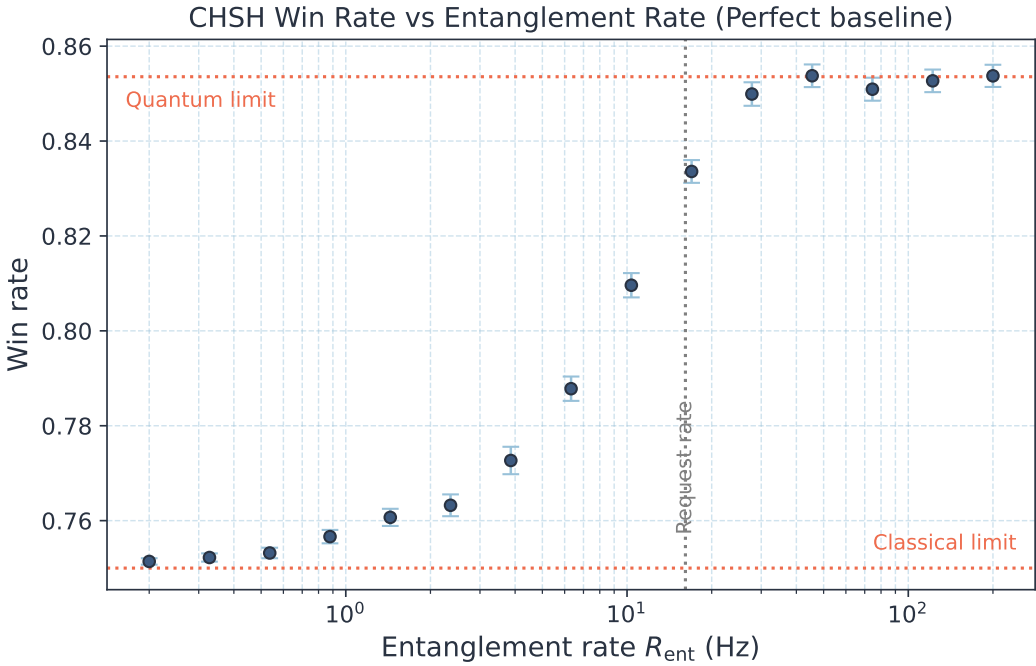


Figure A.2: End-to-end entanglement availability as a function of the entanglement generation rate under perfect metropolitan hardware and unit backbone fidelity. In this regime, the win rate is determined only by whether entanglement is available when a request arrives, and therefore directly reflects the availability of the quantum strategy. Once the entanglement rate exceeds the request rate by a sufficient margin, entanglement is effectively always available. Error bars indicate ± 1 standard error.

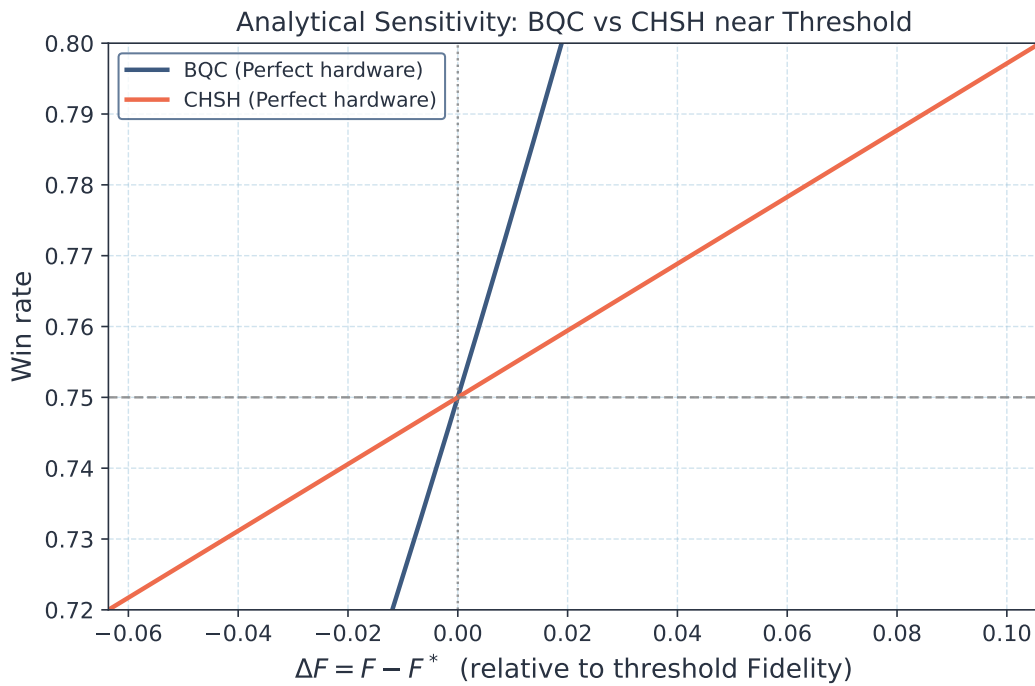


Figure A.3: Comparison of the sensitivity of the perfect metropolitan hardware CHSH-game winning probability and the VBQC test-round success probability to changes in backbone fidelity near the threshold value of 0.75. The VBQC curve exhibits a substantially steeper slope, indicating a stronger dependence on backbone fidelity, whereas the CHSH-game win rate changes more gradually in this threshold region.