HAW HAMBURG

Bachelorarbeit

Justin Hoffmann

Increasing Throughput of Remote Entanglement for Quantum Internet Technologies

Faculty of Engineering and Computer Science
Department Computer Science

Justin Hoffmann

Increasing Throughput of Remote Entanglement for Quantum Internet Technologies

Bachelor thesis submitted for examination in Bachelor´s degree
in the study course *Bachelor of Science Informatik Technischer Systeme*
at the Department Computer Science
at the Faculty of Engineering and Computer Science
at University of Applied Science Hamburg

Supervisor: Prof. Dr. Martin Becke
Supervisor: Prof. Dr. Franz Korf

Submitted on: 10.10.2024

Justin Hoffmann

Thema der Arbeit

Increasing Throughput of Remote Entanglement for Quantum Internet Technologies

Stichworte

Quanteninternet, Quantenkommunikation, Quantennetzwerksimulation, Quantenteleportation, Fernverschränkung, Entanglement Routing, Entanglement Swapping, Fidelity

Kurzzusammenfassung

Diese Bachelorarbeit umfasst eine Netzwerkevaluation unter Verwendung des Quantennetzwerk-Tools *SimQN* und fokussiert sich auf Entanglement-Routing-Algorithmen. Ein proaktiver Referenzalgorithmus dient als Grundlage für die Vergleichsanalyse, während der vorgestellte Ansatz die Netzwerktopologie virtuell erweitert, um Leistungssteigerung zu erzeugen. Ziel war es, den Durchsatz für die Verteilung von Entanglement-Ressourcen zwischen End-Nodes zu erhöhen. Es wurden zwei Experimente durchgeführt — eines opportunistisch und eines in wachsenden, zufällig-generierten Topologien — wobei Daten zu Durchsatz, Fidelity und Netzwerkverkehr gesammelt wurden. Bedeutende Leistungssteigerung ist im opportunistischen Setup beobachtet worden und manifestierte sich ebenfalls im skalierten Versuch. Die Ergebnisse zeigen kompetitive Fidelity-Werte, reduzierten Netzwerkverkehr und erhöhten Durchsatz, was die Vorteile des virtuellen Paradigmas hervorhebt. Abschließend empfiehlt die Arbeit mehr Standardisierung aufgrund der hohen Dynamik des Forschungsfeldes und regt an, virtuelle Konnektivität als Designprinzip für zukünftige Entwicklungen des Quanteninternets zu etablieren.

Justin Hoffmann

Title of Thesis

Increasing Throughput of Remote Entanglement for Quantum Internet Technologies

Keywords

Quantum Internet, Quantum Communication, Quantum Network Simulation, Quantum Teleportation, Remote Entanglement, Entanglement Routing, Entanglement Swapping, Fidelity

Abstract

This bachelor's thesis evaluates a quantum network using the simulation tool *SimQN* centered around entanglement routing algorithms. A proactive reference algorithm serves as a basis for comparative analysis, while a virtual approach integrates topological augmentation to enhance performance. The primary goal was to increase throughput for remote entanglement distribution among end nodes. Two experiments were conducted — one opportunistic and one involving large-scale random topologies — with data collected on throughput, fidelity, and network traffic. Significant performance improvements were observed in the opportunistic experiment and were evident in the large-scale setting. The findings demonstrate competitive fidelity values, reduced network traffic, and increased throughput, reinforcing the advantages of the virtual routing paradigm. The thesis concludes with a call for standardization in this dynamic field, advocating for virtual entanglement-enabled connectivity as a standard design principle for future quantum internet development.

Contents

List of Figures

1 Introduction

The *Quantum Internet* (QI) envisions a future global network of quantum computers [15]. Quantum computers supplement the most primitive and elementary unit of information, the *bit*, with a new type of data carrier, the *qubit*. In *quantum information* processing, quantum mechanical systems encoded in qubits are manipulated [44]. The quantum data plane — as opposed to the classical data plane — is best characterized by those properties lacking counterparts in classical information science. Inherent challenges, such as high fragility to the environment, making long transmission distances and storage times infeasible, as well as the intrinsic inability to copy, amplify, or measure quantum data without altering the qubit's state, fundamentally challenge classical networking paradigms [38]. However, a defining phenomenon of quantum information processing is distributed, non-local correlation, i.e., *quantum entanglement* — prominently deemed *"spooky action at a distance"* [9].

QI technologies leverage these innate features in synergy with a classical communication backbone to enhance conventional capabilities. For example, the impossibility of measuring a qubit without altering its state has massive implications for cryptography. *Quantum Key Distribution (QKD)* protocols exploit the physical inability of third parties to listen in on a conversation without leaving a trace in the form of an altered quantum state. Similarly, in *Blind Quantum Computing (BQC)*, those attributes are utilized in a scenario where clients delegate computational tasks to untrusted and more powerful remote servers. Those protocols ensure the clients' computation results to stay hidden from the untrusted party [60].

To streamline QI development and take advantage of hindsight gained from years of classical internet evolution, several efforts for early standardization are made. Whether it is defining architectural principles for quantum networking [38, 12], suggesting internet architectures and taking inspiration from the traditional layered software stacks [34, 23], laying out multi-stage roadmaps try to predict progress of QI technologies [61], or unifying QI protocol engineering into a modular framework based on a centralized knowledge

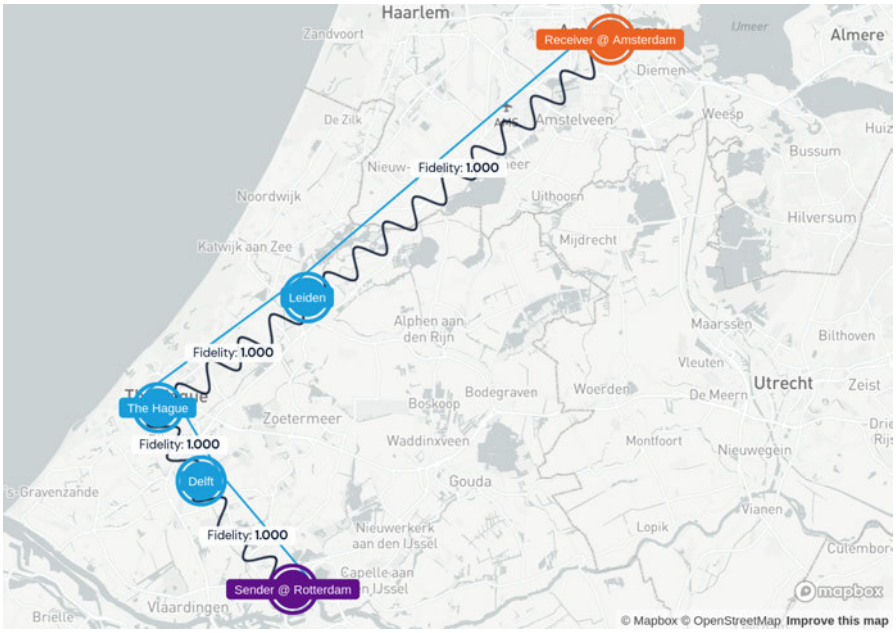

Figure 1.1: Screenshot of the *Quantum Network Explorer (QNE)* interface depicting a real quantum network that connects five dutch cities and offers remote development and deployment of QI services on real hardware [49].

repository [58], the scientific community is determined to not repeat past mistakes. By working with quantum network simulation tools alongside hardware development [21, 22], the creation of QI protocols can progress in parallel, allowing researchers to design and test protocols even before fully functional quantum hardware becomes available [58]. On the other hand, national and international research initiatives work to deploy functional, intermediate-scale quantum networks all around the world [46]. General QKD networks are already functional at short distances on the order of a few hundreds of kilometres [38]. A different example of this is the successful transmission of arbitrary quantum states through entanglement — known as *teleportation* — in a ground-to-satellite quantum network over a distance of 1400 kilometers funded by the Chinese government [51].

[38] recognizes entanglement shared by remote network nodes as the primary resource of QI technologies. This thesis discusses the process of delivering said entanglement through a quantum network, and its associated challenges and algorithms, which [60] lists as a general requirement for QI applications and protocols.

1.1 Objective

Most QI protocols require end-to-end transmission of quantum data. As described in later chapters, the physical delivery of quantum information presents a range of unique challenges regarding feasibility. An alternative and feasible method of communicating in the quantum data plane involves utilizing quantum entanglement as a transmission medium. Thus, entanglement shared between remote network nodes serves as a primary resource in many QI protocols. The process of distributing such end-to-end entanglement represents a sub-discipline within quantum information science known as entanglement routing. This discipline focuses on optimizing the performance and quality of remote entanglement distribution based on various heuristics. Protocols in this area address the unique properties of quantum information while also considering classical solutions to the routing problem. The research conducted in this thesis is situated within this field of study, which is fundamental to advancing QI technologies.

The thesis presents two distinct entanglement routing protocols, each grounded in a different paradigm. The first protocol is a proactive reference algorithm sourced from the existing literature, while the second represents an original virtual solution developed based on design principles in the field of quantum networking. Here, both protocols undergo quantitative analysis in a network evaluation that primarily focuses on assessing network performance through the metric of remote entanglement throughput. Additionally, the evaluation considers secondary factors, including network traffic and entanglement quality, to provide further analysis of viability. The main objective of this research, however, is to achieve increased throughput by mechanism of the design principles incorporated in the proposed solution. This approach aims to make a case for the QI paradigm of *virtuality*, utilized in the entanglement routing process, and highlights its potential to enhance overall network performance.

1.2 Structure

To assist the reader in navigating the thesis, this section provides an overview of its structure, highlighting the content and purpose of each chapter in relation to the overall research:

- **Chapter 2 — Preliminaries:** This chapter outlines the fundamental principles of quantum information, low-level quantum communication protocols, and the design principles of quantum internet architectures.
- **Chapter 3 — Related Work:** This chapter situates the thesis within the existing literature on remote entanglement routing and distribution, providing context through a general overview of the field.
- **Chapter 4 — Methodology:** This chapter details the proposed implementation, the comparative network evaluation framework, and the design of the two experiments conducted.
- **Chapter 5 — Evaluation:** In this chapter, the results of the experiments are presented and analyzed.
- **Chapter 6 — Discussion:** This chapter discusses the key findings of the evaluation, relating them back to the broader body of literature and examining their significance.
- **Chapter 7 — Conclusion:** The final chapter summarizes the contributions and insights of the thesis.

2 Preliminaries

This chapter provides insight to the key concepts of quantum information, namely *qubits* and *superposition*, *quantum gates* and *circuits* as the fundamentals of quantum computing, phenomena such as *entanglement* and *decoherence*, as well as foundational quantum internet protocols such as *teleportation* and *entanglement swapping*. Note that this section merely covers the mathematics underlying quantum computing and communication, not the physical implementation on how to build quantum computers. Also be aware that in this chapter, sections 2.1 to 2.6 serve as a compendium for the mechanisms and mathematical abstractions of quantum information, while 2.7 addresses the architectural principles and design requirements of QI technologies.

2.1 Qubits

In essence, *quantum information* explores how information is processed within quantum mechanical systems. These are physical structures operating at atomic or sub-atomic scales. The *qubit (**qu**antum **bit**)* serves as an abstraction of these systems and forms the cornerstone of quantum information science. Qubits are elementary units of quantum information, analogous to *bits* in classical information. Bits can encode any physical system with a binary state space, mostly by harnessing electrical principles. Similarly, qubits are represented by the properties of quantum mechanical systems and their quantum state. Examples include the spin of an electron or the polarization of a photon. However, understanding qubits as a purely mathematical concept first aids in grasping their foundational principles independent of their physical realizations [44, p. 13].

Dirac notation, sometimes *bra-ket* notation, is specifically intended for syntactically unifying the linear algebra behind quantum mechanics [26]. It is a mathematical tool to denote quantum states without some ambiguities that hitherto existed in the field. The *ket* is a column vector, denoted $|x\rangle$ for a state with label x, where a value at the n-th

index of the vector represents a relation to the probability of the system being in the n-th state of the state space when measured. The way this applies to a single qubit, i.e., a binary state space, is

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tag{2.1}$$

for the qubit being in the concrete state *0* and

$$|1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tag{2.2}$$

for the qubit being in the concrete state *1*. Generally, these are the *orthonormal* — meaning orthogonal and normalized — basis states of a single qubit $|0\rangle$ and $|1\rangle$ and are the forms assumed by the system upon measurement [44, p. 13]. This notation is standard for quantum information and will be used throughout the thesis.

2.1.1 Single Qubits

As a consequence of 2.1 and 2.2 showing what are essentially two concrete bit states, the relationship of quantum and classical information, the former being a superset of the latter, becomes clear. However, the fundamental difference between a qubit and its classical counterpart is how quantum mechanical principles affect its definition. Qubit states exist as a continuum of the physical system's basis states. This is called *superposition*. An n-bit system is in one of 2^n states, while an n-qubit system can be thought of being in all 2^n states at the same time with certain probability [52, p. 60]. This is expressed as a linear combination of the basis states, labeled ψ by convention, e.g.,

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \tag{2.3}$$

for a single qubit system. In other words, a superposed state is the sum of all basis vectors multiplied by some coefficients [44, p. 13].

The coefficients α and β are complex numbers referred to as *amplitudes*. The square of an amplitude's absolute value is equivalent to the probability of the qubit being in that coefficient's basis state upon measurement. This is how α and β relate to the probability of the state being $|0\rangle$ or $|1\rangle$ respectively. Naturally, these probabilities have to sum up

to 1 [44, p. 16]. The normalization condition asserts

$$|\alpha|^2 + |\beta|^2 = 1. \tag{2.4}$$

Different amplitudes can correspond to the same measurement outcome probabilities. Take the *Hadamard basis* as demonstration [52, p. 19]. These vectors are denoted

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \tag{2.5}$$

and

$$|-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}. \tag{2.6}$$

In 2.5, the amplitude of $|1\rangle$ is $\frac{1}{\sqrt{2}}$, while it is $-\frac{1}{\sqrt{2}}$ in 2.6. Although these amplitudes have opposite signs, called *phase*, they correspond to the same probability of $\frac{1}{2}$. Hence, the system has equal probability for both measurement outcomes [44, p. 14].

2.1.2 Bloch Sphere

To obtain additional intuition for superposition, studying the *Bloch sphere* provides a geometric representation of qubits [44, p. 15]. The Bloch sphere parameterizes the qubit state into θ and ϕ. These angles visually map the state to the diagram. As depicted in figure 2.1, θ is the polar angle, with $|0\rangle$ and $|1\rangle$ as north and south pole respectively. This corresponds to the inclination towards one of the two basis states, therefore relating to measurement outcome probability. At the same time, ϕ is the angle changing the position around the equator, controlling the amplitude phases of the state [52, p. 85]. Taking into account 2.4, this results in a unit sphere, where an arbitrary point on the sphere is a valid qubit state defined

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle, \quad 0 \leq \theta \leq \pi, \quad 0 \leq \phi \leq 2\pi. \tag{2.7}$$

Note that ϕ can perform a full rotation around the equator, while θ is limited to half a rotation between the poles. Consider the state in 2.5 for further illustration. In figure 2.1, $|+\rangle$ aligns with the x-axis. From 2.7 it follows that θ is $\frac{\pi}{2}$ by solving

$$\cos\frac{\theta}{2} = \frac{1}{\sqrt{2}} \tag{2.8}$$

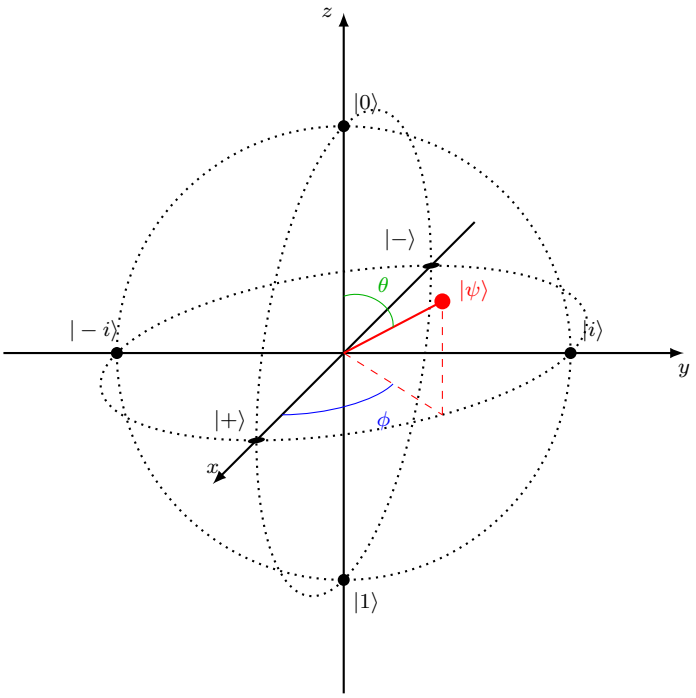

Figure 2.1: Graphical representation (Bloch sphere) of an arbitrary qubit state $|\psi\rangle$, defined by angles ϕ and θ [custom graphic created with TikZ based on [44, p. 15]]

for θ, which is a quarter rotation around the y-axis. The vector sits halfway between $|0\rangle$ and $|1\rangle$ which is interpreted as a probability of $\frac{1}{2}$ for each basis state. In the same way

$$e^{i\phi}\cos\frac{\theta}{2} = \frac{1}{\sqrt{2}} \tag{2.9}$$

can be solved for ϕ equaling 0, meaning no rotation around the equator. The positive amplitudes of 2.5 reflect this. Now consider the states $|i\rangle$ and $|-i\rangle$, which are $\frac{|0\rangle \pm i|1\rangle}{\sqrt{2}}$ and align with the y-axis, showing how deviating from the xz-plane changes the complex component and relative phase of the state [52, p. 23]. The Bloch sphere is a useful diagram, that represents all possible states of a qubit as a sphere. It will reappear in section 2.2 to effectively visualize quantum gate operations.

2.1.3 Multiple Qubits

In most cases, quantum states derive their benefit from being composed of multiple qubits. Adding a second qubit to the general qubit state in 2.3, which can be done via tensor product [44, p. 71], leads to

$$\begin{aligned} |\psi\rangle &= (\alpha_0|0\rangle + \alpha_1|1\rangle) \otimes (\beta_0|0\rangle + \beta_1|1\rangle) \\ &= \alpha_0\beta_0|00\rangle + \alpha_0\beta_1|01\rangle + \alpha_1\beta_0|10\rangle + \alpha_1\beta_1|11\rangle \\ &= \gamma_{00}|00\rangle + \gamma_{01}|01\rangle + \gamma_{10}|10\rangle + \gamma_{11}|11\rangle. \end{aligned} \tag{2.10}$$

It is apparent that the set of basis states doubles in size due to the state encompassing all possible bit configurations. For a pair of qubits, there are four measurement outcomes reflected in the ordered orthonormal basis $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. This aligns with the previous assertion of an n-qubit system having 2^n basis states. Each of these components is associated with an amplitude γ_x, where the label $x \in \{0,1\}^2$ stems from the binary string specifying the bit configuration of the basis state [44, p. 16]. It is evident, that something akin to the Bloch sphere model in figure 2.1 fails to visualize these dimensions and is no longer applicable to an n-qubit system with $n > 1$.

Logically, this scales to

$$|\psi\rangle = \sum_{x\in\{0,1\}^n} \gamma_x|x\rangle = \begin{bmatrix} \gamma_{00\ldots0} \\ \gamma_{00\ldots1} \\ \gamma_{00\ldots10} \\ \vdots \\ \gamma_{11\ldots1} \end{bmatrix}, \quad 0 \leq \gamma_x \leq 1 \tag{2.11}$$

for a general n-qubit system. Intuitively, for a specific state $|x\rangle$, the measurement probability of that outcome is $|\gamma_x|^2$. Extending 2.4 to 2.11 results in

$$\sum_{x\in\{0,1\}^n} |\gamma_x|^2 = 1. \tag{2.12}$$

Consider an equally weighted superposition of a three-qubit system as a tangible example of how measurement probability can be derived from a three-qubit superposition:

$$|\psi\rangle = \frac{1}{\sqrt{8}}(|000\rangle + |001\rangle + |010\rangle + |011\rangle + |100\rangle + |101\rangle + |110\rangle + |111\rangle). \tag{2.13}$$

The probability for each configuration, e.g., measuring all 0's, is $|\gamma_{000}|^2$, i.e., $\frac{1}{8}$. Alternatively, exactly two zeros can be measured with probablity $\sqrt{|\gamma_{001}|^2 + |\gamma_{010}|^2 + |\gamma_{100}|^2}$, i.e., $\frac{3}{8}$.

Ultimately, when multiple qubits combine to form a single quantum state, complexity increases as, for example, the concept of partial measurement becomes relevant. The effects of measuring a subset of the state's qubits are discussed in sections 2.2 and 2.3.

2.2 Computation

Quantum computation embodies the computational framework grounded in the principles discussed earlier and explores how the practical application of quantum information concepts can potentially yield faster solutions to some problems than what is possible classically. One of the earliest instances of this is [55]. This work shows significant speedup for finding prime factors of an integer using the quantum computational paradigm, resulting in important implications for some encryption schemes based on factoring large numbers, such as *RSA* [53].

2.2.1 Gates

To establish a universal toolbox for quantum computation, the most elementary building block for quantum algorithms has to be defined, which is an operation on qubits known as a *quantum gate (operation)*. Initially, quantum gates are established in an algebraic manner. A demonstrative example for a single qubit operation is the X gate, quantum equivalent to logical negation [44, p. 18]. In a classical sense, for a one bit input, 0 outputs 1 and 1 outputs 0. With Dirac notation, negation of 2.1 yields 2.2, and vice versa. Therefore, the operation is equivalent to exchanging both amplitudes so that 2.3 transforms into

$$|\psi'\rangle = \beta|0\rangle + \alpha|1\rangle. \tag{2.14}$$

This is expressed as matrix multiplication such that

$$|\psi'\rangle = X \cdot |\psi\rangle = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \beta \\ \alpha \end{bmatrix}. \tag{2.15}$$

This can be applied to the Bloch sphere. To get from $|0\rangle$ to $|1\rangle$, as well as from $|i\rangle$ to $|-i\rangle$, it takes a half rotation around the x-axis. It follows that the X gate is equivalent to swapping the two basis state's amplitudes and simultaneously constitutes a rotation of π radians around the x-axis.

This begs the question if quantum gates are just matrices. It is not as straightforward as it may seem. To examine what exactly constitutes a quantum gate, it is important to incorporate some additional ideas. In Dirac notation, the *conjugate transpose* of a ket is called a *bra*, denoted $\langle x|$ for a state with label x [26]. The conjugate transpose, indicated by †, is formed by transposing a matrix and applying the *complex conjugate* to all entries [44, p. 18]. The complex conjugate of a number, denoted *, keeps the real part unchanged and inverts the sign of the imaginary part [44, p. 62, p. 70]. The bra matching 2.3 is

$$\langle\psi| = \begin{bmatrix}\alpha\\ \beta\end{bmatrix}^\dagger = \begin{bmatrix}\alpha^* & \beta^*\end{bmatrix}. \tag{2.16}$$

For instance,

$$\langle 0| = \begin{bmatrix}1 & 0\end{bmatrix} \tag{2.17}$$

and

$$\langle i| = \begin{bmatrix}\frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}}\end{bmatrix} \tag{2.18}$$

are the complementary bras to $|0\rangle$ and $|i\rangle$ respectively. Recall that qubit amplitudes hold the constraint formulated in 2.4. According to [44, p. 81], the same condition is represented by

$$\langle\psi|\psi\rangle = \begin{bmatrix}\alpha^* & \beta^*\end{bmatrix}\begin{bmatrix}\alpha\\ \beta\end{bmatrix} = \alpha^*\alpha + \beta^*\beta = 1. \tag{2.19}$$

The inner product of ket and bra has to be 1. For an initial state $|\psi\rangle$, any state $|\psi'\rangle$ that results from a gate operation on that initial state must also satisfy the condition relating to its amplitudes and probabilities [44, p. 18]. Consequently, after a gate operation, such as in 2.15,

$$\langle\psi'|\psi'\rangle = 1 \tag{2.20}$$

must be satisfied. For a matrix U, the condition

$$U^\dagger U = UU^\dagger = I, \tag{2.21}$$

ensures 2.20, illustrated by

$$\begin{aligned}\langle\psi'|\psi'\rangle &= \langle U\psi|U\psi\rangle \\ &= (U\psi)^\dagger(U\psi) \\ &= \psi^\dagger U^\dagger U\psi \\ &= \psi^\dagger I\psi \\ &= \psi^\dagger\psi \\ &= \langle\psi|\psi\rangle = 1.\end{aligned} \tag{2.22}$$

Any matrix that fulfills 2.21 is called *unitary* [44, p. 18]. This mathematical expectation is essential for any matrix to qualify as a quantum gate. In geometric terms, any operation applied to the qubit that does not satisfy this criterion would position the new state off the Bloch sphere's surface. On top of that, all unitary matrices are their own inverse and twofold application yields the initial state again [44, p. 21], e.g.,

$$X^2|0\rangle = I|0\rangle = |0\rangle. \tag{2.23}$$

A particular single qubit gate worth discussing is the *H (Hadamard)* gate [44, p. 19]. This operation takes a vital role in the context of quantum internet communication, seen in later sections of this chapter. It is defined as

$$H = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{2.24}$$

Applying this gate outputs $|+\rangle$ for $|0\rangle$ and $|-\rangle$ for $|1\rangle$. It puts a non-superposed state into superposition with equal outcome probability. On the Bloch sphere, the H gate functions as half a rotation around the diagonal between the x- and z-axis, as depicted in figure 2.2. Taking into account the rotation on the Bloch sphere, it becomes clear that applying H twice to any qubit puts its into the initial state.

To summarize, any unitary $2x2$-matrix signifies a valid quantum gate operation on a single qubit and can be visualized as a vector rotation on the Bloch sphere.

2.2.2 Circuits

Continuing with the composition of these building blocks, the most fundamental tool for developing quantum computing algorithms and communication protocols is the *quantum*

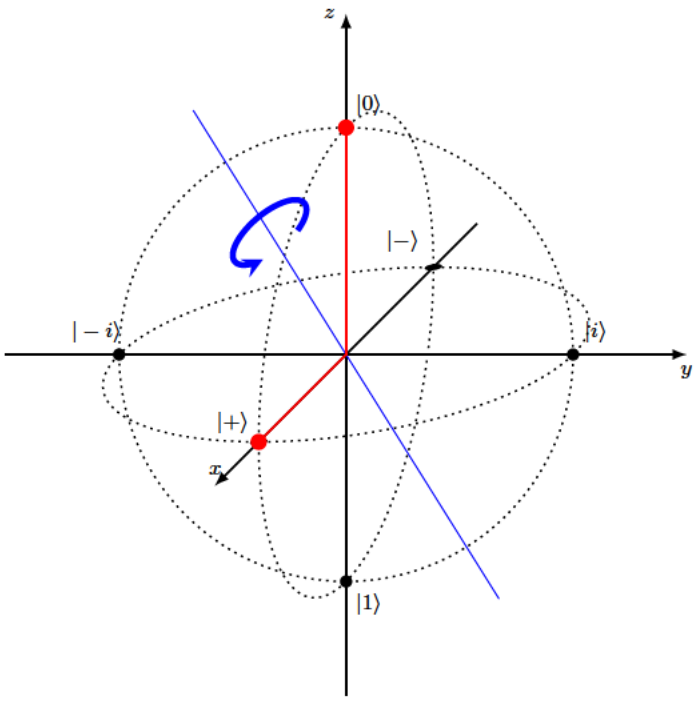

Figure 2.2: Bloch sphere of an H gate operation on the state $|0\rangle$, which represents half a rotation around the diagonal between the x- and z-axis [custom graphic created with TikZ]

circuit. This framework integrates the application of gate operations on individual or multiple qubits, as well as qubit measurement within a single sequential model. In

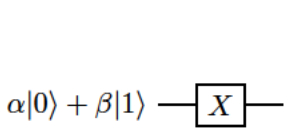

(a) X gate operation

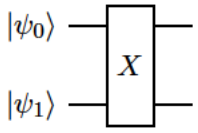

(b) X gate on two qubits

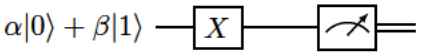

(c) X gate with measurement

Figure 2.3: Collection of basic quantum circuits constructed with the X gate [custom graphic created with TikZ based on [35] and [44, p. 24]]

2.3a you can see the simple circuit of 2.15. For each circuit, one or more qubits are

arranged into parallel wires, on which gate operations are placed as a labeled box on the wire. This is a discrete sequence. Operations that align vertically execute in the same timeslot. Portraying inputs and outputs next to the wires is also an option [44, p. 25]. Gates can span multiple wires, thus multiple qubits, seen in 2.3b. More complex circuits are found in section 2.3 as an analysis of entanglement and 2.5 to define the teleportation protocol.

For a gate operation U, the tensor product $U^{\otimes 2}$ yields a matrix suitable for n-qubit application [44, p. 74]. In the case of 2.3b it is

$$X \otimes X = X^{\otimes 2} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \otimes \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}. \tag{2.25}$$

This inverts the amplitudes of an arbitrary two-qubit state as follows:

$$\begin{aligned} (X \otimes X)|\psi\rangle &= (X \otimes X)(\alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle) \\ &= \delta|00\rangle + \gamma|01\rangle + \beta|10\rangle + \alpha|11\rangle. \end{aligned} \tag{2.26}$$

2.2.3 Measurement

In 2.3c, the operation to the right of the X gate denotes a measurement operation and the wire behind it a classical output. Quantum measurements are characterized by a set of operators $\{M_m\}$, where m denotes the potential measurement outcome [44, p. 84]. In other words, the mathematical application of these operators changes a state into what it would be after measuring a specific m. These operators are defined as the outer product of a basis state, i.e., multiplying ket and bra in that order [44, p. 67]. For example, the measurement operators for basis $\{|0\rangle, |1\rangle\}$ are

$$M_0 = |0\rangle\langle 0| = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \tag{2.27}$$

and

$$M_1 = |1\rangle\langle 1| = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}. \tag{2.28}$$

For a qubit state $|\psi\rangle$,

$$p(m) = \langle\psi|M_m^\dagger M_m|\psi\rangle \tag{2.29}$$

provides the outcome probability of m, of which the square root is used to re-normalize the state after applying measurement gate M_m. Therefore, the measurement of a state constitutes $\frac{M_m|\psi\rangle}{\sqrt{p(m)}}$ [44, p. 16, p. 85]. For further illustration, consider an example where 0 is the measurement outcome of $|+\rangle$:

$$M_0|+\rangle = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} 1 \cdot \frac{1}{\sqrt{2}} + 0 \cdot \frac{1}{\sqrt{2}} \\ 0 \cdot \frac{1}{\sqrt{2}} + 0 \cdot \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix} = \frac{1}{\sqrt{2}}|0\rangle. \tag{2.30}$$

The $|1\rangle$ component is lost. Additionally, normalization transforms the amplitude of $|0\rangle$ into 1:

$$\frac{\frac{1}{\sqrt{2}}|0\rangle}{\frac{1}{\sqrt{2}}} = |0\rangle. \tag{2.31}$$

The state is now $|0\rangle$.

Finally, as mentioned in section 2.1, for multiple qubit joint states, the behavior of partial measurement becomes relevant, i.e., only measuring a subset of the state's qubits. Based on [52, p. 59], to get the measurement operator $M_m^{(j)} = |m\rangle\langle m|$ of an outcome m for qubit j in an n-qubit system, $M_m^{(j)}$ has to be placed at the j-th index in an n-component tensor product, while using the identity operator for indeces $i \neq j$:

$$M_m^{(j)} = I_1 \otimes I_2 \otimes \ldots \otimes I_{j-1} \otimes |m\rangle\langle m| \otimes I_{j+1} \otimes \ldots \otimes I_n. \tag{2.32}$$

Consider the example of measuring 0 for the first qubit within a generic two-qubit system, as depicted in 2.10. According to 2.32, the measurement operator of outcome 0 for the first qubit in a two-qubit state is given by

$$M_0^1 = |0\rangle\langle 0| \otimes I_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \tag{2.33}$$

The entire measurement of 0 in 2.10 comes out to

$$M_0|\psi\rangle = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \gamma_{00} \\ \gamma_{01} \\ \gamma_{10} \\ \gamma_{11} \end{bmatrix} = \begin{bmatrix} \gamma_{00} \\ \gamma_{01} \\ 0 \\ 0 \end{bmatrix} = \gamma_{00}|00\rangle + \gamma_{01}|01\rangle, \tag{2.34}$$

leaving only those configuration outcomes where the first qubit is 0. Again, the amplitudes have to be normalized to culminate the final state after measurement:

$$|\psi'\rangle = \frac{\gamma_{00}|00\rangle + \gamma_{01}|01\rangle}{\sqrt{|\gamma_{00}|^2 + |\gamma_{01}|^2}}. \tag{2.35}$$

Due to the quantum mechanical phenomenon of non-local correlation, partial measurement might influence the state beyond only the measured component, which is further explored in section 2.3.

2.3 Entanglement

This section presents a unique aspect of quantum computing, *quantum entanglement*. It is unique in a way that the probabilistic behavior of entanglement cannot be reproduced by classical mechanisms [62]. This reality reinforces why quantum computers exist and have distinct applications. Instead of delving into the origins and nature of quantum entanglement, the goal is to focus on analyzing entangled states structurally. With this understanding, by accepting the peculiarities of quantum phenomena and treating them as given, computer scientists can start using the effects of entanglement in algorithms. Earlier, manipulating quantum states through measurement was shown, but this only applied to the component being measured. Entanglement, on the other hand, is all about correlation within the state, especially regarding multiple measurement outcomes [44, p. 16]. The correlation is built into the structure of the quantum state and is not affected by physical separation [10]. This section provides a superficial mathematical overview of quantum entanglement. An intuition for entanglement will form by dissecting the concept using the framework laid out so far.

In general and as shown throughout this section, an entangled state can only be described in its entirety with regards to the state's components, rather than its individual components [52, p. 39, p. 218]. To clarify, imagine a quantum state as a rule for probabilistic

behavior. If the objects are somehow manipulated — by measuring them for example — they have to instantly adjust to stay consistent with the quantum mechanical rules imposed by the state. The quantum state describes correlation and will collapse depending on what the correlation dictates. Consider the following entangled state:

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \tag{2.36}$$

This is called a *Bell state*, *Bell pair*, as well as an *EPR pair*, and plays a vital role in quantum communication [6]. The measurement output of $|\Phi^+\rangle$ will either be $|00\rangle$ or $|11\rangle$, i.e., two 0's or two 1's with probability $\frac{1}{2}$. Even with partial measurement, it logically follows from the state structure that upon measurement outcome m, the other qubit can only output m to behave in line with the behavior put forth by the quantum state. This is because there is no basis state $|mn\rangle$ in $|\Phi^+\rangle$ with $m \neq n$. The measurement outcomes of both qubits are correlated.

There are a few ways this can be analyzed. First, take a look at the difference to the arbitrary two-qubit state in 2.10. It is apparent that 2.36 cannot be factored into a product of its individual qubit components like in the generic state of 2.10. By contrast, a composite n-qubit state that can be decomposed into

$$|\psi\rangle = (\alpha_1|0\rangle + \beta_1|1\rangle) \otimes (\alpha_2|0\rangle + \beta_2|1\rangle) \otimes \ldots \otimes (\alpha_n|0\rangle + \beta_n|1\rangle) \tag{2.37}$$

cannot be entangled [34]. Alternatively, applying a single measurement operator to $|\Phi^+\rangle$ should demonstrate similar correlation. Observe the effect of applying M_0^1 in 2.33 to the Bell state:

$$M_0^1|\Phi^+\rangle = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ 0 \\ \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ 0 \\ 0 \end{bmatrix} = \frac{1}{\sqrt{2}}|00\rangle. \tag{2.38}$$

This leads to $|00\rangle$ after normalization. Measuring one qubit collapses the entire superposition of the state. This also shows how the two components are correlated.

To learn about how to construct a Bell state via circuit, a key non-trivial two-qubit gate operation is presented. The *controlled-not (CNOT)* gate takes two qubits as input, one control bit and one target bit. If the control bit is 0, nothing happens to the target bit. If the control bit is 1, the target bit is inverted [44, p. 20].

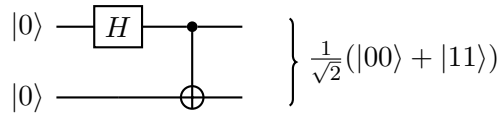

Figure 2.4: Quantum circuit creating Bell states, consisting of an H and $CNOT$ gate on two arbitrary basis states [custom graphic created with TikZ based on [44, p. 26]]

The circuit in figure 2.4 describes the set of operations that yields Bell states [44, p. 26]. This quantum circuit is composed of a $CNOT$ operation on both qubits after an H operation on the control qubit. For any two qubits $|0\rangle$ or $|1\rangle$ as input, it produces one of the four Bell states $|\Phi^+\rangle$, $|\Phi^-\rangle$, $|\Psi^+\rangle$ or $|\Psi^-\rangle$ [52, p. 39]. These differ in amplitude and if the measurement outcomes are equal or opposite. The $CNOT$ operator on the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ is implemented as

$$CNOT = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{2.39}$$

The operator can be used to demonstrate the application of the circuit in figure 2.4. Starting of with $|00\rangle$, the H gate operation, is performed on the first qubit, resulting in

$$(\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle)|0\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|10\rangle. \tag{2.40}$$

The second qubit can now be flipped with $CNOT$ for $|\Phi^+\rangle$:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ 0 \\ \frac{1}{\sqrt{2}} \end{bmatrix} = |\Phi^+\rangle. \tag{2.41}$$

This works for any of the four Bell states. Two-qubit entanglement is called *bipartite* [52, p. 218]. Entangled n-qubit states with $n > 2$ are called *multipartite* [52, p. 225]. However, as seen in later sections, most practical quantum internet technologies require bipartite entanglement.

This section provides a glimpse as to how entanglement might influence an algorithm. It is now more obvious why quantum computer programmers need not concern themselves with physical entanglement realization. The utility of bipartite entanglement and bell states, or EPR pairs, for quantum internet applications is presented in section 2.5.

2.4 Decoherence

So far, the focus has mainly been on closed quantum systems, which are isolated from their surroundings. While these are ideal systems that can offer valuable insight into theoretical information processing, they do not reflect realistic conditions. In practical — or open — quantum systems, particles cannot be perfectly isolated, consequently introducing noise to states through spontaneous and unwanted interactions with the environment [44, p. 353]. This process is called *decoherence*. As demonstrated in this section, decoherence constitutes the loss of quantum coherence and superposition, and the transition from quantum to classical probabilistic behaviour [44, p. 398]. Examples of decoherence mechanisms are variations in the magnetic field [44, p. 345], thermal noise [44, p. 282] and imperfections in the equipment used [44, p. 323]. Furthermore, measurement of qubits itself is viewed as a decoherence process, as it collapses the system into a concrete state, discarding information about the superposition [44, p. 282]. Decoherence

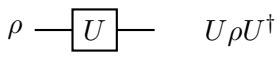

(a) In closed quantum systems, transformations are known unitary matrices.

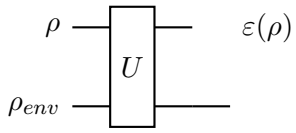

(b) In open quantum systems, environmental influences are out of scope, therefore irreversible.

Figure 2.5: State evolution in open and closed quantum systems [custom graphic created with TikZ based on [44, p. 358]].

triggers an irreversable loss of information [11]. This is demonstrated in figure 2.5. In this context, reversibility is a property coupled to unitary transformations, shown in 2.23. Distinct changes of quantum systems are unitary [44, p. 281]. If the evolution of states is modeled with unitary transformations, as shown in figure 2.5a, some resulting state $\varepsilon(\rho)$ in figure 2.5b would be incomplete with regards to the unknown change induced by its environment. The effect of ρ_{env} on the system ρ captured in the transformation U

is out of scope, making decoherence processes irreversable [44, p. 281]. In other words, there is no way to infer ρ from $\varepsilon(\rho)$ by a unitary transformation [44, p. 358].

The closed systems discussed so far are called *pure states*. These quantum states capture all of their components explicitely, they only deal with known components. All other quantum states are *mixed states* produced by decoherence. They are considered a statistical ensemble of pure states. This introduces uncertainty regarding the behavior of the state, given that it may not fully align with its described state. For pure states, the actual probabilistic behaviour is deterministic [44, p. 100]. Mixed states are modeled with *density operators*, also known as *density matrices*. Density operators are means of describing quantum systems that is alternative yet equivalent to the notational conventions outlined thus far. Consider the density matrix ρ given by

$$\rho = \sum_{i=1}^{m} p_i |\psi_i\rangle\langle\psi_i|, \tag{2.42}$$

which represents a weighted sum of the pure states $|\psi_i\rangle$ that make up the ensemble, with each state occuring with probability p_i [44, p. 99]. Note how ρ is simply $|\psi\rangle\langle\psi|$ for $i = 1$, which are pure states. If the statistical ensemble is made up of two states $|\psi_1\rangle = \alpha_1|0\rangle + \beta_1|1\rangle$ and $|\psi_2\rangle = \alpha_2|0\rangle + \beta_2|1\rangle$ with p_1 and p_2 as their respective probabilities, the density matrix is expressed as

$$\rho = \begin{bmatrix} p_1|\alpha_1|^2 + p_2|\alpha_2|^2 & p_1\alpha_1\beta_1^* + p_2\alpha_2\beta_2^* \\ p_1\alpha_1^*\beta_1 + p_2\alpha_2^*\beta_2 & p_1|\beta_1|^2 + p_2|\beta_2|^2 \end{bmatrix} \tag{2.43}$$

according to 2.42. The example provided in 2.43 illustrates how density operators convey information. The diagonal entries ρ_{00} and ρ_{11} are weighted sums of the probabilities of pure states $|0\rangle$ and $|1\rangle$ and therefore relate to measurement outcome probabilities in the ensemble. The off-diagonal entries are a measure of quantum coherence, in the form of phase relations [44, p. 387]. For further illustration, consider the example of *phase damping*. Phase damping is a type of decoherence and a loss of superposition through, for instance, random photonic scattering in optical fiber [44, p. 290, p. 384]. The change in relative phase constitutes an $R_z(\theta)$ gate operation, which is a rotation around the z-axis on the bloch sphere by a random θ value. Averaging over the possible values of θ results in

$$\rho(t) = \begin{bmatrix} \rho_{00} & \rho_{01}e^{-\gamma t} \\ \rho_{10}e^{-\gamma t} & \rho_{11} \end{bmatrix} \tag{2.44}$$

as the evolution of the density operator [44, p. 384] [57]. Here, quantum coherence is a function of time. Decoherence is conveyed by exponential decline in the off-diagonal matrix entries, as the values approach 0 over time. A diagonal density matrix with no non-zero off-diagonal values, is considered fully decohered and no longer exhibits quantum behaviour [44, p. 387]. It is the process of a pure state becoming more mixed over time, until it is an ensemble of systems displaying only independent classical probabilistic behaviour [44, p. 398]. Geometrically, the bloch sphere would change shape or shrink as a consequence of the state decohering into a less pure state [44, p. 376-378, p. 383].

For practical quantum communications, decoherence remains a central problem. Through spontaneous interaction, particles become entangled with their local quantum neighborhood, information about the superposition is leaked out of scope and therefore lost to unknown quantum territory. Due to the impossibility of perfectly isolating quanta from the environment, decoherence occurs in diverse media, such as different communication channels or quantum memory [44, p. 417]. For this reason, in realistic settings, decoherence is something that has to be accounted for. For computer scientists, it is crucial to acknowledge the existence of decoherence mechanisms and recognize the upstream effects for higher level applications and protocols. These consequences will be discussed in subsequent sections, as they have implications going up the technology stack.

2.4.1 Fidelity

Before moving on to low-level quantum internet protocols, *fidelity* is introduced as a widely used distance measure for quantum state comparison. It employs density operators in its formulation and relates heavily to decoherence, as the fidelity of a state at a point in time to its initial state is often analyzed. The fidelity function F inputs the density operators of two arbitrary states ρ and σ and outputs a value between 0 and 1, with 1 indicating equivalence [44, p. 409]. Results of quantum communication protocols, e.g., a transmitted qubit state, ought to have a fidelity value close to 1 and should only deviate by a small margin to minimize errors. With the trace function definition

$$tr(A) \equiv \sum_i A_{ii} \tag{2.45}$$

[44, p. 75], the formula

$$F(\rho, \sigma) = tr(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}) \tag{2.46}$$

can be simplified into

$$F(\rho, |\psi\rangle\langle\psi|) = \langle\psi|\rho|\psi\rangle \tag{2.47}$$

for most quantum internet applications. This is because a common practice is to compare an arbitrary mixed state ρ to a specified pure state $\rho_\psi = |\psi\rangle\langle\psi|$ [11]. Decreasing fidelity to an initial pure state indicates a loss of information and is therefore used as a measure of decoherence severity — or quantum state quality. In practice, fidelity is calculated from a density matrix reconstruction [56] or algorithmically estimated [29]. Protocols are able to function despite imperfect fidelity and fidelity demands are set by the specific application [38].

2.5 Teleportation

Fundamentally, *quantum teleportation* is an internet protocol for communication based on quantum computing [58]. Conceptually, it leverages entanglement in the form of bipartite *entangled pairs (EPs)* such as EPR pairs introduced in 2.36. Teleportation, as well as the concept of entanglement swapping discussed in section 2.6, are characterized as primitive operations for quantum networks [4]. It is the enabling principle for QI technology, as it poses the most simple entanglement-based communication and therefore one of the few scalable methods of quantum data transmission [11]. As discussed in the previous section, decoherence makes transmission inherently lossy and exponentially less feasible with increasing physical distance. Teleportation as the primary approach for data transfer is a way to overcome decoherence drawbacks. It also constitutes a paradigm shift away from conventional information delivery to embedding transmission into the algorithm itself, without the need for a physical medium and the occupation of network resources [38].

Teleportation is part of a class of transformations called *local operations and classical communication (LOCC)*, a design principle where devices in a quantum network interface with classical information systems. The protocol is split into local measurements of the qubit to manipulate the quantum state and the classical communication of a few bits as control logic to facilitate specific adjustments on the receiver side [44, p. 573]. Furthermore, the sender does not require knowledge of the state or the location of the recipient [6]. Teleportation is a logical reconstruction of a qubit's state. Figure 2.6 is a simplification of this process. Both of the source's qubits decohere into either $|0\rangle$ or $|1\rangle$ during measurement and the entanglement is consumed. The receiver then recovers the

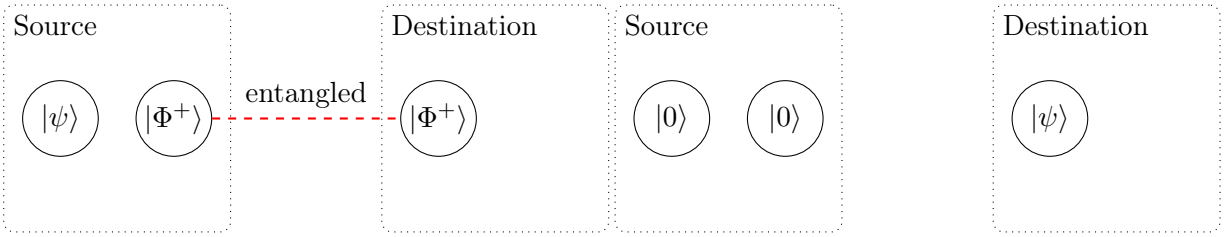

(a) Initially, the source wants to transmit $|\psi\rangle$ to the destination, utilizing a pre-shared entanglement.

(b) After teleportation, the shared entanglement was consumed to reconstruct $|\psi\rangle$ at the destination.

Figure 2.6: Schematic of quantum teleportation [custom graphic created with TikZ].

initial payload state $|\psi\rangle$ with local quantum operations. Despite intuitions one might have about the term 'teleportation', the sender's particle is not physically moved to a different location. Nor is the transmitted quantum state a duplicate of the original. Cloning or copying a quantum state is impossible according to the *no-cloning theorem* [44, p. 25]. Moreover, the correlated action of entangled qubits that are physically seperated might suggest the possibility of communication that is faster than the speed of light. However, there is no way to gain information just from the local measurements alone [44, p. 108].

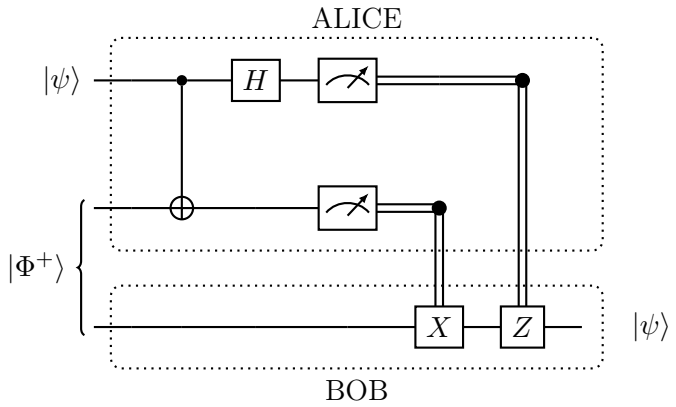

Figure 2.7: Quantum circuit of the teleportation protocol [custom graphic created with TikZ based on [11]].

For this section, two distinct arbitrary communication partners are named *Alice* and *Bob* by convention introduced in [53]. The following calculations can be retraced in figure 2.7 and are based on [11]. Suppose that Alice holds a qubit q_0 with an unknown state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, which she wants to transfer to Bob. As pre-requisite, Alice and Bob

each hold one half of the EP — qubits q_1 and q_2 respectively — with state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \tag{2.48}$$

The initial composite state $|\phi_1\rangle$ of the closed system including all three qubits is

$$\begin{aligned} |\phi_1\rangle &= |\psi\rangle \otimes |\Phi^+\rangle \\ &= \Big(\alpha|0\rangle + \beta|1\rangle\Big) \otimes \frac{1}{\sqrt{2}}\Big(|00\rangle + |11\rangle\Big) \\ &= \frac{1}{\sqrt{2}}\Big[\alpha|000\rangle + \alpha|011\rangle + \beta|100\rangle + \beta|111\rangle\Big]. \end{aligned} \tag{2.49}$$

First, Alice performs a *CNOT* operation on q_1 with q_0 as the control, changing $|\phi_1\rangle$ to

$$|\phi_2\rangle = \frac{1}{\sqrt{2}}\Big[\alpha|000\rangle + \alpha|011\rangle + \beta|110\rangle + \beta|101\rangle\Big]. \tag{2.50}$$

Next, an H gate is performed on q_0. Taking into account the H transformations of the basis states, i.e.,

$$H(|0\rangle) = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad \textit{and} \quad H(|1\rangle) = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle), \tag{2.51}$$

the state $|\phi_2\rangle$ turns into

$$\begin{aligned} |\phi_3\rangle &= \frac{1}{\sqrt{2}}\Big[\alpha H\big(|0\rangle\big)|00\rangle + \alpha H\big(|0\rangle\big)|11\rangle + \beta H\big(|1\rangle\big)|10\rangle \beta H\big(|1\rangle\big)|01\rangle\Big] \\ &= \frac{1}{\sqrt{2}}\Big[\alpha\Big(\frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big)\Big)|00\rangle + \alpha\Big(\frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big)\Big)|11\rangle + \\ &\qquad \beta\Big(\frac{1}{\sqrt{2}}\big(|0\rangle - |1\rangle\big)\Big)|10\rangle + \beta\Big(\frac{1}{\sqrt{2}}\big(|0\rangle - |1\rangle\big)\Big)|01\rangle\Big] \\ &= \frac{1}{2}\Big[\alpha\,|000\rangle + \alpha\,|100\rangle + \alpha\,|011\rangle + \alpha\,|111\rangle + \\ &\qquad \beta|010\rangle - \beta|110\rangle + \beta|001\rangle - \beta|101\rangle\Big] \\ &= \frac{1}{2}\Big[|00\rangle \otimes \Big(\alpha\,|0\rangle + \beta\,|1\rangle\Big) + |01\rangle \otimes \Big(\alpha\,|1\rangle + \beta\,|0\rangle\Big) + \\ &\qquad |10\rangle \otimes \Big(\alpha\,|0\rangle - \beta\,|1\rangle\Big) + |11\rangle \otimes \Big(\alpha\,|1\rangle - \beta\,|0\rangle\Big)\Big] \end{aligned} \tag{2.52}$$

After the H gate application, $|\phi_3\rangle$ is in a position of direct correlation between the four states $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$ and the state of q_2. Finally, by Alice measuring both of her qubits q_0 and q_1, the outcome will be either $|00\rangle$, $|01\rangle$, $|10\rangle$ or $|11\rangle$ with probability $\frac{1}{4}$. These will be control bits for Bob's final operations. By communicating these to Bob over a classical channel, he can perform conditional state rotations, thus completing the reconstruction of $|\psi\rangle$ by leveraging the correlation in $|\phi_3\rangle$. For instance, if Alice measures 1 twice, Bob realizes that the state is $\alpha|1\rangle - \beta|0\rangle$ from which he can recover $\alpha|0\rangle + \beta|1\rangle$ by applying X and Z.

Teleportation has been experimentally demonstrated using atoms [45] and photons [8]. It has worked over 1400 kilometers with good fidelity in a ground-to-satellite quantum network [51]. Teleportation is sometimes referred to as a *low-level* protocol [58], due to its enabling nature of making qubit transmission over large distances feasible. Higher level protocols, such as protocols for key distribution, often utilize teleportation as means of transmission [32].

2.6 Entanglement Swapping

Quantum Teleportation uses a pre-established EP for quantum state transmission. These EPs are generated locally [11], which creates the following dilemma. For feasible transmission in the face of decoherence, an EP is required. However, to distribute the qubits of an EP — each of which is also subject to decoherence — decoherence over large distances must be overcome. The solution for this dilemma is a key functionality of the quantum internet called *entanglement swapping*, introduced in [65], which is another protocol from the $LOCC$ class. Entanglement swapping is the basis for generating end-to-end entanglement, and therefore quantum networking [1].

This protocol leverages the concept of *quantum repeaters*, which are any intermediary quantum nodes between two end points capable of quantum operations and storage [4]. It is a way to produce entanglement between nodes that are non-adjacent. Conceptually, if Alice and Bob both share an EP with the repeater — $|\Phi^+\rangle_{AB}$ for Alice and $|\Phi^+\rangle_{CD}$ for Bob with qubits A, B, C and D— the repeater produces an EP $|\Phi^+\rangle_{AD}$ between Alice and Bob. This is done by teleporting the state of qubit B to Bob using $|\Phi^+\rangle_{CD}$. The entanglement to A is preserved after teleportation. This process is depicted in figure 2.8. Classical communication only takes place between Bob and the repeater to facilitate the same post-processing shown in section 2.5. Two short distance EPs are combined

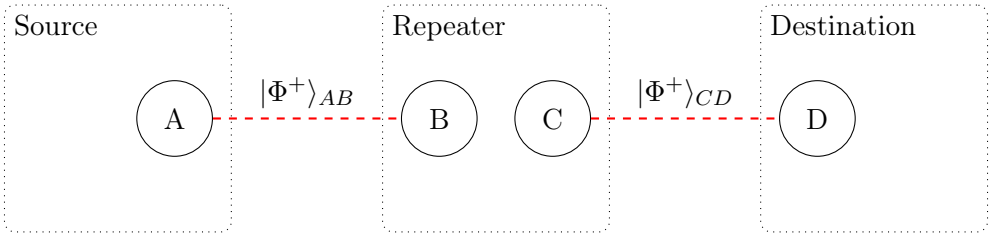

(a) Initially, source and destination each share entanglement with an intermediary node.

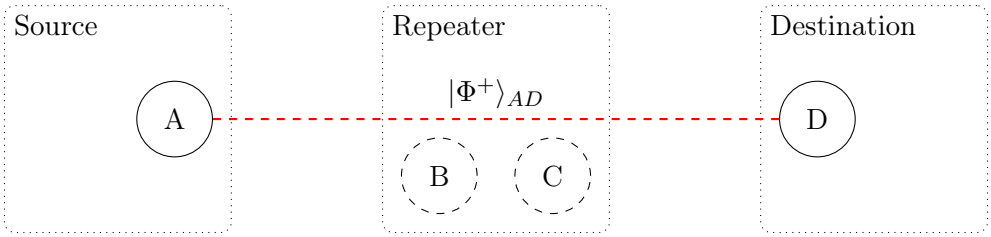

(b) After swapping at the repeater, source and destination share entanglement.

Figure 2.8: Schematic of the entanglement swapping protocol [custom graphic created with TikZ based on [4] and [34]].

into a long distance EP. The process can be carried out iteratively over repeater paths of arbitrary lengths. Evidently, physical qubit transmission between the nodes of an entanglement swapping path has to be feasible to create EPs between adjacent nodes. Entanglement swapping therefore extends the teleportation protocol by exploiting — and dictating [2] — the design of quantum network infrastructure. Entanglement shared by physically adjacent nodes in the network is *elementary* [1]. Entanglement shared by nodes that are not physically adjacent in the network is *remote* [27]. This is the standard vocabulary for the thesis. Entanglement swapping is based on teleportation and represents the mechanism for distributing remote EPs in a quantum network during entanglement routing.

2.7 Quantum Internet

So far, there have been two theoretical frameworks by the *Internet Engineering Task Force* (IETF) discussing QI development — namely RFC 9340 [38] and RFC 9583 [60]. RFC 9430 considers entanglement distribution the enabling service of quantum networks. Similarly, RFC 9583 outlines the requirements for quantum networks, which involve classical

messaging, universal capability for local qubit operations, elementary entanglement generation, remote entanglement distribution, and comprehensive management standards for quantum resources across the network. The goal is to leverage entanglement for new types of distributed protocols that are not possible through classical means. Consequently, new challenges and design principles arise for these networks. Constraints in terms of fidelity and time due to decoherence, as well as paradigm shifts impacting protocol design, networking and connectivity fundamentally shape QI development. In this section, three aspects are examined, each representing a change in paradigm from traditional philosophies.

2.7.1 Networking Approach

There are two general approaches to how data is propagated through a network [59]. *Packet switching* involves breaking up data into packets, which are moved separately in the network. Each packet is stored by a router and forwarded to the next destination based on forwarding table entries. *Circuit switching* sets up a dedicated conversation between two devices in the network, which is based on traditional telephone networks. In other words, the necessary network resources are reserved which allows for constant transmission rates. Qubits are inherently stateless and do not contain additional header information making packet switching approaches difficult [38]. For LOCC protocols, control information is communicated independently with classical channels. Since quantum data cannot simply be read and identified, it introduces the need for intermediate nodes to closely manage and track entanglement resources by memory location, as well as account for the asynchronous protocol instructions. Contrast this to messages in the traditionally packet-switched internet, where state lies in the messages, making forwarding a mostly stateless operation. This differs from the packet-switching approach of the internet. Consequently, [34] argues that quantum networking mirrors connection-oriented circuit switching, and whether entanglement management should be centralized akin to the telephone network. As opposed to classical *store-and-forward* transmission [59], this is referred to as *store-and-swap* transmission, where intermediate nodes turn from forwarding routers into swapping repeaters [38]. Although some attempt to map packet switching onto quantum data transmission [25], most others lean into this paradigm [39, 12, 27].

2.7.2 Protocol Design

In the presented paradigm of 2.7.1, instead of transporting arbitrary quantum states directly, a connection is established by distributing shared EPs between source and destination. The following characteristics of EPs have implications on the design process of QI protocols [34, 38]:

- **Decoherence:** Introduces time constraints of an EP's components in memory — see section 2.4.
- **Bell States:** As a general purpose communication resource with known states, these can be prepared as EPs, distributed in abundance and be accessed and used by any intermediate node along a path — in contrast, arbitrary qubit states cannot be copied due to the *no-cloning theorem.*
- **Metadata:** High demands for closely tracking qubit metadata (storage time, location, or ID) and whereabouts of the second qubit of the EP within the network — without the latter, EPs are useless.
- **Non-Local Correlation:** Manipulation of a single qubit affects the entire quantum state, regardless of physical location of other potential state components — including measurement, which collapses the state, making the entire discipline of classical error correction obsolete.
- **Entanglement Swapping:** Opens up an entire discipline of algorithms that distribute EPs via swapping [1].

As a consequence, the intersection between classical and quantum networking functionalities is narrow. QI protocols are supplementary and assume complete connectivity for classical control information [38]. In fact, the lack of clear overlap between the two makes it unlikely for QI technologies to be functionally autonomous [34]. The relationship can be characterized as bidirectional, resulting in complex interfacing [13]. Some QI protocols provide classical internet functionality, e.g., securing classical communication through key distribution, while others request it, e.g., transmission via teleportation. In general, QI services have the potential to enhance and solve open problems of the classical internet with superposition and entanglement [60]. Leveraging hindsight gained by the incremental evolution of today's internet, standardization can occur early to avoid the same pitfalls [38]. [34] compares three layered models independent of the traditional OSI reference and TCP/IP models [59] and tailored to QI applications. They are based on the principle of

separation of concern and the characteristics unique to quantum networking. Figure 2.9

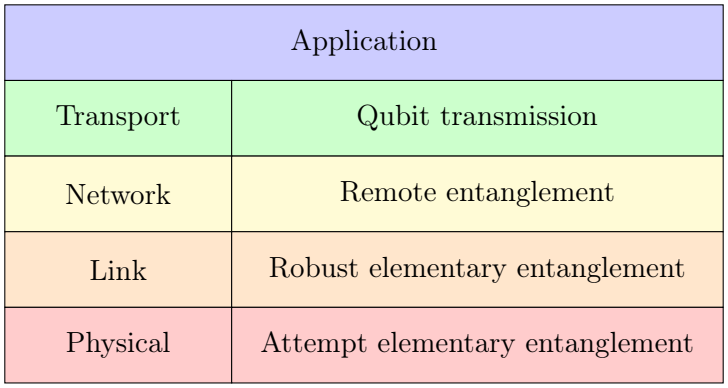

Figure 2.9: Proposal of a potential quantum network stack based [custom graphic created with TikZ based on [23]].

illustrates the hierarchy and each layer's corresponding function of [23]. The proposed architecture puts the focus on abstracting physical details into a hardware abstraction layer, that offers a robust service of perpetually generating ready-to-use elementary EPs [47]. This tackles the necessity for interoperability among different quantum hardware platforms [34], an architectural principle mentioned in [38] and [61]. Accordingly, network layer protocols are simplified into forwarding swap requests to establish a virtual circuit along a path and performing swap operations on said elementary EPs that are readily available [37]. This approach is highly practical for end-to-end quantum communication based on remote entanglement and is in line with the bottom-up service-oriented design philosophy of conventional models.

2.7.3 Connectivity

Due to entanglement assuming the role of a general purpose resource for communication via teleportation, EPs are considered undirected network links that can be used once before having to be replenished [38]. This concept is referred to by various terms, including *virtual quantum links* [54], *virtual circuits* [37], *augmented virtual links* [34], *e-links* [12], and *super links* [27]. The shared aspect here is that the entanglement resource is proactively generated to simulate a communication channel, which impacts the routing process [1]. Informed by [34], four different types of connectivity are illustrated in figure 2.10. Figure 2.10a depicts the physical graph. No nodes in the network share an EP or any other entangled state. In figure 2.10b, entanglement is created over every physical

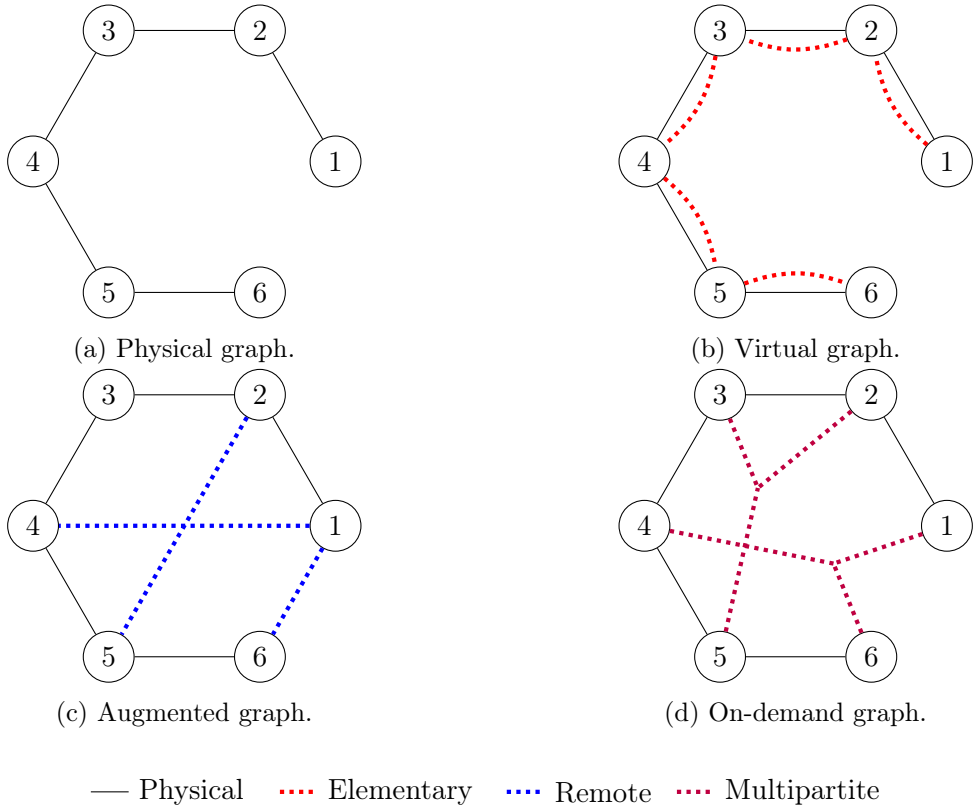

(a) Physical graph.

(b) Virtual graph.

(c) Augmented graph.

(d) On-demand graph.

Figure 2.10: Types of entanglement-enabled connectivity [12] in quantum networks [custom graphic created with TikZ based on [34]].

quantum channel in the network, conforming to the link layer functionality discussed in subsection 2.7.2. Reproducing the physical topology with elementary EPs enables quantum data transmission regardless of the underlying physical conditions. Figure 2.10d portrays an approach that is based on multipartite entanglement. Here, parties extract bipartite EPs from multipartite entangled states shared between multiple nodes, which increases the coordination complexity. Lastly, the augmented graph in 2.10c connects nodes that are disconnected in the physical and virtual graph with bipartite entanglement. It represents the pre-established entanglement-enabled links mentioned above created by swapping elementary EPs belonging to the virtual graph. The separation of concern of the architecture in figure 2.9 becomes evident. [12] argues how scalability in classical internet routing is accomplished by adjusting the frame of reference to a subset of topology information at any given node based on proximity, i.e., *topological depletion*. The shift in paradigm lies in the impact EPs have on the concept of topological neighborhood, by enabling half-duplex unicast between arbitrary nodes in the network for

quantum data. In addition, any one of these links may be used as an intermediary link during the entanglement swapping algorithm to establish remote entanglement. This enables the creation of arbitrary virtual network topologies [1]. This forms the basis for the concept of *topological augmentation* as an opportunistic, virtual strategy for EP distribution in quantum networks.

Without loss of generality and in line with [1], the entanglement-enabled network links introduced in this subsection and depicted in figure 2.10c are called a *virtual links* in the context of this thesis. Furthermore, *topological augmentation* serves as the mechanism intended to achieve throughput improvement in the case study, which is outlined and evaluated in later chapters.

3 Related Work

This chapter contextualizes existing literature relevant to the topic of the thesis. Section 3.2 discusses different options for simulating quantum networks and presents the tool of choice of the thesis. Next, section 3.1 hones in on the protocol implemented in this work's quantum network evaluation. Finally, an identification of the research gap this work intends to fill is presented in section 3.3.

3.1 Protocols

While there are many approaches to the routing phase of *virtual* entanglement routing protocols, [12] conceptualizes the core design philosophy for distributing EPs over a virtual *(logical)* topology. Its suggested multi-level toy model holds close resemblance to the approach theorized in [54], where ring and sphere networks are subdivided into multiple levels of logical topology with decreasing granularity. To reach a destination in the hierarchical structure, the algorithm steps down levels of subdivisions as needed, similar to the level-zero through level-two overlay networks of [12]. In [12], a central problem of the toy model's level-two overlay network remains unaddressed, the selection of its super-nodes. [27] defines this as the *super-link selection (SLS) problem*, where the graph of figure 2.10b is augmented with virtual links *(super-link)* that are pre-calculated over node pairs based on a statistical heuristic to reduce latency of remote EP generation. This is achieved with a greedy selection, as well as a selection based on k-means clustering. Less sophisticated approaches to pre-distribution of EPs throughout the network include random selection and selection based on node degree [17, 48]. Virtual techniques evaluated as part of a quantum network simulator are found in [27], [24], and [36].

Although fidelity is considered a metric of entanglement quality, certain approaches dismiss it as a heuristic tool [28]. *Purification* protocols are algorithms that aim to increase qubit fidelity. Some protocols include purification as a distinct step of entanglement

distribution [64] — considered an architectural principle according to RFC 9340 [38] — while others hope to accomplish higher EP fidelity as an emergent property of the routing algorithm by using a fidelity-based heuristic [37] or avoiding fidelity drivers [16].

Furthermore, many entanglement routing approaches delegate elementary entanglement generation — called the *external forwarding phase* [1] — to the link layer of [23], depicted in 2.9, and build the algorithm on a readily-available virtual entanglement network corresponding to the physical topology [27, 28].

The *internal forwarding phase* contains the entanglement swapping protocol itself [1]. For entanglement routing protocols with an internal phase that is memory-based — i.e., modular EPs of the path are stored in quantum memory for some time before being consumed in a swapping operation — there are several swapping schemes an algorithm can employ to influence distribution performance. The most trivial approach is sequentially swapping entanglement from left to right in a path [39]. Other approaches abstract the swapping path into a binary-tree structure to greatly reduce time complexity, where individual operations may be executed in parallel [28, 14]. In general, a swapping scheme can deploy arbitrary heuristics for the order of entanglement swaps throughout the path [18]. Additionally, there is swapping order calculations based on reinforcement learning [31].

Finally, with the introduction of new network heuristics and metrics for quantum networks, classical algorithms for path computation can be modified. As mentioned above, the requirements of quantum applications differ in terms of performance and fidelity. This can be reflected in the chosen routing metric for fidelity [64], entanglement cost [40], throughput [28], or, as in this study, hop count [43].

This thesis entails a network evaluation based on the throughput of remote entanglement during two entanglement routing protocols. *Proactive* routing defines path computation as a process that occurs before elementary EPs are present, as opposed to *reactive* routing, i.e., elementary EPs already being present at the time of path computation as part of a logical topology (see figure 2.10b) [1]. The first routing protocol that is part of the evaluation conducts a proactive routing phase, as well as a memory-based, sequential forwarding phase, and bases its path computation on a shortest-path metric with uniform edge weights. This algorithm — referred to as *PS (proactive-sequential)* — will serve as reference in the evaluation [39] and has been simulated using the selected tool [21].

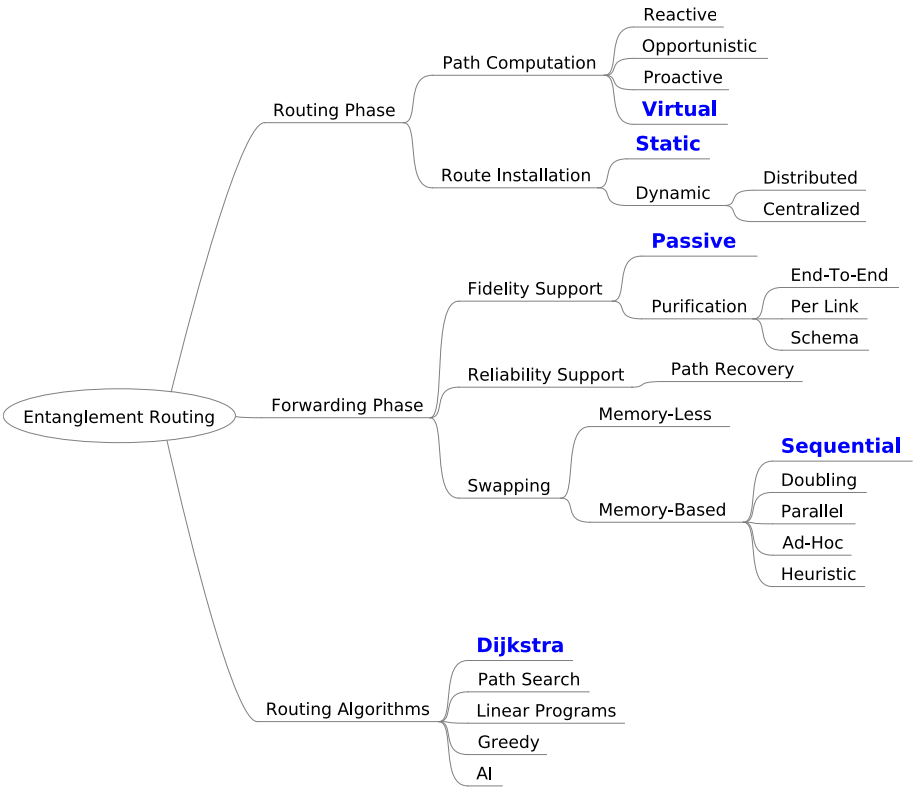

Figure 3.1: Entanglement routing protocol taxonomy [1] *(categorization of VS marked in blue).*

The virtual-link-based solution of this work — referred to as *VS (virtual-sequential)* — differs from the reference algorithm in terms of its routing phase being primarily virtual instead of proactive — see figure 3.1. It utilizes topological augmentation [12] based on the network's community structure [7] to induce better entanglement distribution performance throughout the network. There is often a trade-off between fidelity and performance in entanglement routing, necessitating different approaches for different application requirements [1]. Although there are no purification efforts taken in the protocol, fidelity is passively analyzed. For example, it is expected that decoherence from longer memory wait times of the proposed solution can be offset by an overall reduction of time spent in fibre or swapping operations. The argument is that similar average fidelity of distributed remote entanglement to the reference algorithm suffices to convey the paradigm's benefit in the presence of performance improvement. For a similar reason, a simple shortest-path algorithm based on hop distance and centralized global topology is implemented in the routing phase of the proposed solution, as well as the reference

algorithm. As the emphasis of the proposed solution remains on virtual entanglement routing, the internal forwarding phase implements a sequential swapping scheme, akin to the reference algorithm. The external phase acknowledges the link layer service from [23], specifically adopting the assurance that elementary EP generation is robust and will not fail. However, elementary EP generation remains integrated into the forwarding process, similar to reactive schemes. In other words, the algorithm differentiates between the graph level [12] of the next hop and reacts accordingly.

3.2 Tools

As of September 2024, there are numerous platforms intended for quantum network simulation, with varying use cases and purposes [5]. When simulating network layer protocols, the focus should not be on detailed lower layer hardware characteristics to fit the principle of separation of concern in [23]. NetSquid [22] is a general purpose simulator for quantum network modeling. It is also the most referenced simulator in scientific research of quantum networks and the most versatile when it comes to arbitrarily complex and realistic quantum network simulation. This would deem *NetSquid* as the plausible option, however, it is not tied to any particular layer of a network stack model, making its learning curve rather steep. Although there are community efforts to extend the base package for simpler large-scale bootstrapping [50], there are no built-in features for routing or topology generation. In contrast, *SimQN* is a discrete-event-based quantum network simulation tool with similar intended use cases as *NetSquid* [5], intended to compensate for the drawbacks related difficulty of use. It is also a modular quantum network simulator intended for large-scale network evaluations based on discrete-event simulation and development in Python. It provides users with various options for quantum state simulation, as well as parallelization tools, and has demonstrated a simulation performance 8.02 times better on average than *NetSquid*. The design was partially motivated by mocking a standardized network architecture and layered protocol stack design. This increases user accessibility due to a lack of need to address state of the art bottom-layer implementations. For example, [23] defines a widely used link-layer service for robust EP generation over physical links. Until now, the need for users to either implement or comprehend this service has been a barrier to scaling network analysis. Coupled with network-layer utilities for topology generation and routing, as well as qubit and entanglement backends for varying performance or precision requirements, this significantly reduces bootstrap complexity. Hardware and decoherence details of network entities such as nodes, memories

and channels — classical or quantum — can be specified. Concurrent node applications specify behavior and communication protocols. Several case studies addressing the entanglement routing problem have been conducted with *SimQN* since its release in 2024, however, none of them specifically involve virtual methods [20, 19, 63, 41].

3.3 Research Gap

This thesis proposes a sequential, virtual protocol *VS (virtual-sequential)* for remote bipartite entanglement distribution. VS aims to compete with the sequential, proactive approach *PS (proactive-sequential)*, which is sourced from [39, 21]. By adopting the design philosophies for quantum internet addressing outlined in [12], the goal is to translate these principles into a concrete algorithm. To validate the model, the first network evaluation using *SimQN*, that focuses specifically on virtual entanglement routing through topological augmentation, is conducted. While this evaluation incorporates a secondary review of entanglement quality *(fidelity)*, it primarily emphasizes an analysis of entanglement-related throughput *(EP/s)*.

4 Methodology

This chapter outlines the design and implementation of the proposed solution for entanglement distribution in quantum networks that is part of the network evaluation of this thesis. Section 4.1 details its theoretical foundation and operational design, as well as the structure of the simulated quantum network, while section 4.2 elaborates on the experimental design of the performance evaluation, including metrics, proof of concept and the simulation at scale.

4.1 Implementation

VS is a virtual, sequential entanglement routing protocol based on topological augmentation that establishes remote entanglement as a resource for quantum internet applications and technologies. The quantum network evaluation outlined in this chapter is implemented using *SimQN*. [1] divides entanglement routing protocols into a *routing* and *forwarding phase*. The VS implementation will follow this outline in sections 4.1.1 and 4.1.2 respectively.

4.1.1 Routing Phase

In [12], the practical strategy for figure 2.10 is presented. The network is divided into different sub-graphs, namely a physical graph in figure 4.1a and two separate overlay graphs. For the level-one overlay graph of figure 4.1b, each end-node is clustered and mapped to a super-node that manages an entanglement connection to the super-nodes of the other clusters, similar to figure 2.10b. Note that this structure resembles the proposal put forth in [34], which involves central entanglement management nodes that resemble a circuit-switched network. The level 2 overlay graph of figure 4.1c describes a standalone peer-to-peer topology of all super-nodes, connected through pre-configured entanglement. End-to-end EP distribution between these clusters occurs by entangling

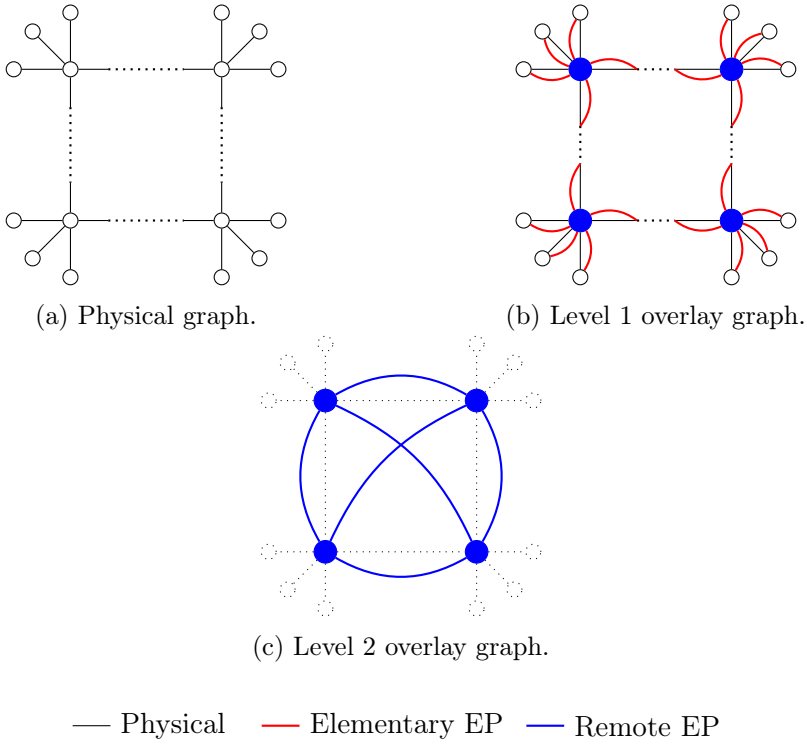

(a) Physical graph.

(b) Level 1 overlay graph.

(c) Level 2 overlay graph.

Figure 4.1: Different graphs used in routing phase VS [custom graphic created with TikZ based on [12]].

the source to a super-node, which can then entangle the source to the super-node serving the destination, and finally, to the destination itself. To complete the virtual topology, the super-nodes connected in figure 4.1c have to be selected through clustering [12], akin to the SLS problem introduced in [27], where virtual links are selected by minimizing an objective function based on a cost budget using k-means and greedy selection. However, this optimization technique necessitates quantitative metrics, particularly for k-means, which relies on a Euclidean metric space [42]. This is not applicable to this situation, in which statistical performance metrics are neglected in the routing process. Instead, there are more fitting ways to learn the structure a network solely based on topology, namely [7]. The *Louvain method* is an algorithm for detecting modular node communities in large networks given a single resolution parameter. The extraction is centered around locally maximizing a modularity function based on edge density, before aggregating the inspected nodes and repeating the process. This groups the network into densely connected areas, with low density between those groupings. The implementation employs the *community_louvain* package [3], which is designed for usage with *NetworkX* [30]. It

offers built in *best_partition* functionality, that returns the set with the highest modularity, trivializing its parameterization. Note, that the rationale for picking *Louvain* is simplicity and ease of bootstrapping, one of the design principles of *SimQN* [21]. All that is left to do is to determine each cluster's centroid node by degree, which will then serve as super-node for graph 4.1c.

To put 4.1 into a practical *routing phase* [1], given three graphs defined as $G_0 = (V, E_0)$, $G_1 = (V, E_1)$, and $G_2 = (V, E_2)$, where V is the set of vertices common to all three graphs and E_0, E_1, and E_2 are the edge sets respective to each graph. The combined graph $G = (V, E)$ with

$$E = E_0 \cup E_1 \cup E_2 \tag{4.1}$$

serves as the global topology for off-line routing table calculation at each node. Here, each edge (u, v) is mapped to its original graph as a tuple (u, v, g), where $g \in \{0, 1, 2\}$. For edges $e_a = (u_a, v_a, g_a)$ and $e_b = (u_b, v_b, g_b)$, where $u_a = u_b$ and $v_a = v_b$ but $g_a \neq g_b$, the edge with the highest value of g is retained in the combined graph. This prevents the protocol from supplying the physical graph with elementary EPs when an elementary EP is already present. This maps onto the design choices for quantum internet addressing presented in [12], specifically the concept of an *e-type* column as part of the addressing scheme. Finally, path computation is facilitated using built in routing functionality by *SimQN*, specifically a shortest-path routing through a network with uniform edge weights.

4.1.2 Forwarding Phase

Continuing with the *forwarding phase* [1]. Here, entanglement swapping (section 2.6) is conducted. This is done in a sequential manner, i.e., from left to right in the path. In *SimQN*, protocols are defined by installing concurrent applications to nodes. Applications are sets of asynchronous message handlers that specify communication behavior. In this context, two distinct applications are implemented to each node:

- **Maintenance application:** Performs the reference protocol PS [39] to supply the pre-selected augmented virtual links to the network as per a designated send rate.
- **Distribution application:** Performs the proposed protocol VS (and PS when no virtual links are present) to establish remote entanglement.

This blend of functionality represents the concept of a network layer entanglement manager suggested in [23].

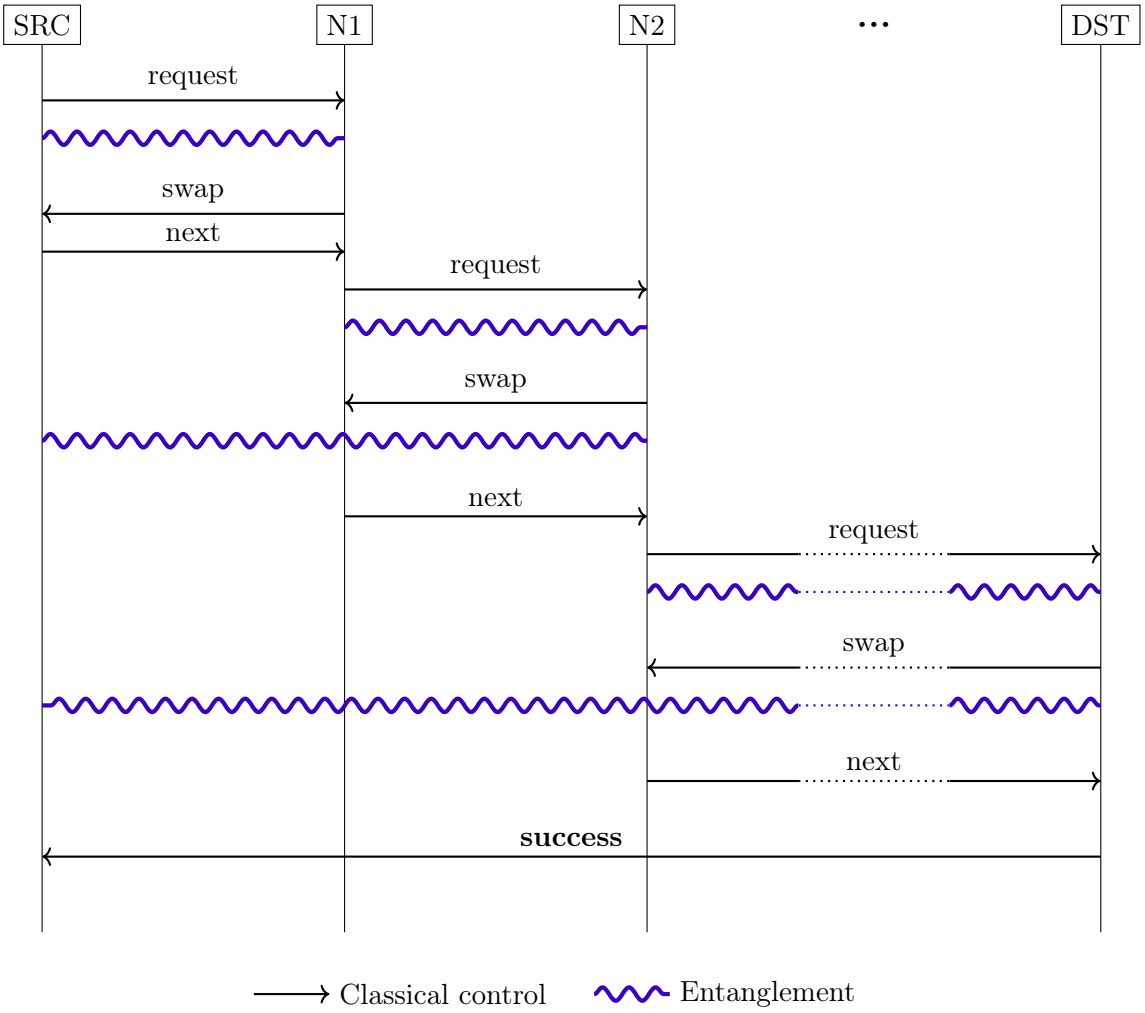

Figure 4.2: Forwarding phase of VS within an arbitrary path [custom graphic created with TikZ based on [39] and [21]].

As mentioned in [38, 34], LOCC protocols require careful coordination through classical control information and bookkeeping of metadata. This is also a consequence drawn from the inherent statelessness of qubits discussed in section 2.7.1. Note, that *SimQN* offers functionality for the synchronization of this process. For example, physical qubit transmission can be attached to classical control information. For every EP distribution, each node receives and records the following information:

- **Transmit ID:** Unique identifier of distribution.

- **Source:** Source of distribution.
- **Destination:** Destination for distribution.
- **Alice Location:** Node that currently holds the first qubit of the EP — this does not change in a sequential scheme.
- **Alice ID:** Unique identifier of the first qubit of the EP — used for storing to and retrieving qubit from quantum storage.
- **Charlie Location:** Node that currently holds the second qubit of the EP.
- **Charlie ID:** Unique identifier of the second qubit of the EP.

The distributed algorithm is characterized by four classical control message types and their response behavior, depicted in figure 4.2:

- **request:** The message itself is based on g of the next edge in the path. If g is 0, the next hop is physical, i.e., no elementary is present yet. It is the task of the sender to now physically transmit a qubit of an EP to the next node and establish a virtual connection. Both parties store their respective qubit in memory. If g is 1 or 2, the next hop is virtual and this step is skipped. Upon obtaining this message, the receiver will allocate the entanglement resource to the specific transmission, update its metadata on the distribution, and send back *swap*. For virtual next hops, the protocol waits for the availability of entanglement resources, that are concurrently established by the maintenance app. Blocked distributions are supplied with a *first-come-first-serve* strategy.
- **swap:** For any node that is not the source node of the distribution, after receiving this message, a swapping operation is executed and the locations in the distribution data updated. The role of the current node is that of an intermediary repeater node that extends the EP range by one hop. The *next* command is sent to the next node in the path.
- **next:** First, the protocol checks, if the destination node in the distribution data and the current node match. In that case, *success* is sent to the source node of the distribution. If not, one iteration of the sequential protocol is complete and the process starts over with *request* to the next node in the path.
- **success:** The distribution was successful, remote entanglement has been established.

Observe, that although the proposed solution is primarily qualified as *virtual*, the implementation works for both presence and absence of elementary EPs, such as in figure 2.10b. In the absence of both augmented virtual and elementary entanglement, it is functionally equivalent to the reference algorithm.

4.2 Experiment Design

SimQN offers a high degree of customizability in terms of network design, such as different concrete and abstract models for quantum entangled states. In this experiment, all EPs take the form of *Bell diagonal states* [33], an API provided by *SimQN*. These are statistical ensembles of *Bell states* (2.36), where decoherence expresses itself in the form of decreasing probability of the state being $|\Phi^+\rangle$. Furthermore, these experiments do not entertain memory-related congestion, thus sufficient quantum memory is assumed. Based on [27], uniform link length is set to 20 km. The quantum memory model of [28] is copied, with coherence times of up to 2 seconds. Two experiments are conducted to test VS against PS. Each test amounts to a fixed simulation time of 50 seconds. Both are conducted with fixed send rates for remote distributions and virtual link distributions of 10 Hz and 50 Hz respectively.

SimQN provides a simulation monitor functionality for data collection in its *simulator* module. The following data is monitored during simulation:

- **Throughput (EP/s):** Rate of successful remote EP distribution [39] — primary objective for improvement in this thesis.
- **Latency (s or ms):** Average, total, minimum and maximum time it takes to distribute a remote EP based on [27] — provides additional information for emergent data trends.
- **Traffic:** Total classical messages sent, qubits physically transmitted and average number of swap operations per EP distribution — provides further justification for or against VS in terms of protocol efficiency.
- **Fidelity:** Average fidelity of distributed remote EPs.

The first experiment is a functional proof of concept in a growing line topology setup with a single fixed virtual link. Here, the first and last nodes will serve as source and destination for continuous endpoint distribution, while the second and second-to-last

nodes will act as source and destination for the virtual link. Beginning with a length of 5 nodes, network size increases incrementally. With each additional node, VS routing will bypass one more node compared to PS routing. This setup is illustrated in figure 4.3. Considering this configuration is opportunistic for virtuality, it is expected that VS will outperform PS in terms of throughput, fidelity and traffic.

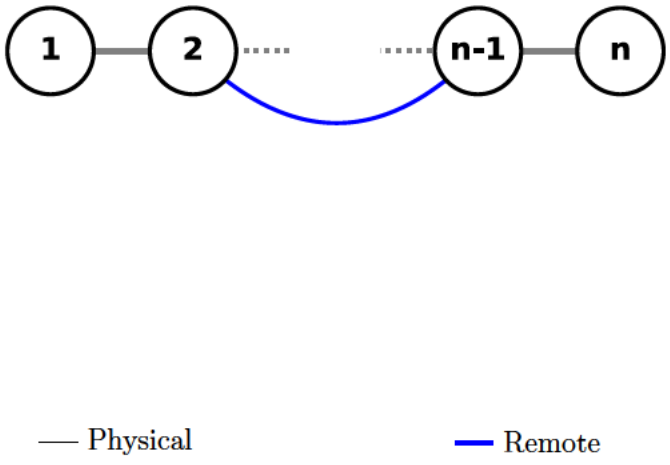

Figure 4.3: Virtual link in an arbitrary line topology positioned at the second and second-to-last node to skip every node inbetween for an opportunistic setup regarding VS.

Secondly, both methods are applied at scale on large random topologies, generated with the network utilities provided by *SimQN*. Here, 20% of possible source-destination pairs for remote EP distribution are selected randomly as static continuous distribution sessions. Starting with 50 nodes, the network grows in steps of 50 up to 500 nodes. For each node count, 3 rounds of data for a varying set of distribution sessions are recorded, of which the median is saved for evaluation. The expected outcome is a manifestation of the results found in the first experiment at scale.

5 Evaluation

This chapter provides an evaluation of the features — throughput, latency and traffic — put forth in 4.2. Section 5.1 details the data from the prototype study, offering validating trends in an opportunistic setup. Section 5.2 demonstrates bottleneck behavior and the role of entanglement as a network resource. In section 5.3, observations of full-scale simulation data display how these trends extrapolate to larger networks and assess whether the results meet the research objectives of the study.

5.1 Proof of Concept

As per section 4.2, the experimental setup introduced in figure 4.3 acts as a functional concept validation. It is an opportunistic network setup for the principle of topological augmentation through virtual entanglement links. The observed trends in the selected features are expected to manifest in large-scale quantum network data.

Begin by reviewing figure 5.1a. With growing network size on the x-axis and distribution throughput on the y-axis, this graph paints the first clear picture of how effective virtual links are in terms of performance. The reference protocol PS (*no vlinks*) is significantly outperformed by the approach VS (*vlinks*) in the routing process. For PS, an exponential decrease of throughput with growing path length can be observed. For VS, notable improvement compared to PS can already be seen for path length 5, which corresponds to skipping a single node with a virtual link. The curve is also a lot more constant with a subtle downward trend. The slight downward slope can be attributed to how the concurrent applications are set up. In the beginning of the simulation, distributions are forced to wait for network resources to be established.

Fidelity is established as a measure of decoherence and therefore the quality of a distributed quantum state. For PS, the bulk of decoherence is caused by the high amount of read, write and swapping operations, as well as physical optical fibre traversal of qubits.

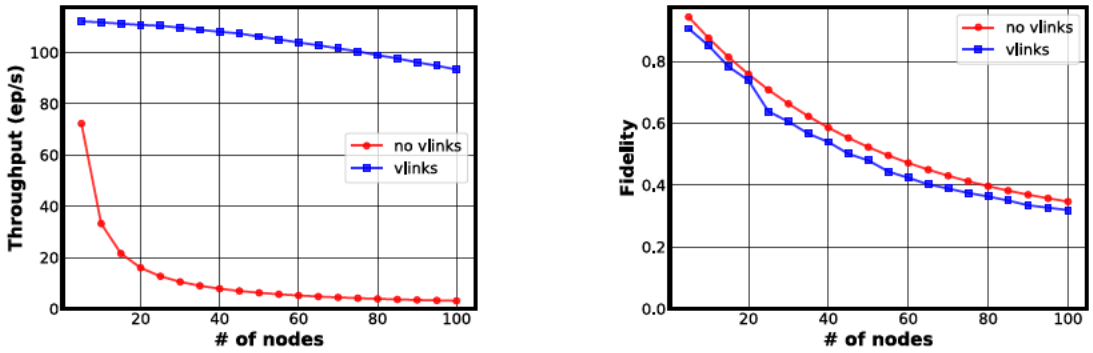

(a) Throughput. (b) Fidelity.

Figure 5.1: Average throughput and fidelity of EP distribution with VS routing *(vlinks)* and PS routing *(no vlinks)* as a proof of concept.

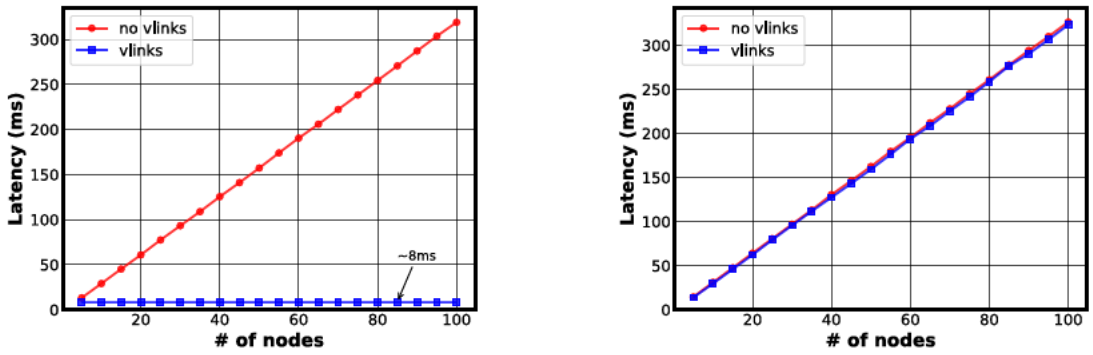

(a) Minimum. (b) Maximum.

Figure 5.2: Proof of concept latency data for VS *(vlinks)* and PS *(no vlinks)*.

On the other hand, VS assumes near perfect fidelity for pre-distributed entanglement, the number of swapping operations is generally lower, and using virtual links results in less optical fibre traversal. However, the effects of qubits waiting in memory for a distribution or, in some cases, for an entanglement resource, seem to outweigh the theoretical advantages virtual links might have in terms of decoherence. As seen in figure 5.1b, average fidelity of both approaches is comparable. PS seems to yield marginally better fidelity for distribution. Note that these differences are subtle and might not express themselves in a large-scale simulation. Considering the performance advantages discovered so far, identical or even lower fidelity outcomes at scale would not skew the net benefit of virtual links.

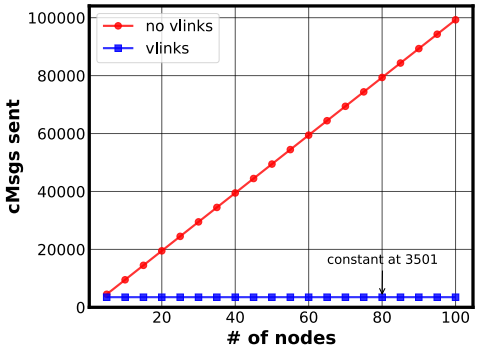

(a) Classical communication aggregate.

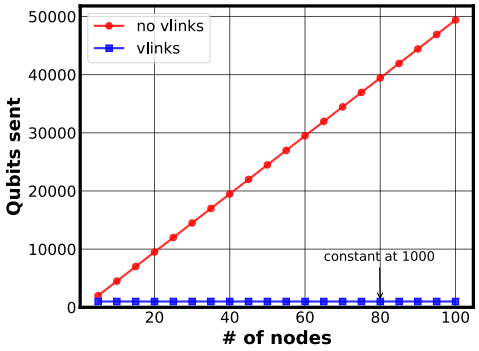

(b) Quantum communication aggregate.

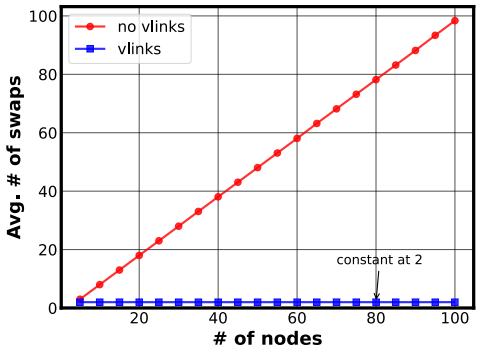

(c) Average number of entanglement swapping operations per EP distribution.

Figure 5.3: Proof of concept traffic data for VS *(vlinks)* and PS *(no vlinks)* .

Figure 5.2 provides data based on the distribution latency of remote entanglement in the prototype setup. These results help substantiate the throughput and fidelity results presented so far. For example, the maximum distribution latency values in figure 5.2b — the worst case scenario — for both PS and VS in the distribution are equivalent. This is reflective of the warm-up period of the simulation where distribution and maintenance effectively run in parallel. Both processes start simultaneously and the distribution has to wait for an available virtual link. As mentioned, this impacts the average throughput of the simulation in the form of a subtle downward slope. Furthermore, minimum latency remains constant using VS while it increases linearly for PS. This is because virtual link utilization essentially limits the swapping path to a length of 4. The idealized and static experiment circumstances lead to minimal difference in latency between the features presented for PS.

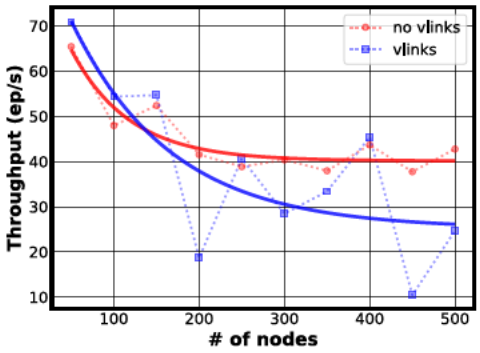

Figure 5.4: Throughput for 20% utilization for VS *(vlinks)* and PS *(no vlinks)*.

Finally, review figure 5.3 for additional traffic data. The effective swapping path length of 4 has already been pointed out. This realization is mirrored in the traffic data. Figure 5.3a shows constant total classical messages sent with increasing path length for VS compared to PS scaling linearly. The same applies to qubits physically sent through fibre and the average amount of swapping operations per distribution. Figure 5.3a shows constant total classical messages sent with scaling path length.

5.2 Entanglement as a Resource

Before moving on to a full-scale network simulation of, a point has to be made about the nature of remote entanglement in the context of topological augmentation. In this section, *network utilization* is defined as the percentage of possible unique source-destination pairs that take part in an active EP distribution session. For instance, for a network with 10 nodes, 5 sessions are considered 100% network utilization. As mentioned, remote entanglement represents a network resource in the form of a single-use unidirectional network link. When using these links to shortcut a given network topology for higher performance, the potential for bottlenecks has to be acknowledged. Due to the randomized setup, the virtual link maintenance might not be able to keep up with the demand within its node cluster. Severe congestion will be the consequence. Unfavourable scenarios like these will impact aggregate feature data. Essentially, virtual links have to be produced faster than they are consumed by concurrent distribution sessions.

The effects of static rates for scaling network size take form in figure 5.4. Initially, the rate of virtual link distribution is five times as high as the distribution rate of the

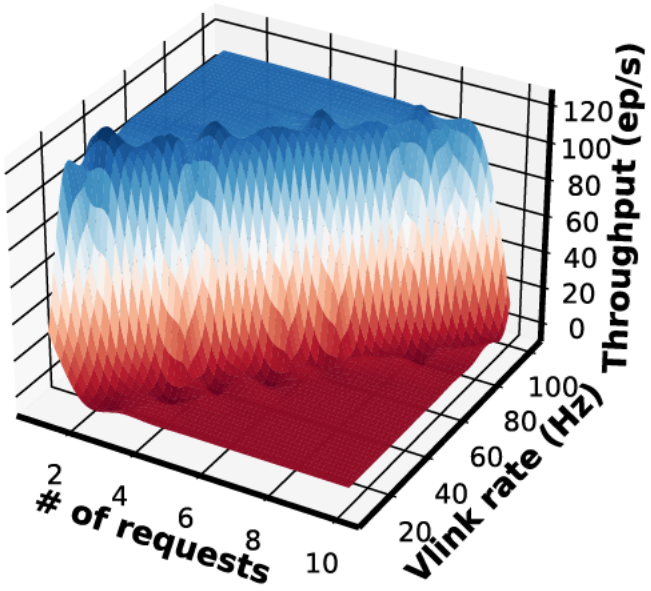

Figure 5.5: Impact of virtual link distribution rate on throughput for the opportunistic setup in figure 4.3, where throughput drops drastically, once an insufficient amount of virtual links is produced.

individual sessions. Conducting the large scale experiment presented in chapter 4 for 20% network utilization yields worse performance outcomes for VS. High variance in data points and bad throughput scaling indicate network congestion and a lack of entanglement resources.

Figure 5.5 illustrates the relationship between the distribution rate of virtual links and the number of sessions a virtual link serves for the setup in figure 4.3. Virtual link establishment has to be at least n-times as fast as state distribution with n being the number of sessions a virtual link serves. Rates below that impact throughput dramatically.

As part of a rough estimation, the ratio of virtual links to distribution sessions can be taken into account in terms of throughput within large-scale simulation data. Figure 5.6 depicts that a generally higher number of virtual links, compared to distribution sessions within the network, boosts throughput significantly. To further confirm the effects of congestion and entanglement bottlenecks, 20% and 80% network utilization will be analyzed at scale. These numbers are chosen to serve as arbitrary high and low

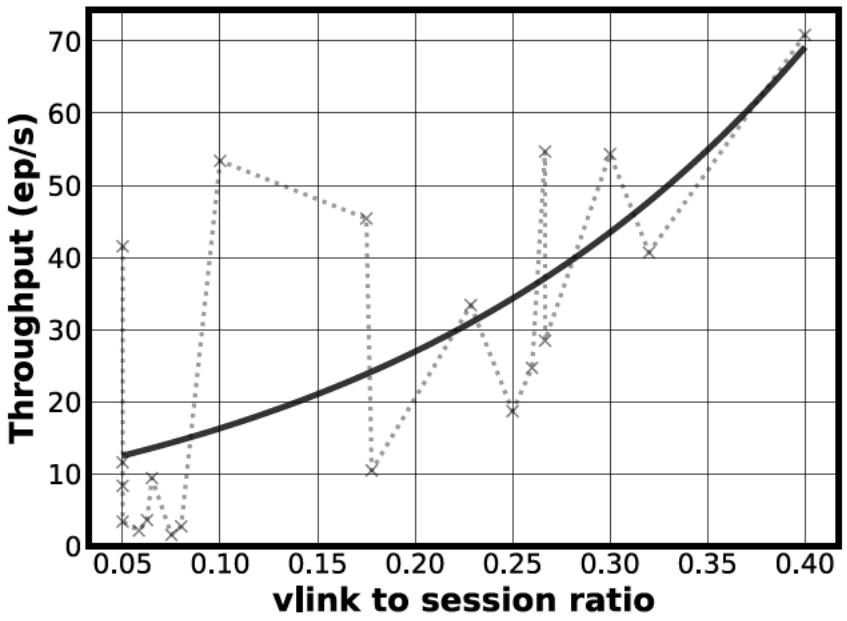

Figure 5.6: $\frac{vlink\ number}{session\ number}$ against throughput in scaled data with static virtual link send rate, in which it is made clear, that a higher amount of virtual links in relation to sessions results in a trend of higher throughput.

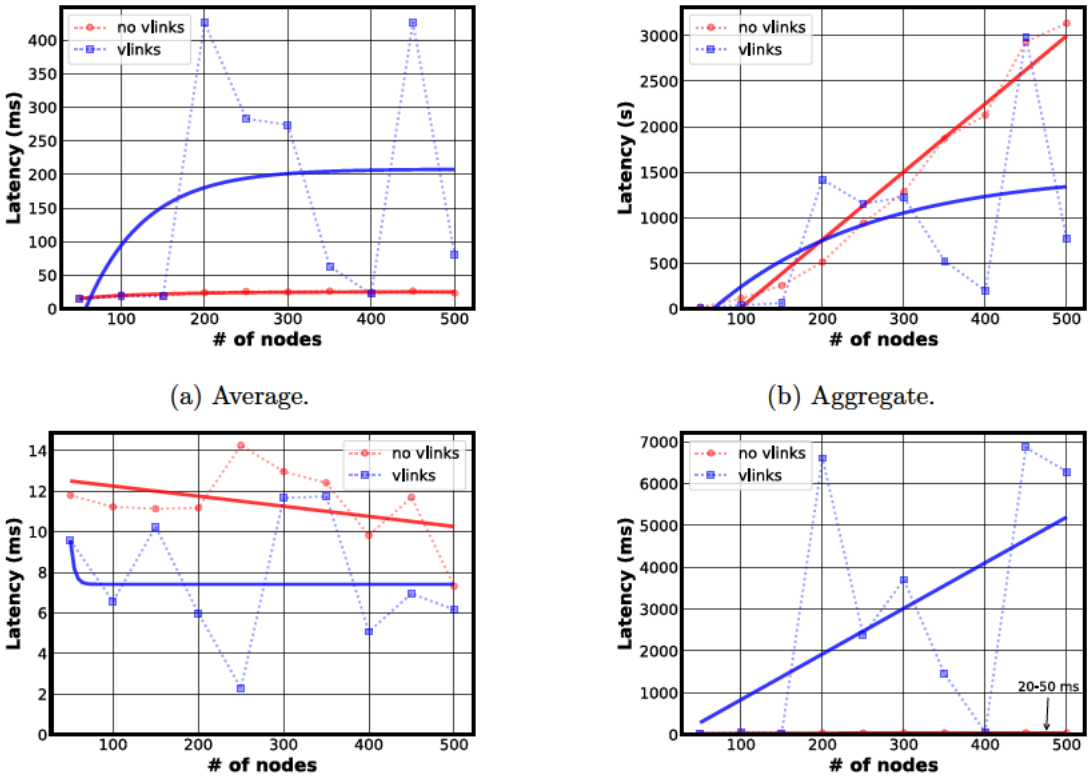

(a) Average.

(b) Aggregate.

(c) Minimum.

(d) Maximum.

Figure 5.7: Latency for 20% utilization for VS *(vlinks)* and PS *(no vlinks)*.

levels of network utilization. Recorded negative effects are expected to get worse with increasing network utilization.

5.2.1 20% Utilization

Refer to figure 5.4 for the average throughput during 20% network utilization. Similar to the throughput graph, the four latency graphs in 5.7 solicit high variance for VS due to the chance of unfavourable bottleneck scenarios. Maximum latency in figure 5.7d can get as slow as ~ 7 s for a single distributed EP using VS, while PS only takes $20-50$ ms in the worst case. Although bottlenecks seem to drastically impact distribution performance, equivalence of average latency for both approaches can be observed in figure 5.7a for some data points, hinting at favourable scenarios. In figure 5.7b, it can be observed that those worst-case outliers are not sufficient to sway the aggregate in a negative direction for VS,

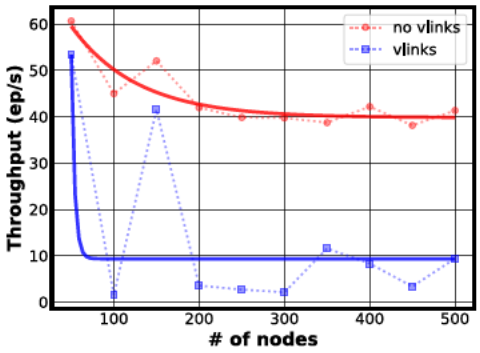

Figure 5.8: Throughput for 80% utilization for VS *(vlinks)* and PS *(no vlinks)*.

going as as low as 200 s, as opposed to PS linearly scaling up to 3000 s. Additionally, the minimum latency still performs better for VS than for PS in all cases, as seen in figure 5.7c. This bolsters the proposition of congestion sensitivity in unfavourable scenarios.

5.2.2 80% Utilization

Consider figure 5.8. As expected, performance plummets as network utilization further increases to 80%. Throughput drops to under 10 EP/s for most data points, making VS more than four times less efficient than PS. The same can be said for latency in figure 5.9, where high variance in the data can be found in a similar manner. The average latency in figure 5.9a is double of that in the previous section. Meanwhile, the maximum latency doubled also up to 14 s for a single distribution, and — despite high variance — does not get near the $20 - 50$ ms latency of the PS approach. Unlike the previous section, the aggregate latency does not withstand congestion and accumulates to much larger levels in figure 5.9b. Unexpectedly, the minimum radically surpasses standard PS for 80% utilization, as well as both approaches in the previous section with 20% utilization, going as low as ~ 0.04 ms. This might stem from the larger number of distributions, which increases the chance of highly favourable sessions to emerge.

This section opens up a non-trivial problem of managing entanglement resources as part of a producer-consumer dynamic. For an initial approach, this issue can be tackled by scaling distribution rate for each virtual link with network utilization. Assuming that on average distribution sessions are equally distributed among virtual links (see figure 5.6),

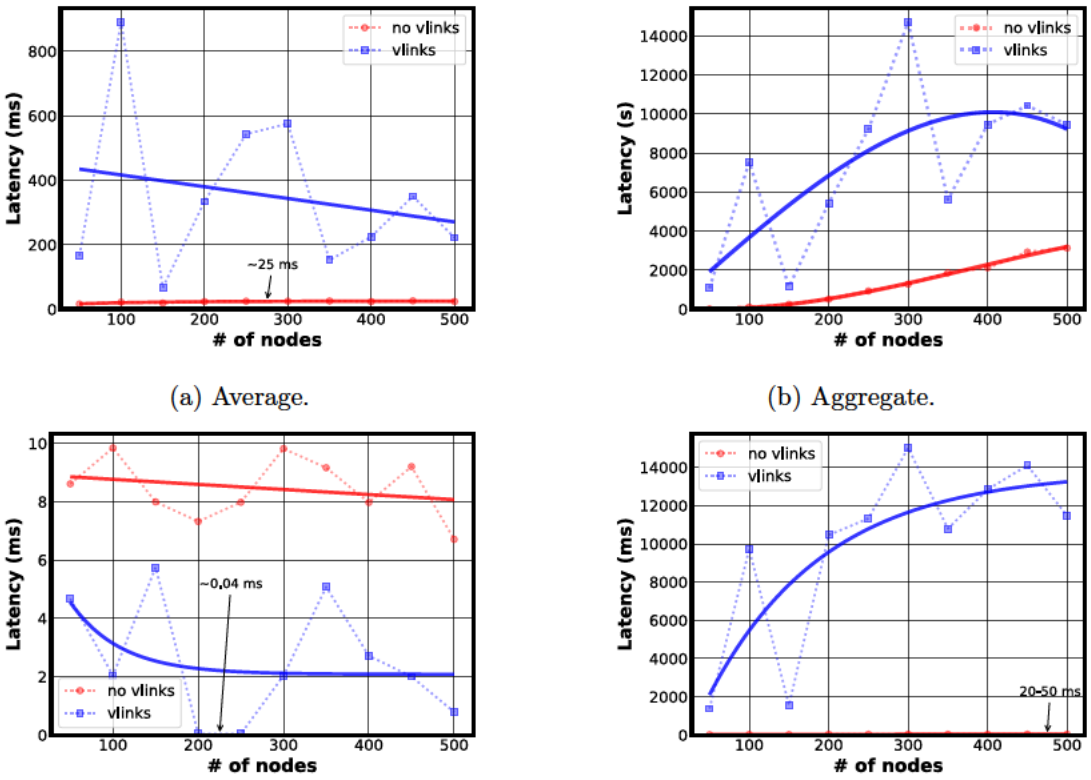

(a) Average.

(b) Aggregate.

(c) Minimum.

(d) Maximum.

Figure 5.9: Latency for 80% utilization for VS *(vlinks)* and PS *(no vlinks)*.

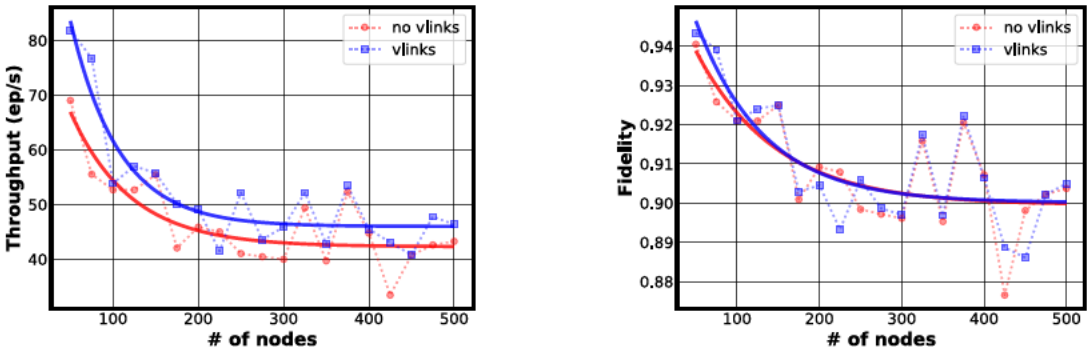

(a) Throughput.

(b) Fidelity.

Figure 5.10: Average throughput and fidelity of EP distribution with VS routing *(vlinks)* and PS routing *(no vlinks)* at scale.

a sufficient distribution rate can be calculated as follows:

$$vlink\ rate = distro\ rate * \frac{session\ count}{vlink\ count}. \tag{5.1}$$

Distro rate is the uniform rate of EP distribution for each session, in this case, 10 Hz. It is expected that despite the strategy being crude, improvements manifest in the experiment at scale. Moreover, this method is justified by the discoveries found in figure 5.6, where a higher ratio already yields greater results.

5.3 Distribution at Scale

This section evaluates the experiment at scale put forth in section 4.2. To attack the resource scarcity showcased in section 5.2, 5.1 is implemented.

Starting with figure 5.10a, the primary research objective of this thesis is met. Provisional topological augmentation through virtual links yields better results in terms of throughput than EP distribution on the physical topology. Note that this improvement occurs despite the naive approach in 5.1 for dealing with congestion. Distribution sessions are selected randomly to meet 20% network utilization. This can lead to highly unfavourable scenarios for virtual link entanglement as shared resources. The employment of virtual links compensates for this on average. This affirms the general utility and net benefit of virtual entanglement routing.

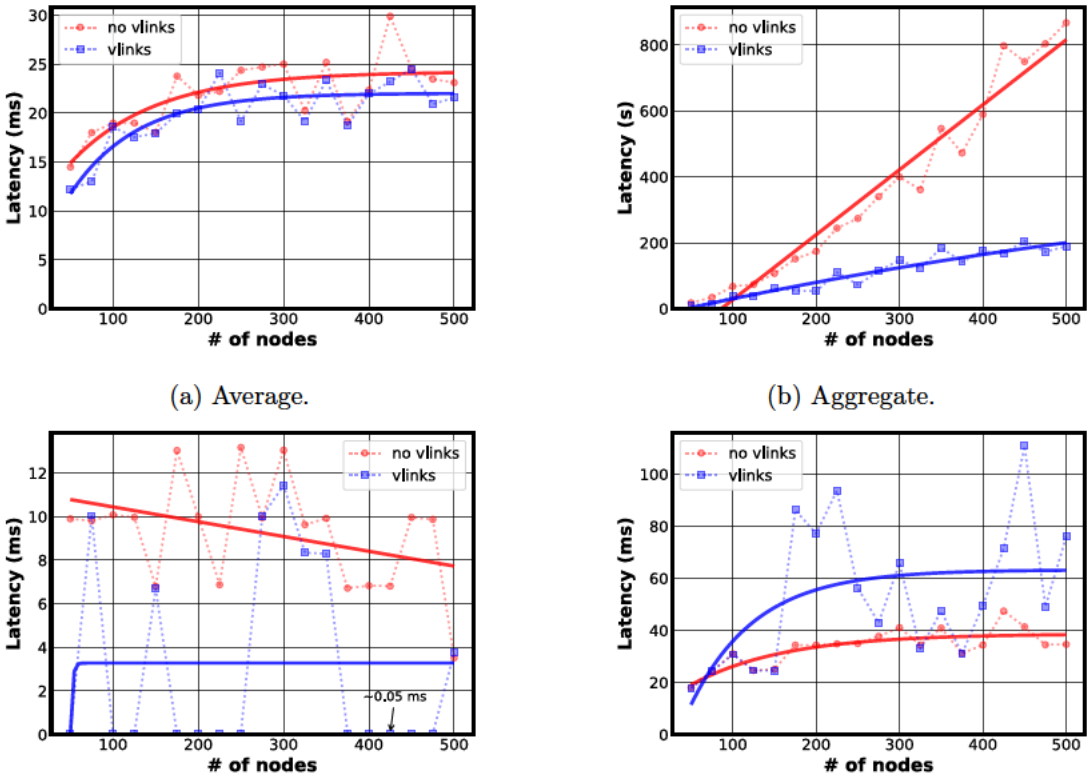

(a) Average.

(b) Aggregate.

(c) Minimum.

(d) Maximum.

Figure 5.11: Latency at scale for VS *(vlinks)* and PS *(no vlinks)*.

Established earlier, the effects of slight fidelity differences in the concept validation might also be visible at scale. However, figure 5.10b exhibits satisfactory results. Average fidelities of both approaches show no significant differences. In combination with performance improvements, this makes a strong case for VS.

Again, latency data provides a more detailed view of the system's behavior. High data variance in figures 5.11d and 5.11c, as well as on average higher maximum latency values in 5.11d, indicates that bottlenecks are still present. However, the marginal minimum latency seen in 5.9c of $\sim$ 0.05ms is achieved through very favourable distribution scenarios. Assuming an average minimum latency of $\sim$ 10ms for PS, VS is 200 times faster than PS in the best case. Figures 5.11a and 5.11b showcase general improvement for the average and aggregate as well.

An additional case for VS routing can be made with the same traffic data features used in section 5.1. Classical and quantum communication is notably reduced for VS, displayed

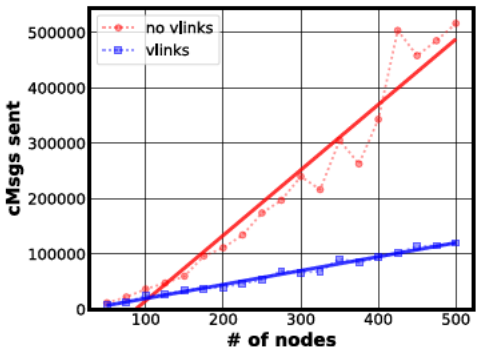

(a) Classical communication aggregate.

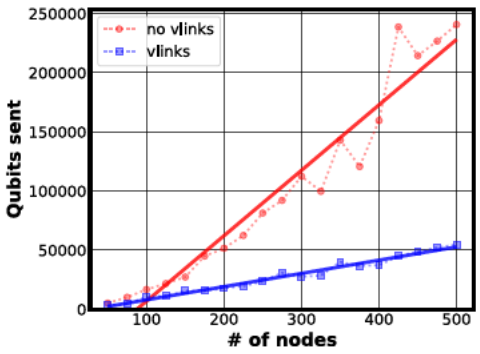

(b) Quantum communication aggregate.

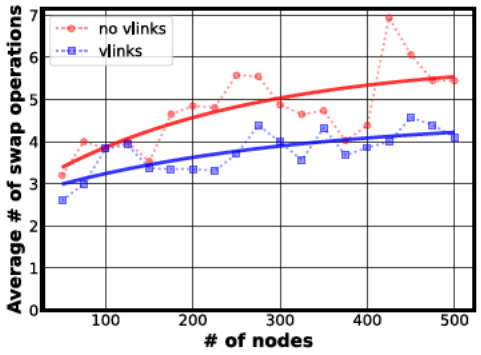

(c) Average number of entanglement swapping operations per EP distribution.

Figure 5.12: Traffic data at scale for VS *(vlinks)* and PS *(no vlinks)*.

in figures 5.12a and 5.12b. Furthermore, number of swaps per distribution can be brought down by $2 - 3$ swap operations on average in figure 5.12c. This outcome is due to the graph diameter scaling logarithmically, as opposed to linearly in the opportunistic setup in section 5.1.

6 Discussion

This chapter interprets the research of this thesis and explores implications, limitations, and future directions. The observations of chapter 5 are first compiled into key findings in section 6.1. This following section 6.2 critically analyzes these findings in reference to the existing literature presented in chapter 3.

6.1 Key Findings

To initiate the chapter, the key findings of chapter 5 are listed:

- **Validated proof of concept:** The proposed solution VS demonstrates conceptual viability in an opportunistic setup.
- **Increased throughput at scale:** The proposed solution VS increases remote EP throughput in large-scale quantum network environments and achieves reduced distribution latency by a factor of 200 in optimal scenarios — the primary research objective of the thesis is met.
- **Generated competitive fidelity:** The proposed solution VS solicits very similar fidelity outcomes for distributed EPs in principle and at scale, as well as drastically reduces swapping operations, which is a driver for decoherence.
- **Reduced network traffic:** The proposed solution VS is associated with a severe reduction in network traffic of both classical messages and qubits sent through physical media.
- **Identified resource constraints:** The proposed solution VS ran into significant challenges in regards to virtual entanglement as a shared resource and emerging bottleneck behavior for unfavourable distribution scenarios, which can be tackled with more sophisticated resource management and broader experiment design scope.

6.2 Contextual Analysis

As a disclaimer, absolute values in evaluation results are discussed in regards to hardware specifics. For example, in [21], a distribution rate of 10 Hz in a 400-node random topology yields an average throughput of 100 EP/s. For a network with no congestion, and an average path length, this result is expected, as throughput and distribution rate are directly proportional. The average path length of a 400-node network can be derived from figure 5.12c, as an average amount of swapping operations of ~ 5 corresponds to a path length of 6. In section 5.1, the same algorithm solicits only 80 EP/s under the same parameters in a shorter 5-node swapping path. Note, that [21] provides no information about individual link length, node processing delay, network edge density, or read/write delays for quantum memory. All of these aspects influence performance and account for major differences in absolute numbers.

Similarly, in [39], throughput is examined using the reference algorithm based on different quantum memory congestion scenarios in a 6-nodes dumbbell topology, i.e., a 4-node swapping path. Throughput relates to memory allocation before distribution and ranges from 20 to 280 EP/s. Scaling these results for an appropriate absolute value comparison is not trivial, even within the literature body itself.

Furthermore, detailed descriptions of decoherence models in memory, as well as physical communication channels are missing in [21]. Especially quantum memory coherence times of the real state-of-the-art hardware seem to constantly change and improve, which impacts remote entanglement fidelity. [39] demonstrates a logarithmic drop in fidelity of 40% at a path length of 10. In figure 5.1b the curve is less steep.

[28] utilizes a heuristic swapping scheme based on expected path performance. The routing algorithm is applied to the elementary logical topology in figure 2.10b. It exhibits similar performance statistics in terms of throughput as the distribution conducted at scale in figure 5.10a, with drastically better scaling in regards to network size. Even the competition for network resources when routing over virtual links, that severely influences performance in a randomized scenario, is rarely addressed. [27] suggests a similarly crude solution to this problem by simply calculating the expected traffic at each virtual link, despite not addressing it as an inherent issue. A more thorough experiment, e.g., averaging over more simulation rounds, generating multiple topologies for each node count, or discarding highly unfavourable scenarios, might lead to an amplification of the functional results in the large scale experiment.

In conclusion, the process of finding comparable simulation parameters is highly non-trivial in a field as dynamic as quantum networking. This is why simulation results in absolute values should always be taken with a grain of salt and are highly contextual, even when using similar tools. This does not prevent analyses within a reference frame, as it is done in the context of this thesis, by directly comparing different algorithms in the same setup. Performance trends and patterns of the same implementation should manifest despite a difference in parameterization.

The thesis makes a strong case for virtual entanglement routing. Results in figures 5.1a verify its conceptual superiority. Moreover, these advantages are found to scale to larger networks in randomized scenarios in figure 5.10a, answering how throughput can be increased for remote entanglement, which is ultimately distributed for a diverse field of QI technologies and applications. The entanglement routing approaches that fall into this category and are associated with analyses based on quantum network simulation conform to this line of reasoning [27, 24, 36].

Virtuality in the context of entanglement distribution yields results that are more favourable than solutions that do not utilize the unique aspects of entanglement-enabled connectivity [12]. Keep in mind, once the paradigm shifts away from heavy dependence on the underlying physical network through abstraction and separation of concern in a potential layered model [23], entanglement routing might always be conducted on a solely virtual graph. The question of whether virtual connectivity is a viable approach might become obsolete. Rather, the discussion might move towards methods of shaping network topology in a opportunistic way and how that can be incorporated into standardization for practical quantum networks. For example, most clustering techniques to tackle the optimization problem of finding suitable virtual links remain unexplored in this context. Topological augmentation could become a foundational mechanism for QI technologies.

7 Conclusion

QI development plans to enhance classical solutions regarding conventional internet protocol and application categories beyond current standards. Remote entanglement represents a unique concept of quantum information processing, that functions as the mechanism to exceed conventional capabilities, as well as the fundamental building block for QI technologies. Due to teleportation as a way to feasibly transmit quantum data in the face of decoherence, entanglement functions as a communication medium. As an extension of teleportation, successive entanglement swapping over repeater nodes establishes the distributed correlation consumed in the teleportation process. The distribution of entanglement resources — in the form of bipartite EPs — to non-adjacent nodes describes the scientific discipline of entanglement routing. This opens up a plethora of different algorithms to facilitate EP delivery within a quantum network, that differ in entanglement swapping schemes, path calculation heuristics, and approaches to topology. Based on the application requirements, entanglement routing protocols have to be either time efficient or deliver entanglement quality standards in the form of high fidelity. The thesis proposed a virtual entanglement routing algorithm VS, that is based on the design principle of topological augmentation, in order to increase throughput of remote entanglement. Using a quantum network simulation tool called *SimQN*, a quantitative evaluation of VS was carried out in reference to the algorithm PS, which had been sourced from the literature body. Here, two different experiment designs were outlined. This thesis found that VS severely outperforms PS in terms of remote EP throughput based on the number of nodes missing in the virtual path that would otherwise be part of a physical path, while having similar statistics in terms remote EP fidelity and greatly reducing network traffic. These results were found to also manifest when scaling both algorithms to large, randomly-generated networks. This presents a positive outlook for virtual entanglement routing schemes, similar to what is observed in related virtual approaches.

For QI development, it is vital to take into account the unique design principles derived from inherent quantum mechanical properties. The field of quantum networking is still

in its infancy and characterized by a vast landscape of simulative case studies, theoretical algorithm proposals and even some practical deployments of QI protocols, such as QKD. The dynamic nature of the literature body creates a lot redundancy and conflicting information in terms of algorithmic approaches, hardware details, and absolute values of simulation results, which creates an opaque state of the art, which could ultimately stump progression speed of QI research. Essentially, the quantum networking field suffers from a lack of real standardization to source from. Therefore, a concerted effort towards establishing standardized frameworks in collaborative, rather than competitive research scenarios, is crucial to accelerate advancements in this promising field.

Bibliography

[1] Amar Abane, Michael Cubeddu, Van Sy Mai, and Abdella Battou. 2024. Entanglement Routing in Quantum Networks: A Comprehensive Survey. https://doi.org/10.48550/arXiv.2408.01234 arXiv:2408.01234 [cs.ET]

[2] Guus Avis, Robert Knegjens, Anders S. Sørensen, and Stephanie Wehner. 2024. Asymmetric node placement in fiber-based quantum networks. *Physical Review A* 109, 5 (May 2024). https://doi.org/10.1103/physreva.109.052627

[3] Thomas Aynaud. 2010. Community detection for NetworkX. https://github.com/taynaud/python-louvain Accessed: September 25, 2024.

[4] Koji Azuma, Sophia E. Economou, David Elkouss, Paul Hilaire, Liang Jiang, Hoi-Kwong Lo, and Ilan Tzitrin. 2022. Quantum repeaters: From quantum networks to the quantum internet. *Reviews of Modern Physics* (2022). https://doi.org/10.48550/arXiv.2212.10820

[5] Oceane Bel and Mariam Kiran. 2024. Simulators for Quantum Network Modelling: A Comprehensive Review. https://doi.org/10.48550/arXiv.2408.11993

[6] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. 1993. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical review letters* 70 13 (1993), 1895–1899. https://doi.org/10.1103/PhysRevLett.70.1895

[7] Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008 (2008), P10008. https://doi.org/10.48550/arXiv.0803.0476

[8] Dirk Bouwmeester, Jian-Wei Pan, Klaus Mattle, Manfred Eibl, Harald Weinfurter, and Anton Zeilinger. 1997. Experimental quantum teleportation. *Nature* 390 (1997), 575–579. https://doi.org/10.48550/arXiv.1901.11004

[9] Max A. Bredig. 1973. The Born-Einstein Letters. *Science* 180 (1973), 1118 – 1118. https://doi.org/10.1126/science.180.4091.1118-b

[10] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. 2013. Bell Nonlocality. https://doi.org/10.48550/arXiv.1303.2849

[11] Angela Sara Cacciapuoti, Marcello Caleffi, Rodney Van Meter, and Lajos Hanzo Hanzo. 2019. When Entanglement Meets Classical Communications: Quantum Teleportation for the Quantum Internet. *IEEE Transactions on Communications* 68 (2019), 3808–3833. https://doi.org/10.1109/TCOMM.2020.2978071

[12] Angela Sara Cacciapuoti, Jessica Illiano, and Marcello Caleffi. 2023. Quantum Internet Addressing. *IEEE Network* 38 (2023), 104–111. https://doi.org/10.48550/arXiv.2306.05982

[13] Angela Sara Cacciapuoti, Jessica Illiano, Seid Koudia, Kyrylo Simonov, and Marcello Caleffi. 2022. The Quantum Internet: Enhancing Classical Internet Services One Qubit at A Time. *IEEE Network* 36 (2022), 6–12. https://doi.org/10.1109/MNET.001.2200162

[14] Marcello Caleffi. 2017. Optimal Routing for Quantum Networks. *IEEE Access* 5 (2017), 22299–22312. https://doi.org/10.1109/ACCESS.2017.2763325

[15] Marcello Caleffi, Daryus Chandra, Daniele Cuomo, Shima Hassanpour, and Angela Sara Cacciapuoti. 2020. The Rise of the Quantum Internet. *Computer* 53 (2020), 67–72. https://doi.org/10.1109/MC.2020.2984871

[16] Kaushik Chakraborty, David Elkouss, Bruno Rijsman, and Stephanie Wehner. 2020. Entanglement Distribution in a Quantum Network: A Multicommodity Flow-Based Approach. *IEEE Transactions on Quantum Engineering* 1 (2020), 1–21. https://doi.org/10.48550/arXiv.2005.14304

[17] Kaushik Chakraborty, Filip Rozpędek, Axel Dahlberg, and Stephanie Wehner. 2019. Distributed Routing in a Quantum Internet. *ArXiv* abs/1907.11630 (2019). https://doi.org/10.48550/arXiv.1907.11630

[18] Alena Chang and Guoliang Xue. 2022. Order Matters: On the Impact of Swapping Order on an Entanglement Path in a Quantum Network. *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*

(2022), 1–6. https://doi.org/10.1109/INFOCOMWKSHPS54753.2022.9798254

[19] Lutong Chen, Kaiping Xue, Jian Li, Ruidong Li, Nenghai Yu, Qibin Sun, and Jun Lu. 2024. Q-DDCA: Decentralized Dynamic Congestion Avoid Routing in Large-Scale Quantum Networks. *IEEE/ACM Transactions on Networking* 32 (2024), 368–381. https://doi.org/10.1109/TNET.2023.3285093

[20] Lutong Chen, Kaiping Xue, Jian Li, Zhonghui Li, Ruidong Li, Nenghai Yu, Qibin Sun, and Jun Lu. 2024. REDP: Reliable Entanglement Distribution Protocol Design for Large-Scale Quantum Networks. *IEEE Journal on Selected Areas in Communications* 42 (2024), 1723–1737. https://doi.org/10.1109/JSAC.2024.3380101

[21] Lutong Chen, Kaiping Xue, Jian Li, Nenghai Yu, Ruidong Li, Qibin Sun, and Jun Lu. 2023. SimQN: A Network-Layer Simulator for the Quantum Network Investigation. *IEEE Network* 37 (2023), 182–189. https://doi.org/10.1109/MNET.130.2200481

[22] Tim Coopmans, Robert Knegjens, Axel Dahlberg, David Maier, Loek Nijsten, Julio de Oliveira Filho, Martijn Papendrecht, Julian Rabbie, Filip Rozpędek, Matthew Skrzypczyk, Leon Wubben, Walter de Jong, Damian Podareanu, Ariana Torres-Knoop, David Elkouss, and Stephanie Wehner. 2020. NetSquid, a NETwork Simulator for QUantum Information using Discrete events. *Communications Physics* 4 (2020). https://doi.org/10.48550/arXiv.2010.12535

[23] Axel Dahlberg, Matthew Skrzypczyk, Tim Coopmans, Leon Wubben, Filip Rozpędek, Matteo Pompili, Arian Stolk, Przemysław Pawełczak, Robert Knegjens, Julio A. de Oliveira Filho, Ronald Hanson, and Stephanie Wehner. 2019. A link layer protocol for quantum networks. *Proceedings of the ACM Special Interest Group on Data Communication* (2019). https://doi.org/10.48550/arXiv.1903.09778

[24] Dibakar Das, Shiva Kumar Malapaka, Jyotsna L. Bapat, and Debabrata Das. 2020. A Proactive Connection Setup Mechanism for Large Quantum Networks. *2022 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT)* (2020), 1–6. https://doi.org/10.48550/arXiv.2012.13566

[25] Stephen DiAdamo, Bing Qi, Glen Miller, Ramana Rao Kompella, and Alireza Shabani. 2022. Packet switching in quantum networks: A path to the quantum Internet. *Physical Review Research* (2022). https://doi.org/10.48550/arXiv.2205.07507

[26] P. A. M. Dirac. 1939. A new notation for quantum mechanics. *Mathematical Proceedings of the Cambridge Philosophical Society* 35, 3 (1939), 416–418. https://doi.org/10.1017/S0305004100021162

[27] Mohammad Ghaderibaneh, Himanshu Gupta, C. R. Ramakrishnan, and Ertai Luo. 2022. Pre-Distribution of Entanglements in Quantum Networks. *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2022), 426–436. https://doi.org/10.1109/QCE53715.2022.00064

[28] Mohammad Ghaderibaneh, Caitao Zhan, Himanshu Gupta, and C. R. Ramakrishnan. 2021. Efficient Quantum Network Communication Using Optimized Entanglement Swapping Trees. *IEEE Transactions on Quantum Engineering* 3 (2021), 1–20. https://doi.org/10.48550/arXiv.2112.11002

[29] Andr'as Gily'en and Alexander Poremba. 2022. Improved Quantum Algorithms for Fidelity Estimation. https://doi.org/10.48550/arXiv.2203.15993

[30] Aric A. Hagberg, Daniel A. Schult, Pieter Swart, and JM Hagberg. 2008. Exploring Network Structure, Dynamics, and Function using NetworkX. *Proceedings of the Python in Science Conference* (2008). https://doi.org/10.25080/TCWV9851

[31] Stav Haldar, Pratik J. Barge, Sumeet Khatri, and Hwang Lee. 2023. Fast and reliable entanglement distribution with quantum repeaters: Principles for improving protocols using reinforcement learning. *Physical Review Applied* (2023). https://doi.org/10.48550/arXiv.2303.00777

[32] Cinthya Hernández, Miriam Portillo, Erick Sánchez-Gaitán, Francisco Delgado, and Alan Anaya. 2024. Quantum Key Distribution Shared Protocol Using Teleportation and Delayed Measurement. *Journal of Physics: Conference Series* 2701 (2024). https://doi.org/10.1088/1742-6596/2701/1/012111

[33] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. 2007. Quantum entanglement. https://doi.org/10.48550/arXiv.quant-ph/0702225

[34] Jessica Illiano, Marcello Caleffi, Antonio Manzalini, and Angela Sara Cacciapuoti. 2022. Quantum Internet Protocol Stack: a Comprehensive Survey. *Comput. Networks* 213 (2022), 109092. https://doi.org/10.48550/arXiv.2202.10894

[35] Alastair Kay. 2018. Tutorial on the Quantikz Package. *arXiv: Quantum Physics* (2018). https://doi.org/10.48550/arXiv.1809.03842

[36] Alexander Kolar, Allen Zang, Joaquín Chung, Martin Suchara, and Rajkumar Kettimuthu. 2022. Adaptive, Continuous Entanglement Generation for Quantum Networks. *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)* (2022), 1–6. https://doi.org/10.1109/INFOCOMWKSHPS54753.2022.9798130

[37] Wojciech Kozlowski, Axel Dahlberg, and Stephanie Wehner. 2020. Designing a quantum network protocol. *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies* (2020). https://doi.org/10.48550/arXiv.2010.02575

[38] Wojciech Kozlowski, Stephanie Wehner, Rodney Van Meter, Bruno Rijsman, Angela Sara Cacciapuoti, Marcello Caleffi, and Shota Nagayama. 2023. Architectural Principles for a Quantum Internet. RFC 9340. https://doi.org/10.17487/RFC9340

[39] Jian Li, Qidong Jia, Kaiping Xue, David S. L. Wei, and Nenghai Yu. 2022. A Connection-Oriented Entanglement Distribution Design in Quantum Networks. *IEEE Transactions on Quantum Engineering* 3 (2022), 1–13. https://doi.org/10.1109/TQE.2022.3176375

[40] Jian Li, Ming jian Wang, Kaiping Xue, Ruidong Li, Nenghai Yu, Qibin Sun, and Jun Lu. 2021. Fidelity-Guaranteed Entanglement Routing in Quantum Networks. *IEEE Transactions on Communications* 70 (2021), 6748–6763. https://doi.org/10.48550/arXiv.2111.07764

[41] Puli Li and Weihai Li. 2024. Traffic-Aware Routing Algorithm in Quantum Network. *2024 9th International Conference on Computer and Communication Systems (ICCCS)* (2024), 706–711. https://doi.org/10.1109/ICCCS61882.2024.10603087

[42] J. B. MacQueen. 1967. *Some methods for classification and analysis of multivariate observations.* University of California Press, 281–297.

[43] Tu N. Nguyen, Kashyab J. Ambarani, Linh Le, Ivan Djordjevic, and Zhi-Li Zhang. 2022. A Multiple-Entanglement Routing Framework for Quantum Networks. https://doi.org/10.48550/arXiv.2207.11817 arXiv:2207.11817 [cs.NI]

[44] Michael A. Nielsen and Isaac L. Chuang. 2011. *Quantum Computation and Quantum Information: 10th Anniversary Edition* (10th ed.). Cambridge University Press, USA.

[45] Michael A. Nielsen, Michael A. Nielsen, Emanuel Knill, and Raymond Laflamme. 1998. Complete quantum teleportation using nuclear magnetic resonance. *Nature* 396 (1998), 52–55. https://doi.org/10.1038/23891

[46] Peter Kaufmann Xavier Jeannin Tim Chown Ivana Golub Domenico Vicinanza Guy Roberts Rudolf Vohnout Pavel Skoda Josef Vojtech Piotr Rydlichowski, Susanne Naegele-Jackson. 2021. Quantum Technologies Status Overview. https://api.semanticscholar.org/CorpusID:231827750

[47] Matteo Pompili, Carlo Delle Donne, Ingmar te Raa, Bart van der Vecht, Matthew Skrzypczyk, Guilherme Ferreira, Lisa de Kluijver, Arian J. Stolk, Sophie L. N. Hermans, Przemyslaw Pawelczak, Wojciech Kozlowski, Ronald Hanson, and Stephanie Wehner. 2021. Experimental demonstration of entanglement delivery using a quantum network stack. *npj Quantum Information* 8 (2021), 1–10. https://doi.org/10.48550/arXiv.2111.11332

[48] Shahrooz Pouryousef, Nitish K. Panigrahy, and Donald F. Towsley. 2022. A Quantum Overlay Network for Efficient Entanglement Distribution. *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications* (2022), 1–10. https://doi.org/10.48550/arXiv.2212.01694

[49] QuTech. [n. d.]. Quantum Network Explorer. https://www.quantum-network.com/ Accessed: 2024-10-03.

[50] QuTech. 2020. NetSquid Snippets. https://netsquid.org/snippets/. Accessed: September 25, 2024.

[51] Ji-Gang Ren, Ping Xu, Hai-Lin Yong, Liang Zhang, Sheng-Kai Liao, Juan Yin, Wei-Yue Liu, Wen-Qi Cai, Meng Yang, Li Li, Kui-Xing Yang, Xuan Han, Yong-Qiang Yao, Ji Li, Hai-Yan Wu, Song Wan, Lei Liu, Ding-Quan Liu, Yao-Wu Kuang, Zhi-Ping He, Peng Shang, Cheng Guo, Ru-Hua Zheng, Kai Tian, Zhen-Cai Zhu, Nai-Le Liu, Chao-Yang Lu, Rong Shu, Yu-Ao Chen, Cheng-Zhi Peng, Jian-Yu Wang, and

Jian-Wei Pan. 2017. Ground-to-satellite quantum teleportation. *Nature* 549, 7670 (Aug. 2017), 70–73. https://doi.org/10.1038/nature23675

[52] Eleanor Rieffel and Wolfgang Polak. 2011. Quantum Computing: A Gentle Introduction. The MIT Press. https://api.semanticscholar.org/CorpusID:63758232

[53] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 26 (1978), 96–99. https://doi.org/10.1145/359340.359342

[54] Eddie Schoute, Laura Maninska, Tanvirul Islam, Iordanis Kerenidis, and Stephanie Wehner. 2015. Shortcuts to quantum network routing. *ArXiv* abs/1610.05238 (2015). https://doi.org/10.48550/arXiv.1610.05238

[55] P.W. Shor. 1994. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 124–134. https://doi.org/10.1109/SFCS.1994.365700

[56] Abhishek Shukla, Boo Carmans, Michael Petrov, Daan Vrancken, and Milos Nesladek. 2024. High fidelity quantum state tomography of electron-^{14}N nuclear hybrid spin register in diamond using Rabi oscillations. https://doi.org/10.48550/arXiv.2408.13349 arXiv:2408.13349 [quant-ph]

[57] Harpreet Singh, Harpreet Singh, Arvind, and Kavita Dorai. 2020. Using a Lindbladian approach to model decoherence in two coupled nuclear spins via correlated phase damping and amplitude damping noise channels. *Pramana* 94 (2020). https://doi.org/10.48550/arXiv.2007.12972

[58] Shraddha Singh, Mina Doosti, Natansh Mathur, Mahshid Delavar, A. N. Mantri, Harold Ollivier, and Elham Kashefi. 2023. Towards a Unified Quantum Protocol Framework: Classification, Implementation, and Use Cases. https://doi.org/10.48550/arXiv.2310.12780

[59] Andrew S. Tanenbaum and David Wetherall. 2011. Computer networks, 5th Edition. https://api.semanticscholar.org/CorpusID:31375788

[60] Chonggang Wang, Akbar Rahman, Ruidong Li, Melchior Aelmans, and Kaushik Chakraborty. 2024. Application Scenarios for the Quantum Internet. RFC 9583. https://doi.org/10.17487/RFC9583

[61] Stephanie Wehner, David Elkouss, and Ronald Hanson. 2018. Quantum internet: A vision for the road ahead. *Science* 362 (2018). https://doi.org/10.1126/science.aam9288

[62] Reinhard F. Werner. 1989. Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model. *Physical review. A, General physics* 40 8 (1989), 4277–4281. https://doi.org/10.1103/PhysRevA.40.4277

[63] Zirui Xiao, Jian Li, Kaiping Xue, Zhonghui Li, Nenghai Yu, Qibin Sun, and Jun Lu. 2024. A Connectionless Entanglement Distribution Protocol Design in Quantum Networks. *IEEE Network* 38 (2024), 131–139. https://doi.org/10.1109/MNET.2023.3321044

[64] Yangming Zhao, Gongming Zhao, and Chunming Qiao. 2022. E2E Fidelity Aware Routing and Purification for Throughput Maximization in Quantum Networks. *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications* (2022), 480–489. https://doi.org/10.1109/INFOCOM48880.2022.9796814

[65] Marek Żukowski, Anton Zeilinger, Michael A. Horne, and Artur Ekert. 1993. "Event-ready-detectors" Bell experiment via entanglement swapping. *Physical review letters* 71 26 (1993), 4287–4290. https://doi.org/10.1103/PhysRevLett.71.4287

Erklärung zur selbständigen Bearbeitung

Hiermit versichere ich, dass ich die vorliegende Arbeit ohne fremde Hilfe selbständig verfasst und nur die angegebenen Hilfsmittel benutzt habe. Wörtlich oder dem Sinn nach aus anderen Werken entnommene Stellen sind unter Angabe der Quellen kenntlich gemacht.

Ort Datum Unterschrift im Original