

Hardware and Protocol Optimization in Quantum-Repeater Networks

Horta Ferreira da Silva, F.

DOI

[10.4233/uuid:45895388-2e1c-41de-88c3-fa06d6ab29ea](https://doi.org/10.4233/uuid:45895388-2e1c-41de-88c3-fa06d6ab29ea)

Publication date

2023

Document Version

Final published version

Citation (APA)

Horta Ferreira da Silva, F. (2023). *Hardware and Protocol Optimization in Quantum-Repeater Networks*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:45895388-2e1c-41de-88c3-fa06d6ab29ea>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Hardware and Protocol Optimization in Quantum-Repeater Networks

Hardware and Protocol Optimization in Quantum-Repeater Networks

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen
op maandag 4 december 2023 om 15.00 uur

door

Francisco FERREIRA DA SILVA

Master of Science in Engineering Physics,
Instituto Superior Técnico, University of Lisbon, Portugal,
geboren te Vila Nova de Gaia, Portugal.


Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. S.D.C. Wehner,	Technische Universiteit Delft, promotor
Prof. dr. ir. R. Hanson,	Technische Universiteit Delft, promotor

Onafhankelijke leden:

Prof. dr. G. A. Steele,	Technische Universiteit Delft
Prof. dr. J. Laurat,	Sorbonne Université, France
Prof. dr. E. Kashefi,	Sorbonne Université, France
Dr. P. Pawelczak,	Technische Universiteit Delft
Prof. dr. A.F. Otte,	Technische Universiteit Delft, reservelid



Printed: on recycled paper by Ridderprint | www.ridderprint.nl

Cover: created by Alexander Naughton | www.alexandernaughton.com

Copyright ©2023 by Francisco Ferreira da Silva
ISBN 978-94-6483-573-1

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

Contents

Summary	xi
Samenvatting	xiii
1 Reader's guide	1
2 Introduction	3
3 Preliminaries	9
3.1 Operational principles of quantum repeaters	9
3.1.1 Entanglement: Bell states	9
3.1.2 Heralded entanglement generation.	10
3.1.3 Entanglement swapping	10
3.1.4 Entanglement purification	11
3.1.5 Scaling of quantum repeaters.	11
3.1.6 Quantum memories	12
3.1.7 Multiplexing	12
3.2 Other types of repeaters	13
3.3 Quantifying repeater performance	14
3.3.1 Application-derived performance metrics	14
4 Counting quantum-repeater configurations	23
4.1 Introduction	23
4.2 Results	24
4.2.1 Scenario 1: no midpoint stations	24
4.2.2 Scenario 2: with midpoint stations but not counting them	25
4.2.3 Scenario 3: with midpoint stations and counting them	25
4.3 Proofs	25
4.3.1 Scenario 1: no midpoint stations	25
4.3.2 Scenario 2: not counting midpoint stations.	25
4.3.3 Scenario 3: counting midpoint stations.	27
5 Optimizing entanglement generation and distribution using genetic algorithms	29
5.1 Introduction	30
5.2 Methodology.	31
5.2.1 Question	31
5.2.2 Cost	31
5.2.3 Abstract model.	33
5.2.4 Genetic algorithms	34
5.2.5 <i>smart-stopos</i>	34

5.2.6	Process overview	35
5.2.7	Challenges in applying genetic algorithms to quantum systems . . .	35
5.3	Validation	38
5.3.1	Benchmarking genetic algorithms	39
5.3.2	Validating on Werner chains	40
5.4	Evaluation: use cases	42
5.4.1	Results	44
5.5	Conclusions	49
5.6	Data availability	50
5.7	Code availability	50
5.8	smart-stopos	50
5.9	Genetic algorithms	51
5.10	Abstract model validation	53
5.10.1	Matching to nitrogen-vacancy center model	53
5.10.2	Comparison of nitrogen-vacancy center and abstract models	55
5.11	Werner chains	58
5.12	Computing baseline values in the abstract model	60
5.12.1	Uniform spacing	60
5.12.2	Real network	61
5.13	Search space reduction using previous runs	62
6	Requirements for a processing-node quantum repeater on a real-world fiber grid	67
6.1	Results	68
6.1.1	Quantum-network path	69
6.1.2	Blind quantum computation	72
6.1.3	Minimal hardware requirements	74
6.1.4	Absolute minimal requirements	74
6.2	Discussion	79
6.2.1	Hardware requirements in simplified settings	79
6.2.2	Entanglement without a repeater	82
6.2.3	Outlook	82
6.3	Methods	82
6.3.1	Conditions on network path to enable VBQC	82
6.3.2	Average teleportation fidelity	82
6.3.3	Hardware improvement for VBQC as an optimization problem . . .	83
6.3.4	Optimization parameters	84
6.3.5	Evaluating hardware quality	86
6.3.6	Framework for simulating quantum repeaters	87
6.3.7	Finding minimal hardware improvements	87
6.3.8	Finding absolute minimal hardware requirements	88
6.4	Data availability	88
6.5	Code availability	88
6.6	Setup	88
6.6.1	Fiber network and node placement	88
6.6.2	Repeater protocol	89

6.6.3	Quantum-computing server	91
6.6.4	Processing nodes	92
6.6.5	Color centers	92
6.6.6	Trapped ions	94
6.6.7	Abstract nodes	100
6.6.8	Entanglement generation.	104
6.7	Target metric.	105
6.7.1	Teleportation fidelity.	105
6.7.2	Requirements from VBQC	106
6.7.3	Proving Theorem 1	108
6.7.4	Proving Theorem 2	113
6.7.5	Remote state preparation.	114
6.8	Double-click model.	122
6.8.1	Model assumptions.	122
6.8.2	POVMs.	124
6.8.3	Results without coincidence window.	126
6.8.4	Results with coincidence window	128
6.9	Effect of detection and coincidence time windows	129
6.10	Single-click model	143
6.10.1	Model assumptions.	143
6.10.2	Results	144
6.11	Optimization method.	146
6.11.1	Termination criteria for genetic algorithms.	146
6.11.2	Cost function.	147
6.11.3	Probabilities of no-imperfection	148
6.11.4	Optimizing over tunable parameters	150
6.12	Simulation performance	152
6.13	Framework for simulating quantum repeaters	154
6.13.1	Services	154
6.13.2	SWAP-ASAP protocol	155
6.13.3	Configuring quantum networks	157
6.14	Extra optimization results	158
6.14.1	To move or not to move	158
6.14.2	Architecture comparison	158
6.14.3	Connecting Delft and Eindhoven without a repeater	160
6.14.4	Hardware requirements for repeaters with single and double-click entanglement generation.	162
6.14.5	Hardware improvement costs	163
7	Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber	173
7.1	Introduction	173
7.1.1	Setup.	174
7.1.2	Applications	177
7.1.3	Minimal hardware requirements	178
7.1.4	State-of-the-art parameters.	179

7.1.5	Determining minimal hardware requirements	179
7.2	Impact of number of repeaters on hardware requirements	180
7.2.1	Absolute minimal number of multiplexing modes	180
7.2.2	Minimal hardware requirements for quantum-key distribution.	181
7.2.3	Secret-key rate: quantum-bit error rate and entanglement generation rate	183
7.3	Impact of target on hardware requirements	184
7.3.1	Requirements for different secret-key-rate targets	185
7.3.2	Requirements for secret-key and blind-quantum-computing success rates.	186
7.4	Conclusion.	187
7.5	Data availability	188
7.6	Code availability	188
7.7	Baseline parameters	188
7.8	Repeater placement chosen by optimization method	189
7.9	Optimization method.	191
7.9.1	No-imperfection probabilities	191
7.9.2	Local optimization	192
7.9.3	Performing the optimization	193
7.10	BQC test protocol	193
8	Reducing entanglement-distribution hardware requirements via joint hardware-protocol optimization	199
8.1	Introduction	199
8.2	Methodology.	200
8.2.1	Hardware model	200
8.2.2	Protocols for end-to-end entanglement generation	202
8.2.3	Optimization algorithm	206
8.2.4	State-of-the-art parameters.	207
8.3	Results	208
8.3.1	Optimal hardware cost	208
8.3.2	Optimal hardware parameters	210
8.4	Conclusion.	213
8.5	Fidelity requirements for quantum key distribution	213
8.6	Waiting Time and maximum internode distance	214
8.7	Optimization Algorithm	217
8.8	Validation	219
9	Protocols for generation of high-quality entanglement using reinforcement learning	225
9.1	Introduction	225
9.2	Preliminaries.	226
9.3	Prior work	228
9.4	Our tool	229
9.4.1	Purification	229
9.4.2	Purify, entangle and discard	232

9.5	Conclusions	235
10	Outlook	239
10.1	Summary of results.	239
10.2	Future work	240
	Acknowledgments	243
	Curriculum Vitæ	249
	List of Publications	251

Summary

The future quantum internet promises to enable users all around the world to, among other applications, generate shared secure keys and perform distributed quantum computations. To do so, entanglement must be distributed between remote users. One way of doing this is by sending photons through optical fiber, which allows for reusing some existent classical infrastructure. However, the probability of photons being absorbed in optical fiber grows exponentially with the distance covered, rendering entanglement generation at larger-than-metropolitan scales unfeasible. One possible approach to enable distributing entanglement over larger distances is to employ quantum repeaters, devices that can in theory mitigate the effects of fiber loss by splitting the total distance to be covered into smaller segments. Despite recent advances, the required technology is still under development. In this dissertation we aim to contribute to a swifter realization of fiber-based quantum-repeater networks.

To this end, we introduce a methodology combining quantum-network simulations and genetic-algorithm-based optimizations that allows for determining hardware requirements for quantum repeaters. Using this methodology we translate quantum-network-application-derived performance metrics into specific requirements on the quantum repeaters used to implement the quantum network. This indicates not only how good hardware must be in order to enable given applications, but also in what specific ways state-of-the-art hardware must be improved to do so.

We also investigate the effects of using existing fiber infrastructure for the deployment of near-term quantum networks. Doing so would be a cost-effective way of constructing quantum networks. However, existing infrastructure also imposes constraints, namely on where quantum hardware can be placed. We quantify to what extent such constraints affect quantum-network performance, as well as how these effects can be mitigated by optimizing repeater placement.

Finally, we contribute to answering the question of how to extract the best possible performance out of imperfect hardware. For a given hardware quality, making the right choices with regards to what protocols are executed by the nodes and where nodes are placed can result in significant boosts in performance. We perform a joint hardware-protocol optimization and find that good hardware choices can significantly relax hardware requirements, as well as highlight multiple possible paths to functional quantum-repeater networks. We also provide tools for the discovery of entanglement generation protocols.

Samenvatting

Het toekomstige kwantuminternet belooft gebruikers over de hele wereld in staat te stellen om, naast andere toepassingen, gezamenlijke beveiligde sleutels te genereren en gedistribueerde quantum berekeningen uit te voeren. Om dit te doen moet verstrengeling worden verdeeld tussen uiteenliggende gebruikers. Een manier om dit te doen is door fotonen te sturen via glasvezelkabels, wat het hergebruik van bestaande klassieke infrastructuur mogelijk maakt. Echter neemt de kans dat fotonen worden geabsorbeerd in de glasvezelkabels exponentieel toe met de afgelegde afstand, wat het genereren van verstrengeling op schalen groter dan stedelijke afstanden onhaalbaar maakt. Een mogelijke manier om het verspreiden van verstrengeling over grotere afstanden mogelijk te maken is door het gebruik van quantumrepeaters. Quantumrepeaters zijn apparaten die, in theorie, de effecten van verlies in de kabels kunnen verminderen door de totale te overbruggen afstand in kleinere segmenten te verdelen. Ondanks recente vooruitgang is de technologie hiervoor nog in ontwikkeling. In dit proefschrift streven we ernaar bij te dragen aan een snellere realisatie van op glasvezelkabels gebaseerde kwantumrepeater-netwerken.

Hiertoe introduceren we een methodologie die kwantumnetwerksimulaties en optimalisaties op basis van genetische algoritmen combineert en die het mogelijk maakt om hardwarevereisten voor quantumrepeaters te bepalen. Met behulp van deze methodologie vertalen we prestatie-indicatoren die zijn afgeleid van kwantumnetwerktoepassingen naar specifieke vereisten voor de quantumrepeaters die worden gebruikt om het kwantumnetwerk te implementeren. Dit laat niet alleen zien hoe goed de hardware moet zijn om bepaalde toepassingen mogelijk te maken, maar ook op welke specifieke manieren de state-of-the-art hardware moet worden verbeterd om dit te doen.

Ook onderzoeken we de effecten van het gebruik van bestaande glasvezelkabelinfrastructuur voor de uitrol van near-term kwantumnetwerken. Dit zou een kosteneffectieve manier zijn om kwantumnetwerken te bouwen. Echter legt bestaande infrastructuur ook beperkingen op, met name in de plaatsingsmogelijkheden van de kwantumhardware. We kwantificeren in hoeverre dergelijke beperkingen de prestaties van het kwantumnetwerk beïnvloeden, en hoe deze effecten kunnen worden verminderd door de plaatsing van repeaters te optimaliseren.

Tot slot dragen we bij aan het beantwoorden van de vraag hoe we met imperfecte hardware de best mogelijke prestaties kunnen behalen. Voor een bepaalde hardwarekwaliteit kunnen de juiste keuzes met betrekking tot welke protocollen door de deelnemers worden uitgevoerd en waar quantumapparaten worden geplaatst, leiden tot aanzienlijke verbeteringen in de prestaties. We voeren een gezamenlijke hardware-protocoloptimalisatie uit en vinden dat goede hardwarekeuzes de hardware-eisen aanzienlijk kunnen versoepelen, evenals meerdere mogelijke paden naar functionele quantumrepeater-netwerken kunnen belichten. We geven ook hulpmiddelen voor het ontdekken van protocollen voor het genereren van verstrengeling.

1

Reader's guide

In the latter stages of the writing of this dissertation I attended my sister's wedding, where I had a conversation with a family friend whom I had not seen in a long time. In explaining what I had been up to during the past few years, it became clear that this person was surprised to find that, paraphrasing, *people are still discovering new things*. This might seem almost risible to an academic, which should describe most people reading this dissertation. But it is a natural thought in those who are not. It's completely understandable if someone who ended their formal education in high school believes that calculus (developed in its modern form in the 17th century) and electromagnetism (unified by Maxwell in the 19th century) represent the pinnacle of human knowledge in mathematics and physics. With this in mind, finding out that *people are still new discovering things* might very well then also constitute a discovery in itself.

I imagine that some non-academics might try to read this dissertation. I also imagine that it could prove a challenging read. Part of the challenge is to do with content, and that has no easy solution. But there is also the particular way in which academics communicate with one another. This becomes very natural if one is involved in an academic environment, as I have been for the past four years. However, it is borderline impenetrable to outsiders. I will in this chapter try to explain what is the purpose of a doctoral dissertation, how this particular one is structured and how to read it. If you are an academic, there is nothing new for you in this chapter.

A doctoral dissertation is the culmination of a PhD. A PhD, short for *doctor of philosophy*¹, is an academic degree. Other examples of academic degrees that might be more readily recognized are bachelor's degrees and master's degrees. The main difference between these and a PhD is that PhDs require that original research be produced. In other words, this means that to receive a PhD one has to *discover new things*. The doctoral dissertation is the document in which a PhD candidate reports on the new things they have discovered.

Chapters 4 to 9 of this dissertation consist of new things. Even though a doctoral dissertation is individual, modern scientific work is typically collaborative. All of the new

¹Philosophy does not in this case refer to the discipline of philosophy, but to the broader sense derived from the original Greek work, meaning "pursuit of knowledge".

things in this dissertation were discovered in collaboration with others. Their names are written at the beginning of each chapter. Chapters 2 and 3 do not describe new things. Chapter 2 attempts to place the work on which this dissertation reports in context. Chapter 3 gives a brief introduction to the technical knowledge required to understand this dissertation. Chapter 3 is likely not the best resource to acquire the required knowledge; I included it due to convention and for completeness.

Academics communicate their results to each other via scientific articles, commonly known as *papers*. Chapters 4 to 9 of this dissertation are either based on papers or will be the basis of future papers. Papers are not typically read from start to finish, as one would a novel. Instead, one typically starts by reading the title and maybe the abstract (the short summary of the paper immediately following the title). For most papers, one does not read further. If the paper turns out to be interesting and relevant enough, an academic might proceed to read section titles, look at plots and a summary of results. This is the last step for the vast majority of papers. For the very few that one can absolutely not afford to not understand fully, a proper readthrough follows. This likely involves skimming through the easy or unimportant parts and many rereads of the hard, important parts. Identifying which are which is a skill painstakingly developed through reading many papers. If you wish to read this dissertation, I suggest you do so as an academic would: start with the titles and the abstracts. If something catches your eye, dive deeper. If not, have a look at the acknowledgements! Those are always my favorite part of a dissertation.

2

Introduction

The development of quantum mechanics was one of the greatest scientific achievements of the twentieth century. It was created to interpret phenomena unexplainable by previous theories and it describes the behavior of physics at small scales, particularly atomic and subatomic. One key property of quantum mechanics is that of entanglement, through which the properties of separate quantum systems become linked. This renders their independent description inherently incomplete [1].

Quantum mechanics is not only of theoretical interest. In fact, the insights delivered by quantum mechanics have already led to what is sometimes retroactively referred to as a quantum revolution [2]. In less grandiose terms, this means that there are very impactful technologies whose working principles depend on quantum mechanics, such as lasers and transistors.

Transistors are the (tiny) building blocks of modern computers. Earlier computers were mostly built with vacuum tubes, devices that could control the flow of electricity. The development of the transistor allowed for the same principle to be implemented on a much smaller scale. This proved to be fundamental to the scaling and consequent growing usefulness of computers: if we can all now have smart devices in our pockets, it is because transistors can be made incredibly small (for reference, modern smartphones have about 16 billion transistors [3]).

Although the working principle of transistors relies on quantum mechanics, it does not require that single quantum systems, such as atoms, be individually controlled. Doing so presents a much greater challenge, but also has the potential for enabling exciting novel technologies: quantum computers and quantum networks. The advent of these technologies has sometimes been referred to as the second quantum revolution [2].

Quantum computers manipulate the states of quantum systems to, at least in theory, perform certain computations more efficiently than what is possible with *normal computers* (usually called *classical computers* in the quantum community) [4–7]. The best-known example of a quantum-network application is that of quantum key distribution (QKD), through which users can be provided with a mathematically-secure key that can be used for cryptographic protocols [8–10]. Other examples include secret sharing [11], improved

telescope image resolution [12], more accurate clock synchronization [13] and secure remote access to quantum computers [14, 15].

The technologies of the second quantum revolution can be said to be in their vacuum-tube state. There have been proof-of-principle experiments of both quantum computers [16] and quantum networks [17], as well as the first small commercial products [18]. These are however bulky, hard to operate and in general not very useful.

The main task of a quantum network is to provide potentially-remote users with entanglement, which they can consume to execute applications [19]. This marks a shift in how we think of entanglement: besides a fundamental phenomenon of nature, it is a resource to be created, managed and consumed [20]. Entanglement is a remarkably hard resource to generate. This dissertation aims to contribute to the development of quantum networks by investigating how entanglement can be best be generated and distributed. To better understand how, let us examine exactly why entanglement is so hard to come by.

Entanglement between remote quantum systems (hereinafter quantum nodes) is typically established through the use of photons [21]. One possible way this can be done is as follows: first, entanglement is generated between each of the quantum nodes and a photon. Then, the photons are sent, typically through optical fiber, to a station placed somewhere between the two nodes. There, the photons are made to interact and are measured. This process, known as an entanglement swap, ensures that the quantum nodes are entangled. Photons are used for this purpose as they travel fast and are not very interactive. However, the probability that photons are successfully transmitted through optical fiber decays exponentially with distance [22]. This renders transmitting them, and hence establishing entanglement, over long distances challenging.

Photons are also sent through optical fibers for classical communication. The problem of absorption is in that case solved by amplification. The same approach is not directly possible for quantum communication due to a quantum-mechanical effect known as the no-cloning theorem [23, 24]. A potential solution is to employ *quantum repeaters* to split the distance separating the two nodes that want to establish entanglement into smaller segments [24–26]. Entanglement is in this scenario first established along the smaller segments, and then connected at the repeaters through entanglement swapping. One way in which quantum networks can be conceptualized is then as consisting of end nodes that wish to establish entanglement that they can consume for different applications, and of quantum repeaters which are used to provide the end nodes with entanglement [19].

Since they were first proposed, much work has been done on quantum repeaters, both regarding the physical systems used to implement them, and on different architectural principles (see e.g., [24]). There have been demonstrations of quantum repeaters, as well as of small-scale quantum-repeater networks [17, 27, 28]. Nevertheless, they are still far from achieving the performance necessary to enable large-scale quantum networks and, eventually, the ultimate goal of a worldwide quantum internet.

In this dissertation we aim to contribute to the realization of quantum-repeater networks based on optical fiber. We investigate what are the requirements on such networks and identify how they might be realized. We further look into how these requirements depend on the repeaters' working principles and on constraints imposed by existing fiber networks. We list below the specific research questions that we aim to answer in this dissertation.

What are the minimal requirements on quantum-repeater hardware?

Current quantum-repeater hardware is not good enough to enable most quantum-networking applications, making clear that experimental progress is still required. However, the questions of (i) how much progress exactly is needed and (ii) along which directions do not have clear answers. By (ii) we mean that it is not clear whether improving, for example, the quality of the repeaters' memories or their emission efficiency would have the greatest impact on the achievable performance. An answer to this question could result in a more efficient allocation of effort and resources, contributing to a swifter realization of quantum-repeater networks. Our first contribution to answering this research question is made in Chapter 5. In it, we propose a method for framing the question as an optimization problem, as well as an approach for solving it using genetic algorithms and simulations of quantum-repeater networks. In Chapters 7, 6 and 8 we apply this method to different scenarios. In Chapter 6 we address the question in-depth. We consider two particular types of quantum hardware, a single-repeater setup using real-world fiber infrastructure and determine requirements for executing a particular application, namely a simple form of blind quantum computing. Finally, in Chapter 7 we consider relatively simple hardware models and investigate hardware requirements for two applications, quantum key distribution and blind quantum computing, while also considering the effect of different placements of repeaters. Finally, in Chapter 8 we answer the question in combination with an optimization over different protocols that can be employed by the repeater nodes, with the goal of investigating how hardware requirements can be minimized through good protocol choices.

How can we best make use of the hardware that we do have?

As already discussed, current quantum-repeater hardware is imperfect. It is certainly true that improving the hardware will, in general, translate to improvements in performance. However, for a given hardware quality, choices can be made regarding protocol usage and repeater placement in order to extract the best performance possible. Deciding which protocols are appropriate is not always trivial, especially because the best choice is hardware quality-, network configuration- and application-dependent. The same holds for how many repeaters to place and where to place them.

In Chapter 6 we investigate different choices of protocols for remote entanglement generation, highlighting different possible paths to functional quantum repeaters. In Chapter 7 we investigate how (i) employing different numbers of repeaters and (ii) optimizing over their placement affects hardware demands. In Chapter 8 we optimize over hardware and protocol parameters in parallel, finding that a good choice of protocols can result in significantly lower hardware requirements. Finally, in Chapter 9 we explore how two nodes with imperfect memories can best share entanglement of a given quality by choosing when to attempt entanglement generation, perform purification or discard old entanglement.

How much do fiber-network-imposed constraints affect the performance of quantum-repeater networks, and how can we best minimize their impact?

Reusing preexistent fiber infrastructure can significantly reduce the cost of deploying quantum-repeater networks in the real world [29]. However, the fact that nodes in this infrastructure are unevenly spaced and loss is not uniform can cause a loss in performance [30]. In this dissertation, we have made efforts to quantify how large of an impact these effects have and how they can best be minimized. In Chapter 4 we determine how

many different ways there are of placing quantum repeaters and heralding stations given a specific number of locations where equipment being placed, which is a prerequisite for investigating the deployment of quantum-repeater chains in existing fiber paths. In Chapters 5, 6 and 7 we investigate hardware requirements under the assumption of constrained placement of network nodes. In Chapters 6 and 7 we optimize over such constraints, and in Chapter 6 we explicitly compare how much larger the requirements are on a real-life fiber grid when compared to an idealized scenario where all nodes are equally spaced.

References

- [1] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, Rev. Mod. Phys. **81**, 865 (2009).
- [2] J. P. Dowling and G. J. Milburn, *Quantum technology: the second quantum revolution*, Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **361**, 1655 (2003).
- [3] AnandTech, *The apple 2022 fall iphone event live blog*, (2022).
- [4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, tenth ed. (Cambridge University Press, USA, 2011).
- [5] R. P. Feynman, *Simulating physics with computers*, Int J Theor Phys **21**, 467 (1982).
- [6] P. W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM review **41**, 303 (1999).
- [7] A. W. Harrow, A. Hassidim, and S. Lloyd, *Quantum algorithm for linear systems of equations*, Physical review letters **103**, 150502 (2009).
- [8] C. H. Bennett, G. Brassard, and N. D. Mermin, *Quantum cryptography without Bell's theorem*, Phys. Rev. Lett. **68**, 557 (1992).
- [9] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Theoretical Computer Science Theoretical Aspects of Quantum Cryptography – Celebrating 30 Years of BB84, **560**, 7 (2014).
- [10] A. K. Ekert, *Quantum cryptography based on Bell's theorem*, Phys. Rev. Lett. **67**, 661 (1991).
- [11] M. Hillery, V. Bužek, and A. Berthiaume, *Quantum secret sharing*, Physical Review A **59**, 1829 (1999).
- [12] D. Gottesman, T. Jennewein, and S. Croke, *Longer-Baseline Telescopes Using Quantum Repeaters*, Phys. Rev. Lett. **109**, 070503 (2012).
- [13] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, *A quantum network of clocks*, Nature Physics **10**, 582 (2014).

- [14] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal blind quantum computation*, 2009 50th Annual IEEE Symposium on Foundations of Computer Science, 517 (2009), 0807.4154.
- [15] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, *Verifying BQP Computations on Noisy Devices with Minimal Overhead*, PRX Quantum 2, 040302 (2021).
- [16] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, *et al.*, *Quantum supremacy using a programmable superconducting processor*, Nature 574, 505 (2019).
- [17] M. Pompili, S. L. Hermans, S. Baier, H. K. Beukers, P. C. Humphreys, R. N. Schouten, R. F. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, *et al.*, *Realization of a multinode quantum network of remote solid-state qubits*, Science 372, 259 (2021).
- [18] J. Chow, O. Dial, and J. Gambetta, *Ibm quantum breaks the 100-qubit processor barrier*, IBM Research Blog (2021).
- [19] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead*, Science 362, eaam9288 (2018).
- [20] W. K. Wootters, *Quantum entanglement as a quantifiable resource*, Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 356, 1717 (1998).
- [21] S. Hermans, M. Pompili, L. dos Santos Martins, A. Rodriguez-Pardo Montblanch, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study*, New Journal of Physics (2023).
- [22] F. Mitschke, *Fiber Optics: Physics and Technology* (Heidelberg ; New York, 2010).
- [23] W. K. Wootters and W. H. Zurek, *The no-cloning theorem*, Physics Today 62, 76 (2009).
- [24] W. J. Munro, K. Azuma, K. Tamaki, and K. Nemoto, *Inside Quantum Repeaters*, IEEE Journal of Selected Topics in Quantum Electronics 21, 78 (2015).
- [25] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum repeaters: the role of imperfect local operations in quantum communication*, Physical Review Letters 81, 5932 (1998).
- [26] N. Sangouard, C. Simon, H. De Riedmatten, and N. Gisin, *Quantum repeaters based on atomic ensembles and linear optics*, Reviews of Modern Physics 83, 33 (2011).
- [27] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, *et al.*, *Experimental demonstration of memory-enhanced quantum communication*, Nature 580, 60 (2020).
- [28] S. Langenfeld, P. Thomas, O. Morin, and G. Rempe, *Quantum repeater node demonstrating unconditionally secure key distribution*, Physical review letters 126, 230506 (2021).

- [29] J. Rabbie, K. Chakraborty, G. Avis, and S. Wehner, *Designing quantum networks using preexisting infrastructure*, npj Quantum Information **8**, 5 (2022).
- [30] G. Avis, R. Knegjens, A. S. Sørensen, and S. Wehner, *Asymmetric node placement in fiber-based quantum networks*, arXiv preprint arXiv:2305.09635 (2023).

3

Preliminaries

In this chapter we introduce some concepts that are key to understanding the rest of the dissertation. We focus on quantum repeaters, particularly their operational principles and metrics for evaluating their performance.

3.1 Operational principles of quantum repeaters

As discussed in Chapter 2, quantum repeaters can be employed to generate entanglement between distant quantum-network nodes. The particular types of repeaters we study do this by employing a combination of three basic actions: heralded entanglement generation, entanglement swapping and entanglement purification [1–7]. Heralded entanglement generation is the process through which neighbouring nodes establish entanglement with one another, entanglement swapping is the process through which existing entangled links are fused together to create entangled links spanning a longer distance and entanglement purification protocols consume multiple lower-quality entangled links to probabilistically generate fewer higher-quality ones. Entanglement purification is not strictly necessary for long-distance entanglement generation (in fact, we do not consider it in Chapters 5, 6 and 7 of this dissertation), but it can be useful in combating the quality decay associated with the use of multiple repeaters. We will start by describing in more detail what an entangled link looks like.

3.1.1 Entanglement: Bell states

The state $|\psi\rangle$ of a qubit, the basic unit of quantum information, can be represented as

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad (3.1)$$

in which $|0\rangle$ and $|1\rangle$ are vectors and α and β are complex numbers that are normalized such that $|\alpha|^2 + |\beta|^2 = 1$. If both α and β are non-zero, $|\psi\rangle$ is said to be in a superposition. Two-qubit states can be superpositions of the two-qubit base states $|0\rangle \otimes |0\rangle$, $|0\rangle \otimes |1\rangle$, $|1\rangle \otimes |0\rangle$ and $|1\rangle \otimes |1\rangle$. \otimes is often omitted for the sake of brevity, i.e., $|0\rangle \otimes |0\rangle \equiv |00\rangle$. We can now define *Bell states* [8, 9]:

$$|\phi_{ij}\rangle = (X^i Z^j \otimes \mathbb{1}_2) \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle), \quad (3.2)$$

with X and Z being the Pauli X and Z operators, defined respectively as $|0\rangle\langle 1| + |1\rangle\langle 0|$ and $|0\rangle\langle 0| - |1\rangle\langle 1|$, $\mathbb{1}_2$ the two-dimensional identity matrix, $\mathbb{1}_2 = |0\rangle\langle 0| + |1\rangle\langle 1|$, and i, j being either 0 or 1. The four Bell states define a basis for two-qubit states. They are also entangled states, in the sense that it is impossible to rewrite them as a product of two single-qubit states [10], i.e.,

$$|\phi_{ij}\rangle \neq |\psi_A\rangle \otimes |\psi_B\rangle. \quad (3.3)$$

Furthermore, they are *maximally* entangled, which means that if the one of the qubits is for some reason lost, we will have no information about the state of the other. Every two-qubit maximally entangled state is equivalent to a Bell state in the sense that it can be transformed into a Bell state via single-qubit unitary operations [10]. This is relevant because in the context of quantum communication one often thinks of local operations, such as single-qubit operations on one half of a Bell state, as being "easy" and remote operations, such as entanglement generation or distributed gates, as being "hard". The goal of quantum repeaters is to generate Bell states between distant quantum-network nodes and to do so as fast as possible. One way of doing so begins with heralded entanglement generation.

3

3.1.2 Heralded entanglement generation

The protocols for entanglement generation considered in this dissertation are heralded [1, 2]. Such protocols can broadly be described as follows: two remote nodes perform attempts at entanglement generation. After each attempt they are notified of whether or not they have been successful. This is known as "heralding". If they were not successful, they keep trying. Otherwise, the protocol concludes with the nodes sharing a possibly-imperfect entangled state. More specifically, we consider single- and double-click protocols. The single-click protocol was proposed in [3] and experimentally demonstrated in, among others, [11–13]; for an in-depth discussion see [14]. The double-click protocol was proposed in [4] and experimentally demonstrated in, among others, [15–18]. Other protocols exist, see e.g., [1, 7]. We will here give a high-level description of these protocols. For a more detailed and technical exposition, see Chapter 6. Both protocols start with the two nodes that wish to be entangled generating light-matter entanglement. The resulting photons, which are entangled with each of the nodes, are sent to a heralding station placed somewhere between the two nodes. This station contains both beamsplitters and photon detectors. There, the photons are interfered and measured, which projects the entanglement onto the matter qubits. The nodes are then informed of the outcome of measurement. There are multiple reasons due to which this process can fail. For example, it might be that the detector does not detect a photon even though it has arrived, or that a photon emitted by one of the nodes is not properly captured into the fiber. Even under the assumption of perfect detectors and linear optics, the success probability cannot exceed 50% [19] (although this can be increased if ancillary photons are used [20]). The main reason why the process can fail is the same reason why quantum repeaters are needed in the first place: photons are absorbed when travelling in optical fiber.

3.1.3 Entanglement swapping

Entanglement swapping is the second and last of the strictly-necessary building blocks used by the quantum repeaters investigated in this dissertation. In order to define it, we

start by defining a Bell state measurement (BSM) as a measurement in which a two-qubit state is projected onto the Bell basis (for a primer on measurement in quantum mechanics see [9]). Let qubits 1 and 2 be in the state $|\psi_{ij}\rangle_{1,2}$ and qubits 3 and 4 in state $|\psi_{kl}\rangle_{3,4}$. Assume further that a BSM is performed on qubits 2 and 3 and that the outcome of this BSM is $|\psi_{mn}\rangle$. The state that qubits 1 and 4 share after the measurement is:

$$\frac{(|\psi_{mn}\rangle\langle\psi_{mn}|)_{2,3} |\psi_{ij}\rangle_{1,2} |\psi_{kl}\rangle_{3,4}}{\left|(|\psi_{mn}\rangle\langle\psi_{mn}|)_{2,3} |\psi_{ij}\rangle_{1,2} |\psi_{kl}\rangle_{3,4}\right|} = |\psi_{i\oplus k, j\oplus l}\rangle_{1,4} |\psi_{mn}\rangle_{2,3}. \quad (3.4)$$

Note that qubits 2 and 3 have been projected to $|\psi_{mn}\rangle$ in accordance with the BSM outcome. Qubits 1 and 4, which initially shared no entanglement, and indeed had potentially never interacted or even been spatially close to each other, are now entangled. This is an entanglement swap, which is nothing more than a BSM applied to two qubits which are part of different entangled pairs. In the context of the quantum repeaters studied in this dissertation, an entanglement swap is performed in two contexts. First, between photons which are entangled with matter qubits at measurement stations. Second, between matter qubits at repeater nodes, which have previously established entanglement with two other remote nodes. Quantum repeaters can then establish long-distance entanglement by first generating entangled links over short distances in a heralded fashion, and then connecting these links into longer ones through entanglement swapping. There have been multiple experimental demonstrations of entanglement swapping using various physical systems [11, 13, 21–24].

3.1.4 Entanglement purification

Entanglement purification is a process through which m entangled links are probabilistically transformed into k links of higher quality (the concept of link quality will be discussed in more detail further on), with $k < m$ [25, 26]. Even if the physical systems used to implement quantum repeater nodes are perfect, the quality of the links decays exponentially with the number of entanglement swaps performed [5, 27]. The use of purification can then be used to combat this. In this work we consider 2-to-1 protocols, i.e., protocols in which 2 entangled links are consumed to, with some non-zero probability, output 1 link of higher quality. The two protocols that we consider are, in particular, DEJMPS [25] and EPL [28, 29]. They are described in detail in Chapter 8.

3.1.5 Scaling of quantum repeaters

Quantum repeaters can be used to combat photon loss in fiber. The transmission efficiency $\eta(L)$, i.e., the probability that a photon is successfully transmitted through a segment of fiber of length L , is given by [30]

$$\eta(L) = 10^{-\frac{\alpha_{\text{att}}}{10} L}, \quad (3.5)$$

where α_{att} is a fiber-specific attenuation coefficient (typically considered to be around 0.2 dB km⁻¹ in the quantum-repeater literature [2, 7]). Besides transmission losses, other imperfections can cause photon loss. These include, for example, the photon not being directed into the fiber successfully. Let us assume, for the sake of simplicity, that these

can be captured into a single parameter p such that the probability p_{surv} of an emitted photon successfully travelling L km is given by

$$p_{\text{surv}}(L) = p\eta(L). \quad (3.6)$$

Let us now imagine that we place N equidistant repeaters between the two nodes that wish to generate entanglement, such that the total length to be covered L is divided into $N + 1$ segments of length $L_0 = L/(N + 1)$ each. All nodes will attempt to perform entanglement generation with both their neighbors simultaneously. If all of them are successful, the repeater nodes can perform entanglement swaps and end-to-end entanglement is established. Otherwise, if there are any failures, all links must be regenerated. Therefore, the probability of successfully establishing end-to-end entanglement p_{e2e} is given by

$$p_{\text{e2e}} = p_{\text{surv}}(L_0)^{N+1} = (p\eta(L_0))^{N+1} = p^{N+1}\eta(L). \quad (3.7)$$

Adding repeaters actually made things worse by keeping the scaling of p_{e2e} with distance the same, but making the constant factor smaller (p^{N+1} vs p). In order for repeaters to be useful an extra ingredient is required: quantum memories.

3.1.6 Quantum memories

A quantum memory is a system that can preserve quantum states over time. Endowing quantum repeaters with quantum memories can boost the entanglement generation rate [2, 5, 7]. To see this, let us again imagine that we place N equidistant repeaters between the two nodes that wish to generate entanglement, such that the total length to be covered L is divided into $N + 1$ segments of length $L_0 = L/(N + 1)$ each. We previously assumed that all links had to be successfully generated simultaneously. But if the repeaters are endowed with quantum memories, this is not true. Links that are successfully generated can be stored in memory while the remaining ones try again. This process can be repeated until all links have been generated, at which point entanglement swaps can be performed to establish end-to-end entanglement. State-of-the-art quantum memories are imperfect [11, 23, 31–33]. Storage implies errors in the quantum state, and these typically become more severe the longer a state is kept in memory. In fact, after enough time has elapsed the state might be of too low quality to be useful at all, in which case discarding it is a good choice. The amount of time after which a state is discarded is known as the cut-off time [34–40]. It implies a trade-off between rate of entanglement generation and quality of the generated entanglement: not employing cut-offs at all results in faster entanglement, whereas discarding often ensures high-quality states. Determining the best cut-off time for a particular situation is in general a non-trivial problem [36]. In Chapters 6 and 7 we optimize over the cut-off time so as to determine minimal hardware requirements to satisfy given network performance metrics. We note also that even though we gave as a simple example the case in which all links are generated and then all swaps are performed, this is not in general the optimal swapping strategy [34].

3.1.7 Multiplexing

The probability of successful entanglement generation between neighbouring nodes decays exponentially with the fiber distance between the nodes, as seen in Equation 3.5.

One possible way of combating this is to perform multiple entanglement generation attempts in parallel, a process known as *multiplexing*. This can be achieved in multiple different ways. For example, photons can be sent at different frequencies [41–43] (spectral multiplexing) or with some time delay between them [44–47] (temporal multiplexing). It is likely that achieving satisfactory levels of performance (see Section 3.3) will require employing a combination of these methods. If a large enough number of multiplexing modes is employed, the probability of successfully generating entanglement can be made arbitrarily close to 1:

$$P(N) = 1 - (1 - P_0)^N, \quad (3.8)$$

with N being the number of multiplexing modes and P_0 the probability of a single mode succeeding. Employing multiplexing somewhat relaxes the requirements on how long the quantum memories must be. However, the quantum memories must still be good enough to store states while waiting for the heralding signal. Furthermore, employing multiplexing requires multimode quantum memories which are challenging to implement [43, 44, 46, 48–52].

3.2 Other types of repeaters

The repeater architecture we have discussed thus far, and which will be considered in this dissertation, is often referred to as the first generation of quantum repeaters [5, 7]. Such repeaters are characterized by the employment of purification to combat errors in the entanglement generation and entanglement swapping processes. Another tactic that can be used for combatting such errors is that of quantum-error-correction [53–55]. In this case, entanglement swaps (and hence end-to-end entanglement) are not performed directly between physical qubits, but instead between logical qubits which result from an encoding of physical qubits [56–59]. Assuming the quality of operations performed at the quantum repeaters is high enough, the process of encoding can result in higher resilience against errors. Repeater architectures that employ quantum memories in combination with error-corrected entanglement swaps are commonly known as second-generation quantum repeaters [2, 7, 60].

Error correction can also be employed to combat loss in photon transmission. This obviates the need for two-way communication inherent to the heralding process (and hence allowing for, in theory, higher entangling rates). For this reason, the repeaters implementing this idea are known as one-way quantum repeaters, or alternatively third-generation quantum repeaters [2, 7, 60–64]. A downside to this type of repeaters is that due to the no-cloning theorem one cannot protect against losses exceeding 50%, meaning that such repeaters would have to be placed very closely together (at intervals of roughly 15 km for standard optical fiber) [7]. This would likely render their deployment very costly.

Techniques such as purification and error correction may help mitigate errors, but they do so at the cost of more stringent requirements on quantum-repeater hardware. For example, one needs to be able to generate and store multiple entangled pairs between two nodes as well as perform more involved operations than simply an entanglement swap. This might make them less suitable for near-term quantum networks. Some discussion of this topic can be found in Chapter 8 as well as in [65]. Further information about different quantum-repeater architectures can be found in, e.g., [2, 7].

3.3 Quantifying repeater performance

We need metrics for determining how well repeaters perform. Some commonly used metrics are derived from information theory, and relate to fundamental bounds on how much information can be transmitted per channel use [66, 67]. This gives rise questions of the type "when does a repeater repeat", i.e., when does employing repeaters allow for transmitting strictly more information than would be possible without one (see, for example, [37, 40]). In this dissertation we focus instead on the rate of entanglement generation and the fidelity of the entangled pairs generated. These can in simple terms be thought of as how fast entanglement is generated and how good it is. The analysis of the fidelity is required as errors can affect the quantum state that is provided by the quantum repeaters. This can better be understood through the use of the density matrix formalism [8, 9]. When a system is in state $|\psi_i\rangle$ with probability p_i its density matrix ρ can be written as

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|, \quad (3.9)$$

with $\sum p_i = 1$ to ensure normalization. In case there are multiple non-zero p_i s the system is said to be in a mixed state, whereas if only one of the p_i s is non-zero it is said to be in a pure state. In the latter case the density matrix formalism is equivalent to the ket formalism we introduced earlier.

We are now ready to introduce the concept of fidelity. If our aim was to distribute a state $|\phi\rangle$ but we instead distributed ρ , the fidelity of ρ to $|\phi\rangle$ is given by [68]:

$$F = \langle\phi|\rho|\phi\rangle. \quad (3.10)$$

The fidelity is then a measure for how different from the target state the state that was actually delivered is. We can now also more concretely define the concept of rate. The entanglement generation rate (or entangling rate) R is the average number of entangled pairs that can be distributed per unit time. Due to its simplicity of measurement it is often reported in quantum-networking experiments [11, 12, 17, 18, 69–71]. If the average waiting time for entanglement distribution is T , the entangling rate is given by:

$$R = \frac{1}{T}. \quad (3.11)$$

3.3.1 Application-derived performance metrics

Rate and fidelity are useful metrics for evaluating quantum-repeater performance. However, they are somewhat incomplete. For one, it is inconvenient that they are two separate numbers, as it renders comparisons difficult. For example, it is not clear whether it is better to distribute entangled pairs of fidelity 0.8 at a rate of 10 Hz, or entangled pairs of fidelity 0.95 at a rate of 1 Hz. The answer to this question is context-dependent and more concretely application-dependent. Entanglement is a resource, and quantum repeaters provide a service in distributing this resource. Entanglement will then be consumed by nodes to perform applications, so it is only natural that performance evaluation done is from an application-centric perspective.

Secret-key rate

One of the oldest and best-known quantum-network application is quantum key distribution (QKD) [72–77]. This is an application through which two users that share entanglement can generate a shared mathematically-secret key that can then be employed for secure communication. They do this by performing measurements on their entangled state. If they measure in the same basis and the entangled state they share is perfect ($F = 1$ to a Bell state) they either expect perfect correlation or perfect anticorrelation between their measurement outcomes (depending on which Bell state they share and in which basis they measure). This then directly produces one bit of secret key. However, if the entangled state is imperfect, some measurements outcomes will not follow the expected (anti) correlation pattern. This is measured by the Quantum Bit Error Rate (QBER), the fraction of measurement outcomes which do not follow the expected pattern. In this situation, a raw key is generated from the measurement outcomes. A secret key can still be distilled in this case using classical error correction, given that the QBER is not larger than a protocol-specific limit. We thus define the secret-key fraction (SKF) as the ratio between the lengths of the secret key and the raw key in the asymptotic limit, i.e., when the length of the raw key goes to infinity. The SKF is a decreasing function of the QBER and is thus a measure of the quality of the entangled state used to generate the key. It can be thought of as the fraction of key that can be extracted from one entangled state. The quality of service provided by a QKD system can then be measured through the secret-key rate (SKR), the amount of secret-key bits generated per unit time:

$$SKR = SKF \cdot R. \quad (3.12)$$

There are multiple QKD protocols (see, e.g., [76]), but we have in this dissertation focused on the entanglement-based version of BB84 [73, 74]. For this protocol, the SKF can be computed as follows [78–80]:

$$SKF = \max(0, 1 - 2h(\text{QBER})), \quad (3.13)$$

where $h(x) = -x \ln x - (1 - x) \ln(1 - x)$ is the binary-entropy function and we have for simplicity assumed that the QBER is the same in both measurement bases. The BB84 SKR is a widely-used performance metric [22, 36, 37, 40, 43, 81–84].

Blind quantum computing

The SKR is an application-derived performance metric, which as discussed gives it advantages over the fidelity and rate of entanglement generation. However, it pertains to a particular application, and it might be that other applications impose different demands. Furthermore, QKD is a single-qubit application, in the sense that the two parties involved never need to hold multiple entangled pairs in memory. This renders it fundamentally different from many other quantum-network applications, which might require multiple entangled qubits to be simultaneously held in memory. Given that near-term quantum memories are noisy, it is likely that this will have a significant effect on the quantum-network requirements. For example, it might be that a quantum network whose nodes have fairly short coherence times will still enable users to perform QKD well, but it will likely struggle to enable applications requiring multiple live qubits. One example of such

an application is blind quantum computing (for more information on blind quantum computing and how it can be used to derive a performance metric, see Chapter 7). This is an application that allows a client to execute a quantum computation on a server while ensuring that the server does not learn what computation is executed. Such an application might be useful in a scenario where the client has access to significantly less quantum processing power than the server. For concreteness and simplicity, we focus on a minimal version of this application where a two-qubit program is executed at the server. In particular, we consider an instance of the protocol introduced in [85–91]. In this protocol, the client randomly alternates between computation rounds and test rounds. In computation rounds, the client is interested in learning the output of the computation. In test rounds, the client knows the expected outcome ahead of time, and compares it with the obtained outcome. Disparities between the two can be due to either noise or a dishonest server. The purpose of test rounds is to ensure the server’s honesty, but this is foiled if the entangled states shared between client and server are too noisy. The test-round success probability under the assumption of an honest server is then an important metric for how well the protocol can be executed, and hence we employ it as a performance metric in Chapter 7.

References

- [1] T. E. Northup and R. Blatt, *Quantum information transfer using photons*, Nature Photon **8**, 356 (2014).
- [2] K. Azuma, S. E. Economou, D. Elkouss, P. Hilaire, L. Jiang, H.-K. Lo, and I. Tzitrin, *Quantum repeaters: From quantum networks to the quantum internet*, (2022), arXiv:2212.10820 .
- [3] C. Cabrillo, J. I. Cirac, P. García-Fernández, and P. Zoller, *Creation of entangled states of distant atoms by interference*, Phys. Rev. A **59**, 1025 (1999).
- [4] S. D. Barrett and P. Kok, *Efficient high-fidelity quantum computation using matter qubits and linear optics*, Phys. Rev. A **71**, 060310 (2005).
- [5] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum repeaters: the role of imperfect local operations in quantum communication*, Physical Review Letters **81**, 5932 (1998).
- [6] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels*, Phys. Rev. Lett. **70**, 1895 (1993).
- [7] W. J. Munro, K. Azuma, K. Tamaki, and K. Nemoto, *Inside Quantum Repeaters*, IEEE Journal of Selected Topics in Quantum Electronics **21**, 78 (2015).
- [8] S. Khatri and M. M. Wilde, *Principles of Quantum Communication Theory: A Modern Approach*, (2020), arXiv:2011.04672 .
- [9] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, tenth ed. (Cambridge University Press, USA, 2011).
- [10] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, Rev. Mod. Phys. **81**, 865 (2009).

- [11] M. Pompili, S. L. Hermans, S. Baier, H. K. Beukers, P. C. Humphreys, R. N. Schouten, R. F. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, *et al.*, *Realization of a multinode quantum network of remote solid-state qubits*, *Science* **372**, 259 (2021).
- [12] P. C. Humphreys, N. Kalb, J. P. Morits, R. N. Schouten, R. F. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network*, *Nature* **558**, 268 (2018).
- [13] M. Riebe, T. Monz, K. Kim, A. S. Villar, P. Schindler, M. Chwalla, M. Hennrich, and R. Blatt, *Deterministic entanglement swapping with an ion-trap quantum computer*, *Nature Phys* **4**, 839 (2008).
- [14] S. Hermans, M. Pompili, L. dos Santos Martins, A. Rodriguez-Pardo Montblanch, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study*, *New Journal of Physics* (2023).
- [15] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded entanglement between solid-state qubits separated by three metres*, *Nature* **497**, 86 (2013).
- [16] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenber, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres*, *Nature* **526**, 682 (2015).
- [17] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network*, *Phys. Rev. Lett.* **124**, 110501 (2020).
- [18] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, *Entanglement of trapped-ion qubits separated by 230 meters*, *Phys. Rev. Lett.* **130**, 050803 (2023).
- [19] J. Calsamiglia and N. Lütkenhaus, *Maximum efficiency of a linear-optical Bell-state analyzer*, *Appl Phys B* **72**, 67 (2001).
- [20] W. P. Grice, *Arbitrarily complete Bell-state measurement using only linear optical elements*, *Phys. Rev. A* **84**, 042331 (2011).
- [21] F. B. Basset, M. B. Rota, C. Schimpf, D. Tedeschi, K. D. Zeuner, S. C. Da Silva, M. Reindl, V. Zwiller, K. D. Jöns, A. Rastelli, *et al.*, *Entanglement swapping with photons generated on demand by a quantum dot*, *Physical Review Letters* **123**, 160501 (2019).
- [22] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, *et al.*, *Experimental demonstration of memory-enhanced quantum communication*, *Nature* **580**, 60 (2020).

- [23] V. Krutyanskiy, M. Canteri, M. Meraner, J. Bate, V. Krcmarsky, J. Schupp, N. Sangouard, and B. P. Lanyon, *Telecom-wavelength quantum repeater node based on a trapped-ion processor*, *Physical Review Letters* **130**, 213601 (2023).
- [24] J.-W. Pan, D. Bouwmeester, H. Weinfurter, and A. Zeilinger, *Experimental entanglement swapping: entangling photons that never interacted*, *Physical review letters* **80**, 3891 (1998).
- [25] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, *Quantum privacy amplification and the security of quantum cryptography over noisy channels*, *Physical review letters* **77**, 2818 (1996).
- [26] W. Dür and H. J. Briegel, *Entanglement purification and quantum error correction*, *Reports on Progress in Physics* **70**, 1381 (2007).
- [27] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller, *Quantum repeaters based on entanglement purification*, *Physical Review A* **59**, 169 (1999).
- [28] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356**, 928 (2017).
- [29] F. Rozpędek, T. Schiet, D. Elkouss, A. C. Doherty, S. Wehner, *et al.*, *Optimizing practical entanglement distillation*, *Physical Review A* **97**, 062333 (2018).
- [30] F. Mitschke, *Fiber Optics: Physics and Technology* (Heidelberg ; New York, 2010).
- [31] C. Bradley, S. de Bone, P. Möller, S. Baier, M. Degen, S. Loenen, H. Bartling, M. Markham, D. Twitchen, R. Hanson, *et al.*, *Robust quantum-network memory based on spin qubits in isotopically engineered diamond*, *npj Quantum Information* **8**, 122 (2022).
- [32] P. Drmota, D. Main, D. Nadlinger, B. Nichol, M. Weber, E. Ainley, A. Agrawal, R. Srinivas, G. Araneda, C. Ballance, *et al.*, *Robust quantum memory in a trapped-ion quantum network node*, *Physical Review Letters* **130**, 090803 (2023).
- [33] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, *Single ion qubit with estimated coherence time exceeding one hour*, *Nature communications* **12**, 233 (2021).
- [34] Á. G. Iñesta, G. Vardoyan, L. Scavuzzo, and S. Wehner, *Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs*, *npj Quantum Information* **9**, 46 (2023).
- [35] S. Khatri, *Policies for elementary links in a quantum network*, *Quantum* **5**, 537 (2021).
- [36] B. Li, T. Coopmans, and D. Elkouss, *Efficient optimization of cutoffs in quantum repeater chains*, *IEEE Transactions on Quantum Engineering* **2**, 1 (2021).

- [37] F. Rozpędek, K. Goodenough, J. Ribeiro, N. Kalb, V. C. Vivoli, A. Reiserer, R. Hanson, S. Wehner, and D. Elkouss, *Parameter regimes for a single sequential quantum repeater*, *Quantum Science and Technology* **3**, 034002 (2018).
- [38] S. Santra, L. Jiang, and V. S. Malinovsky, *Quantum repeater architecture with hierarchically optimized memory buffer times*, *Quantum Science and Technology* **4**, 025010 (2019).
- [39] W. Kozłowski, A. Dahlberg, and S. Wehner, *Designing a quantum network protocol*, in *Proceedings of the 16th international conference on emerging networking experiments and technologies* (2020) pp. 1–16.
- [40] F. Rozpędek, R. Yehia, K. Goodenough, M. Ruf, P. C. Humphreys, R. Hanson, S. Wehner, and D. Elkouss, *Near-term quantum-repeater experiments with nitrogen-vacancy centers: Overcoming the limitations of direct transmission*, *Physical Review A* **99**, 052330 (2019).
- [41] N. Sinclair, E. Saglamyurek, H. Mallahzadeh, J. A. Slater, M. George, R. Ricken, M. P. Hedges, D. Oblak, C. Simon, W. Sohler, *et al.*, *Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control*, *Physical review letters* **113**, 053603 (2014).
- [42] T. Chakraborty, H. van Brug, A. Das, O. Pietx-Casas, P.-C. Wang, G. C. d. Amaral, A. L. Tchegotareva, and W. Tittel, *Frequency multiplexed photon pairs and detection for quantum repeaters*, arXiv preprint arXiv:2205.10028 (2022).
- [43] S. Guha, H. Krovi, C. A. Fuchs, Z. Dutton, J. A. Slater, C. Simon, and W. Tittel, *Rate-loss analysis of an efficient quantum repeater architecture*, *Physical Review A* **92**, 022357 (2015).
- [44] C. Simon, H. De Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, *Quantum repeaters with photon pair sources and multimode memories*, *Physical review letters* **98**, 190503 (2007).
- [45] S. B. van Dam, P. C. Humphreys, F. Rozpędek, S. Wehner, and R. Hanson, *Multiplexed entanglement generation over quantum networks using multi-qubit nodes*, *Quantum Science and Technology* **2**, 034002 (2017).
- [46] D. Lago-Rivera, S. Grandi, J. V. Rakonjac, A. Seri, and H. de Riedmatten, *Telecom-heralded entanglement between multimode solid-state quantum memories*, *Nature* **594**, 37 (2021).
- [47] M. Businger, L. Nicolas, T. S. Mejia, A. Ferrier, P. Goldner, and M. Afzelius, *Non-classical correlations over 1250 modes between telecom photons and 979-nm photons stored in 171yb3+: Y2sio5*, *Nature communications* **13**, 6438 (2022).
- [48] M. Afzelius, C. Simon, H. De Riedmatten, and N. Gisin, *Multimode quantum memory based on atomic frequency combs*, *Physical Review A* **79**, 052329 (2009).

- [49] M.-X. Dong, W.-H. Zhang, L. Zeng, Y.-H. Ye, D.-C. Li, G.-C. Guo, D.-S. Ding, and B.-S. Shi, *Highly efficient storage of 25-dimensional photonic qudit in a cold-atom-based quantum memory*, arXiv preprint arXiv:2301.00999 (2023).
- [50] A. Ortu, J. V. Rakonjac, A. Holzäpfel, A. Seri, S. Grandi, M. Mazzera, H. de Riedmatten, and M. Afzelius, *Multimode capacity of atomic-frequency comb quantum memories*, *Quantum Science and Technology* **7**, 035024 (2022).
- [51] J. V. Rakonjac, D. Lago-Rivera, A. Seri, M. Mazzera, S. Grandi, and H. de Riedmatten, *Entanglement between a telecom photon and an on-demand multimode solid-state quantum memory*, *Physical Review Letters* **127**, 210502 (2021).
- [52] A. Seri, D. Lago-Rivera, A. Lenhard, G. Corrielli, R. Osellame, M. Mazzera, and H. de Riedmatten, *Quantum storage of frequency-multiplexed heralded single photons*, *Physical review letters* **123**, 080502 (2019).
- [53] J. Roffe, *Quantum Error Correction: An Introductory Guide*, *Contemporary Physics* **60**, 226 (2019).
- [54] B. M. Terhal, *Quantum Error Correction for Quantum Memories*, *Rev. Mod. Phys.* **87**, 307 (2015).
- [55] E. Knill and R. Laflamme, *Theory of quantum error-correcting codes*, *Phys. Rev. A* **55**, 900 (1997).
- [56] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, *Quantum repeater with encoding*, *Phys. Rev. A* **79**, 032325 (2009).
- [57] Y. Jing and M. Razavi, *Quantum Repeaters with Encoding on Nitrogen-Vacancy-Center Platforms*, *Phys. Rev. Appl.* **18**, 024041 (2022).
- [58] W. J. Munro, K. A. Harrison, A. M. Stephens, S. J. Devitt, and K. Nemoto, *From quantum multiplexing to high-performance quantum networking*, *Nature Photon* **4**, 792 (2010).
- [59] A. M. Stephens, J. Huang, K. Nemoto, and W. J. Munro, *Hybrid-system approach to fault-tolerant quantum communication*, *Phys. Rev. A* **87**, 052333 (2013).
- [60] S. Muralidharan, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, *Ultrafast and Fault-Tolerant Quantum Communication across Long Distances*, *Phys. Rev. Lett.* **112**, 250501 (2014).
- [61] J. Borregaard, H. Pichler, T. Schröder, M. D. Lukin, P. Lodahl, and A. S. Sørensen, *One-Way Quantum Repeater Based on Near-Deterministic Photon-Emitter Interfaces*, *Phys. Rev. X* **10**, 021071 (2020).
- [62] A. G. Fowler, D. S. Wang, C. D. Hill, T. D. Ladd, R. Van Meter, and L. C. L. Hollenberg, *Surface code quantum communication*, *Phys. Rev. Lett.* **104**, 180503 (2010).
- [63] M. Varnava, D. E. Browne, and T. Rudolph, *Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction*, *Phys. Rev. Lett.* **97**, 120501 (2006).

- [64] K. Azuma, K. Tamaki, and H.-K. Lo, *All-photonic quantum repeaters*, Nature Communications **6**, 1 (2015).
- [65] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpędek, M. Skrzypczyk, L. Wubben, W. de Jong, D. Podareanu, A. Torres-Knoop, D. Elkouss, and S. Wehner, *NetSquid, a NETWORK Simulator for QUantum Information using Discrete events*, Commun Phys **4**, 1 (2021).
- [66] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, *Fundamental limits of repeaterless quantum communications*, Nature communications **8**, 15043 (2017).
- [67] M. Takeoka, S. Guha, and M. M. Wilde, *The squashed entanglement of a quantum channel*, IEEE Transactions on Information Theory **60**, 4987 (2014).
- [68] A. Gilchrist, N. K. Langford, and M. A. Nielsen, *Distance measures to compare real and ideal quantum processes*, Physical Review A **71**, 062310 (2005).
- [69] T. van Leent, M. Bock, R. Garthoff, K. Redeker, W. Zhang, T. Bauer, W. Rosenfeld, C. Becher, and H. Weinfurter, *Long-Distance Distribution of Atom-Photon Entanglement at Telecom Wavelength*, Phys. Rev. Lett. **124**, 010510 (2020).
- [70] J. Yin, Y.-H. Li, S.-K. Liao, M. Yang, Y. Cao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, S.-L. Li, *et al.*, *Entanglement-based secure quantum cryptography over 1,120 kilometres*, Nature **582**, 501 (2020).
- [71] J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld, and H. Weinfurter, *Heralded entanglement between widely separated atoms*, Science **337**, 72 (2012).
- [72] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, *et al.*, *Advances in quantum cryptography*, Advances in optics and photonics **12**, 1012 (2020).
- [73] C. H. Bennett, G. Brassard, and N. D. Mermin, *Quantum cryptography without Bell's theorem*, Phys. Rev. Lett. **68**, 557 (1992).
- [74] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Theoretical Computer Science Theoretical Aspects of Quantum Cryptography – Celebrating 30 Years of BB84, **560**, 7 (2014).
- [75] A. K. Ekert, *Quantum cryptography based on Bell's theorem*, Phys. Rev. Lett. **67**, 661 (1991).
- [76] G. Murta, F. Rozpędek, J. Ribeiro, D. Elkouss, and S. Wehner, *Key rates for quantum key distribution protocols with asymmetric noise*, Physical Review A **101**, 062321 (2020).
- [77] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, *The security of practical quantum key distribution*, Reviews of modern physics **81**, 1301 (2009).

- [78] P. W. Shor and J. Preskill, *Simple proof of security of the bb84 quantum key distribution protocol*, Physical review letters **85**, 441 (2000).
- [79] R. Renner, *Security of quantum key distribution*, International Journal of Quantum Information **6**, 1 (2008).
- [80] B. Kraus, N. Gisin, and R. Renner, *Lower and upper bounds on the secret-key rate for quantum key distribution protocols using one-way classical communication*, Physical review letters **95**, 080501 (2005).
- [81] Y. Jing and M. Razavi, *Quantum repeaters with encoding on nitrogen-vacancy-center platforms*, Physical Review Applied **18**, 024041 (2022).
- [82] L. Kamin, E. Shchukin, F. Schmidt, and P. van Loock, *Exact rate analysis for quantum repeaters with imperfect memories and entanglement swapping as soon as possible*, Physical Review Research **5**, 023086 (2023).
- [83] G. Vardoyan and S. Wehner, *Quantum network utility maximization*, arXiv preprint arXiv:2210.08135 (2022).
- [84] S. Langenfeld, P. Thomas, O. Morin, and G. Rempe, *Quantum repeater node demonstrating unconditionally secure key distribution*, Physical review letters **126**, 230506 (2021).
- [85] A. M. Childs, *Secure assisted quantum computation*, arXiv preprint quant-ph/0111046 (2001).
- [86] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, *Verifying bqp computations on noisy devices with minimal overhead*, PRX Quantum **2**, 040302 (2021).
- [87] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal blind quantum computation*, 2009 50th Annual IEEE Symposium on Foundations of Computer Science , 517 (2009), 0807.4154 .
- [88] T. Morimae and K. Fujii, *Blind quantum computation protocol in which alice only makes measurements*, Physical Review A **87**, 050301 (2013).
- [89] J. F. Fitzsimons and E. Kashefi, *Unconditionally verifiable blind quantum computation*, Physical Review A **96**, 012303 (2017).
- [90] P. Arrighi and L. Salvail, *Blind quantum computation*, International Journal of Quantum Information **4**, 883 (2006).
- [91] T. Kapourniotis, E. Kashefi, D. Leichtle, L. Music, and H. Ollivier, *A framework for verifiable blind quantum computation*, arXiv e-prints , arXiv (2022).

4

Counting quantum-repeater configurations

**Francisco Ferreira da Silva¹, Guus Avis¹,
Kenneth Goodenough¹ and
Stephanie Wehner.**

In this chapter we give formulas for counting in how many different ways quantum repeaters can be distributed over N points. Such formulas might be useful when investigating the deployment of quantum repeaters using existing fiber infrastructure. We consider two distinct cases: (i) the total number of repeaters to be placed is limited only by the number of points N , and (ii) the number of repeaters is taken to be R . We give results for repeaters which require midpoint stations, providing formulas for the number of configurations when each midpoint station placement is (not) counted as a separate configuration. We consider also the simple case in which the repeaters do not require the installation of midpoint stations.

4.1 Introduction

It is likely that near-term deployment of quantum networks will make use of existing fiber infrastructure (see Chapters 6 and 7, as well as [1]). This infrastructure places constraints on where quantum hardware, such as quantum repeaters and midpoint stations, can be placed. For cost-effectiveness, this would ideally happen only at points where fiber emerges from the ground [1]. The placement of the quantum hardware can have significant impact on the performance that the corresponding repeater chain can achieve (see Chapters 6 and 7 as well as [1, 2]). It is therefore of interest to know what are the unique ways in which quantum repeaters can be placed to create a quantum-repeater chain along a fiber path with N points where hardware can be installed.

¹These authors contributed equally.

This chapter is based on the article in preparation: Francisco Ferreira da Silva, Guus Avis, Kenneth Goodenough and Stephanie Wehner. "Counting quantum-repeater configurations".

In this chapter, we calculate two different quantities. In what we call *case 1*, we consider that we can place as many repeaters as allowed by the chain. In *case 2*, we consider that we are restricted to placing R repeaters. Case 2 might be of interest if, for example, in the deployment of a quantum-repeater chain there were only budget to install a limited number of quantum repeaters.

Some types of quantum repeaters require a *midpoint station* between them to function. For example, as discussed in Chapter 3, first- and second-generation quantum repeaters require a heralding station that can perform photonic Bell-state measurements. Other repeaters (see, e.g., [3] for a review) might require a station endowed with entangled-photon-pair sources. In this case, the points in the fiber path in which fiber comes off the ground and quantum hardware can be installed must be used to install not only quantum repeaters but also midpoint stations. Therefore, we will consider three different scenarios:

4

- Scenario 1: no midpoint stations need to be installed;
- Scenario 2: there need to be midpoint stations installed between each pair of quantum repeaters, but we do not count the different ways in which the midpoint stations can be placed;
- Scenario 3: there need to be midpoint stations installed between each pair of quantum repeaters, and we count each different way in which the midpoint stations can be placed as a separate configuration.

We denote the number of configuration corresponding to Case 1 and Scenario i as $C_i(N)$ and corresponding to Case 2 and Scenario i as $C_{i,R}(N)$. Note that the following relation holds between the two:

$$\sum_R C_{i,R}(N) = C_i(N). \quad (4.1)$$

This means that summing over the different configurations corresponding to placing R repeaters gives the total number of configurations. We will now start by giving the results for all the cases and scenarios we consider. The corresponding proofs can be found at the end of the chapter.

We note that some of the results shown below are trivial (e.g., placing R repeaters in N points without midpoint stations). We nevertheless include them for completeness.

4.2 Results

Here we present formulas without derivation.

4.2.1 Scenario 1: no midpoint stations

Case 1:

$$C_1(N) = 2^N. \quad (4.2)$$

Case 2:

$$C_{1,R}(N) = \binom{N}{R}. \quad (4.3)$$

4.2.2 Scenario 2: with midpoint stations but not counting them

Case 1:

$$C_2(N) = \mathcal{F}_N, \quad (4.4)$$

where \mathcal{F}_N is the N^{th} Fibonacci number.

Case 2:

$$C_{2,R}(N) = \binom{N-R-1}{R}. \quad (4.5)$$

4.2.3 Scenario 3: with midpoint stations and counting them

Case 1:

$$C_3(N) = \sum_{R=0}^{\lfloor \frac{N-1}{2} \rfloor} \binom{N}{2R+1} = 2^{N-1}, \quad (4.6)$$

where $\lfloor \cdot \rfloor$ denotes the floor function.

Case 2:

$$C_{3,R}(N) = \binom{N}{2R+1}. \quad (4.7)$$

4.3 Proofs

Here we give proofs for each of the formulas in the Results section.

4.3.1 Scenario 1: no midpoint stations

Case 1: Each of the N points can either be left empty or have a repeater. The total number of configurations can be calculated by considering each point in turn and multiplying the number of options. This leads to a total of 2^N different configurations.

Case 2: Each of the N points can either be left empty or have a repeater, up until R repeaters have been placed. Given that repeaters are taken to be indistinguishable, the number of configurations is given by N choose R .

4.3.2 Scenario 2: not counting midpoint stations

We start by determining what is the maximum number of repeaters that can be placed, as this will be useful in all the following derivations. One midpoint station must be placed between every two repeaters as well as between each of the end nodes and the repeaters. Therefore, if we wish to place R repeaters, we must also place $R+1$ midpoint stations and

the total number of quantum-hardware nodes to be placed N is $2R + 1$. If N is odd we can directly write

$$R(N) = \frac{N-1}{2}. \quad (4.8)$$

If N is even, we can think of the fiber path as a fiber path of $N - 1$ points plus one extra point. This additional node cannot be used to place an additional repeater, as placing one more repeater also requires placing one more midpoint station. We can therefore write

$$R(N) = \lfloor \frac{N-1}{2} \rfloor. \quad (4.9)$$

This is the maximum number of repeaters that can be placed in a fiber path with N points where hardware can be placed, assuming that the repeaters require midpoint stations.

4

Cases 1 and 2: Each of the N points can either be left empty, have a repeater or have a midpoint station. There must be one and only one midpoint station between repeaters and between repeaters and the end nodes. We can think of the repeaters as dividing the fiber path into bins. In each of the bins we must place one and only midpoint station. This is the stars-and-bars problem (see, for example, [4]), which has the following known solution for placing R repeaters:

$$C_{2,R}(N) = \binom{N-R-1}{R}. \quad (4.10)$$

The total number of possible configurations is then given by summing over this expression up until the maximum number of repeaters that can be placed in a path of N points:

$$C_2(N) = \sum_{R=0}^{\lfloor \frac{N-1}{2} \rfloor} \binom{N-R-1}{R}. \quad (4.11)$$

The sequence defined by increasing N in this sum corresponds to the Fibonacci sequence. We will now prove this by induction.

We start by recalling that the Fibonacci sequence can be defined by the following recurrence relation:

$$\begin{aligned} \mathcal{F}_N &= \mathcal{F}_{N+1} + \mathcal{F}_{N+2}, \quad \text{for } N > 1, \\ \mathcal{F}_0 &= 0, \\ \mathcal{F}_1 &= 1. \end{aligned}$$

As a base case, we have

$$\begin{aligned} C_2(N=2) &= \binom{1}{0} = 1 = \mathcal{F}_2, \\ C_2(N=3) &= \binom{2}{0} + \binom{1}{1} = 2 = \mathcal{F}_3. \end{aligned}$$

For the induction step, we will show that $C_2(N+1) = \mathcal{F}_{N+1}$ given that $C_2(N) = \mathcal{F}_N$ and $C_2(N-1) = \mathcal{F}_{N-1}$.

We have

$$C_2(N+1) = \sum_{R=0}^{\lfloor \frac{N}{2} \rfloor} \binom{N-R}{R} = \binom{N}{0} + \sum_{R=1}^{\lfloor \frac{N}{2} \rfloor} \left(\binom{N-R-1}{R} + \binom{N-R-1}{R-1} \right), \quad (4.12)$$

in which we have split the $R=0$ term from the sum and then made use of Pascal's identity.

$$\begin{aligned} C_2(N+1) &= 1 + \sum_{R=1}^{\lfloor \frac{N}{2} \rfloor} \binom{N-R-1}{R-1} + \sum_{R=1}^{\lfloor \frac{N-1}{2} \rfloor} \binom{N-R-1}{R} \\ &= \sum_{R=0}^{\lfloor \frac{N-1}{2} \rfloor} \binom{N-R-1}{R} + \sum_{R=0}^{\lfloor \frac{N}{2} \rfloor - 1} \binom{N-R-2}{R} \\ &= C_2(N) + C_2(N-1) \\ &= \mathcal{F}_N + \mathcal{F}_{N-1} \\ &= \mathcal{F}_{N+1}. \end{aligned}$$

This concludes the proof.

4.3.3 Scenario 3: counting midpoint stations

Cases 1 and 2: Each of the N points can either be left empty, have a repeater or have a midpoint station. We must place R repeaters and $R+1$ midpoint stations. This means that $N-2R-1$ points will be left vacant. In contrast with the previous scenario, which points will be left vacant is now relevant to counting the number of configurations. We can therefore now think of the problem as one of placing vacant stations into different positions in the sequence of repeaters and midpoint stations. There are a total of $2(R+1)$ points where vacant stations can be placed, namely before and after each midpoint station. This is again the 'stars and bars' combinatorics problem, in which we want to distribute identical items (the vacant spots) into distinct bins (the positions in the sequence). The number of ways of distributing n identical items into k distinct bins is

$$\binom{n+k-1}{k}. \quad (4.13)$$

In this case n is the number of vacant spots, $n = N - 2R - 1$, and k is the number of distinct positions in the sequence, $k = 2(R+1)$. Therefore,

$$C_{3,R}(N) = \binom{N}{2R+1}. \quad (4.14)$$

In order to obtain $C_3(N)$, we simply sum over $C_{3,R}(N)$ up until the maximum number of repeaters that can be placed in a fiber path with N points.

$$C_3(N) = \sum_{R=0}^{\lfloor \frac{N-1}{2} \rfloor} \binom{N}{2R+1} = 2^{N-1}. \quad (4.15)$$

References

- [1] J. Rabbie, K. Chakraborty, G. Avis, and S. Wehner, *Designing Quantum Networks Using Preexisting Infrastructure*, npj Quantum Information **8**, 5 (2022), 2005.14715 .
- [2] G. Avis, R. Knegjens, A. S. Sørensen, and S. Wehner, *Asymmetric node placement in fiber-based quantum networks*, arXiv preprint arXiv:2305.09635 (2023).
- [3] W. J. Munro, K. Azuma, K. Tamaki, and K. Nemoto, *Inside Quantum Repeaters*, IEEE Journal of Selected Topics in Quantum Electronics **21**, 78 (2015).
- [4] W. Feller, *An Introduction to Probability Theory and Its Applications* (Wiley, 1950).

5

Optimizing entanglement generation and distribution using genetic algorithms

5

Francisco Ferreira da Silva, Ariana Torres-Knoop, Tim Coopmans, David Maier and Stephanie Wehner.

Long-distance quantum communication via entanglement distribution is of great importance for the quantum internet. However, scaling up to such long distances has proved challenging due to the loss of photons, which grows exponentially with the distance covered. Quantum repeaters could in theory be used to extend the distances over which entanglement can be distributed, but in practice hardware quality is still lacking. Furthermore, it is generally not clear how an improvement in a certain repeater parameter, such as memory quality or attempt rate, impacts the overall network performance, rendering the path towards scalable quantum repeaters unclear. In this work we propose a methodology based on genetic algorithms and simulations of quantum repeater chains for optimization of entanglement generation and distribution. By applying it to simulations of several different repeater chains, including real-world fiber topology, we demonstrate that it can be used to answer questions such as what are the minimum viable quantum repeaters satisfying given network performance benchmarks. This methodology constitutes an invaluable tool for the development of a blueprint for a pan-European quantum internet. We have made our code, in the form of NetSquid simulations and the smart-stopos optimization tool, freely available for use either locally or on high-performance computing centers.

5.1 Introduction

In this chapter we propose a methodology based on genetic algorithms and simulations of quantum repeaters for optimization of entanglement generation and distribution in quantum networks. This allows us to find minimal requirements on quantum-repeater hardware for given performance metrics. Later on in this dissertation, namely in chapters 6, 7 and 8, we apply it to determine minimal hardware requirements in different scenarios.

Contrasting with previous work on repeater chain optimization [1–6], our methodology constitutes a systematic and modular approach to this problem, successfully integrating simulation and optimization tools, as well as allowing for the use of high-performance computing clusters. A high-level overview of how a user interfaces with this process is shown in Figure 5.1.



Figure 5.1: Overview of our optimization process. The user inputs the desired optimization parameters and defines a cost function. Using simulation and optimization tools, our methodology finds a set of parameters optimizing the cost function. For example, the optimization parameters could be parameters defining a quantum repeater model and the cost function could be the inverse of the secret key rate plus a penalty term for parameter values that are much better than a given baseline. The output would then be the values of the parameters defining the quantum repeater model optimizing the cost function.

We performed our simulations using NetSquid [7, 8]. NetSquid can accurately model the effects of time-dependent noise, rendering it well equipped to predict quantum network performance in a physically accurate setting. The tools used in this methodology, which allow for running NetSquid simulations together with an optimization algorithm both locally and on an HPC cluster, are made freely available (see [9]).

This chapter is structured as follows. In Section 8.2 we introduce our methodology, together with the required preliminaries. Section 5.3 concerns the validation of the methodology. This comprises two steps: (i) benchmarking our GA implementation by running it on standard optimization problems and comparing its performance to those found in the literature; and (ii), validating our approach by applying it to a repeater chain where elementary link states are in the Werner form and all noise sources are depolarizing [10]. In this case, analytical results can be found, so we can evaluate how well our optimization method performs.

After validating our methodology, we apply it to some different repeater chain setups, in order to demonstrate its potential usefulness. We present these results in Section 5.4, where we first consider a repeater chain based on real-life fiber data, courtesy of SURF, a classical network provider for Dutch education and research institutions. This showcases the power of our simulation-based approach, as chains of unevenly spaced nodes

are hard to study analytically. We further apply our methodology to chains of varying length, internode distance and number of repeaters and we compare the solutions found with our methodology for each of these different setups. This allows us to investigate how the impact of the parameters varies across setups, thus identifying possible bottlenecks and paths towards scalable quantum repeaters.

5.2 Methodology

In this section we introduce the main contribution of our work, a methodology for the optimization of entanglement generation and distribution. We first present each of the elements that are used in this optimization process. We finalize the section with an overview of how they are integrated to answer the question of what are the minimum requirements on quantum repeaters to achieve a given benchmark.

5.2.1 Question

We aim to answer the question of what the minimum requirements are on the quality of quantum repeaters to achieve a given benchmark by framing it as an optimization problem. To do so, we must first clarify what we mean by requirements and by quality of a quantum repeater. Let us say that a quantum repeater is described, in a given model, by a set of N parameters $\{x_i\}_{i \in \{1, \dots, N\}}$. The meaning of x_j is model-dependent. For example, if we consider a model of a trapped ion system, x_j and x_k could be single-qubit and two-qubit gate error probabilities. We could also, in a more abstract model, combine these two parameters together to obtain a swap quality that quantifies the noise introduced in an entanglement swap operation, which would then be y_j in this model. The quality of a quantum repeater is then a function of the set of parameters describing it. This also helps clarify what we mean by requirements. Suppose we have some fixed network topology and performance metric. To give a concrete example, the topology could be a repeater chain of 10 equally spaced nodes and the performance metric the end-to-end secret key rate. The requirements on the repeaters are then the worst set of parameters that enable attaining some value of the end-to-end secret key rate over a chain of 10 nodes, i.e. the lowest quality repeaters satisfying this metric. The meaning of repeater quality will be made clear in the following section.

5.2.2 Cost

Let us say that we have two repeaters described by a set of parameters $\{y_i\}_{i \in \{1, \dots, N\}}$ and $\{z_i\}_{i \in \{1, \dots, N\}}$, and that the values of these parameters are the same for all but two of them, i.e. $\{y_i\} = \{z_i\} \forall i \in \{1, 2, \dots, N\} \setminus \{j, k\}$. Let us further say that y_j is better than z_j , but z_k is better than y_k . Which of these sets of parameters is the better one? To answer this, we will now introduce the quantity to be optimized, the cost function. We emphasize that our method is completely general and could be applied to any cost function, but for concreteness we focus on a particular one from here on out.

We expect that in an experimental setting a given physical parameter becomes harder to improve the closer to its perfect value it is, so we would like our cost function to reflect this. We start by transforming our parameters so that they all live in the $[0, 1]$ interval, with 1 being the perfect value and 0 the worst possible value. We refer to Appendix 5.9

for details. Denoting x_b as the baseline value of a parameter, i.e. the value from which we are improving, k as the improvement factor and x_{new} as the new improved value, we claim that the following equation reflects this behaviour:

$$x_{\text{new}}(k) = x_b^{\frac{1}{k}}. \quad (5.1)$$

This can be read as: we improve x_b by a factor k to get x_{new} . To see that Equation 5.1 does in fact reflect the desired behaviour, we note the that

$$x_{\text{new}}(k = 1) = x_b, \quad (5.2)$$

$$\lim_{k \rightarrow \infty} x_{\text{new}} = 1. \quad (5.3)$$

Equation (5.2) can be read as: improving a parameter by a factor of 1 is equivalent to not improving it all, whereas Equation (5.3) can be taken to mean that in order to improve a parameter to its perfect value we must improve by a factor of infinity, i.e. there is no such thing as a perfect process.

We can then define the cost associated to x_{new} as the factor k by which we must improve the baseline value x_b to obtain x_{new} . Therefore, solving Equation (5.1) for k , we get

$$k = \frac{1}{\log_{x_b}(x_{\text{new}})}. \quad (5.4)$$

With this in hand, we can finally define the cost associated to a set of parameters. Let us say our model is described by a set of parameters $\{x_i\}_{i \in \{1, \dots, N\}}$, and that the current baseline value of each of these parameters is $\{x_{i_b}\}_{i \in \{1, \dots, N\}}$. A set of values $\{x_{i_c}\}_{i \in \{1, \dots, N\}}$ is mapped to a cost, C , by Equation (5.5). Intuitively, this can be seen as taking the average of the cost associated to each of the parameters.

$$C(x_{1_c}, \dots, x_{N_c}) = \sum_{i=1}^N \frac{1}{\log_{x_{i_b}}(x_{i_c})} \quad (5.5)$$

Note that with this definition, the minimum parameter cost is N , with N being the number of parameters in the model under consideration. Since this cost function is meant to be used for comparing the relative cost of parameter sets of the same model, N is the same for all parameter sets under consideration, and hence it is nothing but a constant shift in each set's cost. One could, for instance, divide the cost by N to normalize it or subtract N from it, making the minimum cost 0. This would however have no impact on the results obtained, since the relative ordering of parameter sets according to their cost would remain the same.

There is still the matter of how the network's target performance metrics are taken into account. Throughout this work we will focus on fidelity F of the end-to-end state with the ideal Bell state and entanglement generation rate R , but we stress that our method is not limited to optimizing for these quantities. More concretely, we will try to answer the question of what are the minimum requirements on repeaters to concurrently achieve certain values of F and R . We are then faced with a multi-objective problem, as we want

to optimize multiple quantities simultaneously, namely end-to-end fidelity, entanglement generation rate and parameter cost. Furthermore, there are trade-offs between these goals. For example, improving the memory lifetime of nodes in a chain has a positive contribution towards end-to-end entanglement fidelity, but a negative one towards parameter cost. There is a multitude of possible ways of approaching such problems [11]. We chose to map our multi-objective optimization problem to a single-objective one by assigning weights to the different objectives and adding them, a process known as scalarization [12]. In this way, the total cost function T_C to minimize becomes a weighted sum of the parameter cost and the thresholds on end-to-end rate and fidelity:

$$T_C(p_{1_c}, \dots, p_{N_c}, F_{min}, R_{min}) = w_1 \Theta(F_{min} - F) + w_2 \Theta(R_{min} - R) + w_3 C(x_{1_c}, \dots, x_{1_N}), \quad (5.6)$$

where the w_i are the weights of each objective, Θ is the Heaviside function and F_{min} and R_{min} are, respectively, the minimum required end-to-end fidelity and end-to-end entanglement generation rate. Using step functions reflects the idea that we are looking for solutions that satisfy performance benchmarks, with no reward given for surpassing them. The weights in Equation (5.6) are hyperparameters of our method, meaning that they are not determined by some algorithm but must instead be chosen. This choice can be of any real number, and it has an impact on which sets of parameters have the lowest costs and hence on the solutions found by the method. For example, if we assign very high values to w_1 and w_2 and a low value to w_3 the best sets of parameters will be those that satisfy the requirements on the end-to-end fidelity and rate without much regard for how costly it is to achieve them. To give a concrete example of what the hyperparameter values might be, for the applications we present in section 5.4, we set w_1 and w_2 to 20000 and w_3 to 1. The parameter cost term, defined in Equation (5.5), depends on the baseline values of the parameters, which must also be chosen. Typically, for the use cases we consider, these will be chosen to reflect what is currently achievable experimentally.

Optimal solutions to this single-objective optimization problem are then solutions to the multi-objective optimization problem.

5.2.3 Abstract model

In order to explore and better understand the methodology we propose, we believe it to be wise to employ a relatively simple model whose behavior we understand. We must however again emphasize that our methodology is completely general in terms of the model used for the quantum repeater hardware.

We consider a simplified five-parameter model for a quantum repeater, the five parameters being denoted by $[F_{EL}, p_{suc}, s_q, T_1, T_2]$. We assume that elementary links states have fidelity F_{EL} with the ideal Bell state upon generation, and that they are generated with a success probability p_{suc} . We assume also that each swap introduces depolarizing noise parametrized by a swap quality s_q and that memory decoherence is described by a T_1, T_2 process, with T_1 (T_2) being the memory's relaxation (dephasing) time. In simple terms, this means that T_2 determines how fast the off-diagonal components of the density matrix decay, whereas T_1 defines how long it takes for a quantum system to relax to its lowest energy state. For more details on T_1, T_2 noise processes see Appendix 5.10, where our parametrization of depolarizing noise is also clarified. We further assume that entanglement swapping, although noisy, is deterministic. We note that this model is quite

abstract. It could, in principle, describe the behaviour of any repeater of the processing node type, examples being NV centers and trapped ions. By this we mean that it is possible to map the parameters in a physically accurate model of an NV center or trapped ion to this smaller set of more abstract parameters. In fact, we did exactly this for NV centers in order to validate this model, as laid out in Appendix 5.10. It is important to note that in this mapping we considered induced dephasing noise instead of the usual memory dephasing. Furthermore, atomic ensemble based repeaters could be described by considering non-deterministic entanglement swaps and enriching the model with a swap success probability parameter, but this lies beyond the scope of this work.

We again highlight that this model was chosen for demonstrative purposes, and that our methodology could just as well be applied to more realistic hardware models, as discussed in Section 5.5.

5.2.4 Genetic algorithms

Evolutionary Algorithms (EAs) have been shown to have an advantage over conventional gradient-based methods in finding global minima in multimodal functions whose search space is not well known [13], although we stress that this is not guaranteed. They are also robust to noise in data and easy to parallelize. There are multiple approaches within the umbrella of EAs, with prominent examples being genetic algorithms (GAs) [14], evolution strategy [15], differential evolution [16] and particle swarm optimization [17, 18]. In this work we have used GAs, a search heuristic inspired by the theory of evolution. We limit ourselves to a high-level overview of GAs. For a comprehensive introduction, we direct the interested reader to [19].

We start with a population of randomly generated individuals. In our case, each individual in a population is a set of values for the parameters of the abstract model introduced in section 5.2.3. The GA generates new individuals in an iterative process, with each iteration being known as a generation. In each generation, the cost function is evaluated for every member of the population, the resulting value being known as the fitness. A subset of the population is then selected according to a fitness-dependent rule, in which higher-fitness solutions are more likely to be chosen. New individuals are then generated through random crossover and mutation operations. The new population is used for the following iteration of the algorithm, meaning that the simulation is run with the new individuals (i.e. sets of abstract model parameters) as input and the cost function is computed using the simulation outputs. The algorithm can terminate after a set number of generations or once some predefined condition is attained. For the examples given in this work, we have chosen to use the first condition and let the algorithm terminate after a preset number of generations, typically 150. Exploration of the search space is assured by the crossover and mutation-driven recombination of solutions, whereas fitness-based selection ensures exploitation of minima.

GAs come in several different flavours. See Appendix 5.9 for details on our particular implementation.

5.2.5 *smart-stopos*

The simulation tools we use are computationally heavy and produce large amounts of data. In order to make good use of them and extract useful information from said data,

we need a systematized way of feeding input parameters to the simulations in batches, run the simulations on a high-performance computing (HPC) cluster using *stopos* [20], feed the outputs to the optimization algorithm and iterate this procedure. To these ends, we made use of *smart-stopos* (freely available at [9]), a set of tools we developed to allow for parameter exploration and optimization, both locally and in an HPC setting. *smart-stopos* can be seen as an addition to *stopos*, extending its capabilities by allowing for the seamless integration of simulation and optimization tasks. We used GAs in this work but in principle any other algorithm could be plugged in, provided that it can be run with only simulation inputs and outputs. Furthermore, we note that we used NetSquid but our methodology could also be made to work with any other quantum network simulator. For more details on the use of *smart-stopos*, we direct the interested reader to Appendix 5.8.

5.2.6 Process overview


We will now show how the tools we introduced can be pieced together to answer the question of what the minimum requirements are on the quality of quantum repeaters. To that end, we show in Figure 5.2 a diagram of the workflow of our methodology.

The process is started by defining the parameters to be optimized and their allowed range of values. This information, together with a termination criterion, is passed to *smart-stopos*, which then randomly generates sets of parameters within the defined ranges. Each of these sets of parameters is fed to the NetSquid simulation, which outputs an end-to-end entangled state and the time its generation took, allowing us to compute the fidelity with the ideal Bell state and the entanglement generation rate. Note that these quantities are stochastic, so throughout this work we average them over multiple runs of the same setup. These metrics, together with the parameter values and the baseline values, are used to compute the cost function, as defined in Equation 5.5. This process is then repeated for each set of parameters. The ensemble of parameter sets and respective costs are given as input to *smart-stopos*, which generates new sets of parameters using our GA. The process repeats until the termination criterion is reached. The final output is the minimum value of the cost function found by the algorithm, which in this case corresponds to an answer to the question of what are the minimum requirements on a quantum repeater.

Figure 5.2 makes the modularity of our approach clear. Any of the building blocks of our process, namely the optimization algorithm used by *smart-stopos*, NetSquid simulation and cost function, can be swapped out without changes to the overall workflow. For example, if we wanted to apply our methodology to a simulation of a repeater chain of trapped ions, all we would have to do would be to replace our abstract model NetSquid simulation for an appropriate trapped ions simulation. Similarly, to answer a different optimization question one just has to redefine the cost function.

5.2.7 Challenges in applying genetic algorithms to quantum systems

We came across some challenges when applying GAs to simulations of quantum systems. Some of these were of a practical nature, and others were more fundamental. We will now give an overview of what these issues were, and how we overcame them.



5

Figure 5.2: Overview of our optimization process. The user inputs the desired optimization parameters, their ranges and a stopping criterion. *smart-stopos* generates sets of parameters in the allowed range and feeds them to the NetSquid simulation. The outputs of the simulation are used to compute the cost associated to each parameter set, which in turn is used by *smart-stopos* to generate new parameter sets. This process is repeated until the stopping criterion is reached. In our particular case, the optimization parameters are the parameters defining the abstract repeater model introduced in 5.2.3, the relevant simulation outputs are the fidelity and generation rate of end-to-end entangled states and the cost function is the one defined in 5.2.2.

Practical challenges

We came across two practical challenges: (i) the size of the parameter space and (ii) the amount of data generated. (i) is due to the complexity of quantum repeater modelling. In general the search space may be big, but in our illustrative example of the abstract model introduced in 5.2.3 it is manageable. We nevertheless introduced a pre-processing procedure for restricting the parameter space, as we believe it would be useful when considering

use cases with larger parameter spaces. This procedure consists of performing sensitivity analysis for each of the five parameters individually, i.e. holding four parameters constant and running simulations varying the fifth one from its baseline value to its perfect one. As an example of how this can reduce the search space, we show in Figure 5.3 the variation of the end-to-end fidelity with the elementary link fidelity when all other parameters are kept at their perfect values. The optimal set of parameters for this setup will certainly contain less-than-perfect values, so the elementary link fidelity of this set will be higher than the one found using this sensitivity analysis, so we can safely restrict the search space for this parameter in GA optimizations runs to the interval $[f_{\text{perf}}, 1.]$, where f_{perf} is the elementary link fidelity that results in an end-to-end fidelity of 0.7 when all other parameters are perfect.




Figure 5.3: Variation of end-to-end fidelity across five equally spaced nodes as the elementary link fidelity is varied and the other four parameters are kept at their perfect values. The value of the elementary link fidelity that results in an end-to-end fidelity of 0.7 is at the intersection of the two lines in the plot, being just above 0.9 in this case.

Another practical challenge is the sheer amount of data that is produced. For each setup we consider we run our simulations for hundreds of different sets of parameters at each optimization step, with each set of parameters being in turn run a hundred times. In order to systematically and efficiently process all of this data, we developed *smart-stopos*, as detailed in Section 5.2.5.

Fundamental challenges

Fundamental challenges occur due to the fact that quantum systems produce inherently non-deterministic outputs. This can be problematic if the cost function has terms that are step functions, which is our case. For a concrete example, let us say that in generation 34

of the optimization procedure, the GA found a set of parameters that result in an entanglement generation rate of 1.05 Hz, just above the desired threshold. In generation 35, this parameter set would again be fed into the simulation. However, this time around, due to statistical fluctuations, the simulation outputs an entanglement generation rate of 0.99 Hz, just below the threshold. Since the cost function defined in Equation (5.6) assigns a very high cost to any solution that does not attain the performance metrics, this solution would in generation 35 have a very high cost function value. This means that it would almost certainly not be chosen as a parent for the following generation, and the algorithm would effectively lose it. This is a problem, as it results in the algorithm losing a good solution and potentially wasting computation time finding it again.

There are several possible solutions to this problem. The one we chose, due to its simplicity, was to run the simulation multiple times for each set of parameters and compute the value of the cost function using the average end-to-end fidelity and entanglement generation rates. Running the simulations multiple times provides some security against statistical fluctuations, although it increases the computation time. We found empirically that running the simulation 100 times for each set of parameters represents a good trade-off between minimizing fluctuations and keeping computation times feasible.

Another possible solution that we also explored was to use a smoother function, such as a sigmoid, instead of a sharp step function. This would in principle address the problem we mentioned of a set of parameters being heavily penalized because its metrics dipped just below the targets due to statistical fluctuations. For a smoother function, such fluctuations would lead to small fluctuations in the value of the cost function. There are however some issues with this solution. Since the function is smoother, it no longer acts as a hard constraint, which is the behaviour we are looking for. What we mean by this is that a solution whose performance metrics are slightly below the targets will only be lightly penalized. It might thus have a lower cost function value than a solution with better, i.e. more expensive, parameters that attains the performance metrics. In less technical terms, this translates as the cost function not being well aligned with the stated optimization goal.

This concludes the introduction of the optimization methodology we propose. The rest of the chapter concerns itself with two questions: (i) is our methodology valid, addressed in Section 5.3 and (ii) what results do we get when we apply it, addressed in Section 5.4.

5.3 Validation

As we stated in the previous section, before we apply our methodology we must validate it. By this we mean that we must verify that the methodology we propose for applying genetic algorithms to simulations of quantum networks can produce meaningful results. We can see this validation as being split into two different steps. One, benchmarking the genetic algorithms i.e., evaluating how well they perform and two, validating that the methodology is sound. The first step will be accomplished by applying our specific implementation of genetic algorithms to the optimization of common benchmarking functions and comparing their performance to that of implementations found in the literature. The second step will consist of applying our methodology to a chain of evenly-spaced nodes generating Werner states, for which analytical expressions for the end-to-end fidelity and entanglement generation rate in terms of repeater parameters can be found. Having these

expressions, we can compute what are the repeater parameters that minimize the cost function. If our GA approach is capable of finding this solution, we have compelling evidence that our methodology would also perform well when applied to the more realistic cases we are interested in, for which analytical results cannot be readily derived.

We have also validated the abstract model we use in our simulations against a more physically accurate model of NV center-based repeaters. These results are shown in Appendix 5.10.

5.3.1 Benchmarking genetic algorithms

In order to evaluate the performance of GAs and how it is affected by the algorithm's hyperparameters, several benchmarking functions have been defined [21]. These are designed to test how well each GA implementation handles cost functions with given properties. For example, if we expect the function we want to optimize to be noisy, i.e. to have the output for a given input randomly oscillate each time the function is called, we should benchmark the GA against a noisy function, such as the quartic function, defined in Equation (5.7)


$$f_q(x) = \sum_{k=1}^{30} (kx_k^4 + \mathcal{N}(0, 1)) \quad -1.28 \leq x_k \leq 1.28, \quad (5.7)$$

where $\mathcal{N}(0, 1)$ is a normal distribution with mean 0 and standard deviation 1. This function, plotted in the bottom half of Figure 5.4, is a unimodal function padded with Gaussian noise. Therefore, a GA that performs poorly on it will also perform poorly on any function with noisy outputs.


Taking this into account, we chose two functions to benchmark our GA implementations. This choice was made by taking into account which of the functions best represented the cost landscape we expect our problem to have. Since the quantum nature of our simulations implies that they will necessarily be noisy in the above-defined sense, we will choose the quartic function as a benchmarking function. Furthermore, we expect that the landscape of the cost function defined in Equation (5.5) will have multiple local minima, corresponding to different sets of parameters that satisfy the imposed constraints on end-to-end fidelity and entanglement generation rate. With this in mind, we also chose Rastrigin's function, defined in Equation (5.8).

$$f_r(x) = 200 + \sum_{i=1}^{20} (x_i^2 - 10 \cos(2\pi x_i)), \quad -5.12 \leq x_i \leq 5.12 \quad (5.8)$$

For illustrative purposes, the 2-dimensional version of Rastrigin's function is shown on the top half of Figure 5.4. It can be seen that it has a very bumpy landscape, with a global minimum at $\mathbf{0}$, in the center of the plotted region. Its many local minima render it a challenging benchmark for GAs. We applied our GA implementation to both of these functions, with the results being plotted in Figure 5.5. The hyperparameters used for these optimization runs were chosen according to the guidelines given in [21] and population selection was done using the Roulette Wheel method [22]. For an explanation of the Roulette Wheel method we point the interested reader to Appendix 5.9. By best value we mean the lowest value of the cost function achieved by any of the parameter sets in the population.



(a) Rastrigin's function.



(b) Quartic function.

Figure 5.4: Plot of the 2-dimensional versions of (a) Rastrigin's function and (b) quartic function. The multiple minima of Rastrigin's function and the noisy landscape of the quartic function can be clearly seen.

Similarly, by average value we mean the average of the costs of all parameter sets in the population. We see that, for both functions, the average cost and the best cost at each generation approach their global minimum, 0. Furthermore, the performance of our implementation is in line with that of those in [21], which indicates that our GA is capable of handling both noisy and multimodal functions. We note that convergence requires significantly more generations for Rastrigin's function than for the quartic function. This reflects the well-known fact [21] that multimodal functions are challenging for GAs. We must also note that we could, by further tuning some of the algorithm's hyperparameters, obtain a marginally better performance on these benchmarking functions. However, since our goal is only to verify that our implementation is correct and performs reasonably well for the type of cost landscapes that we expect to encounter, we abstain from doing so.

5.3.2 Validating on Werner chains

The previous section focused on benchmarking the performance of the GA, but the question of whether applying GAs to repeater chain optimization problems can produce good

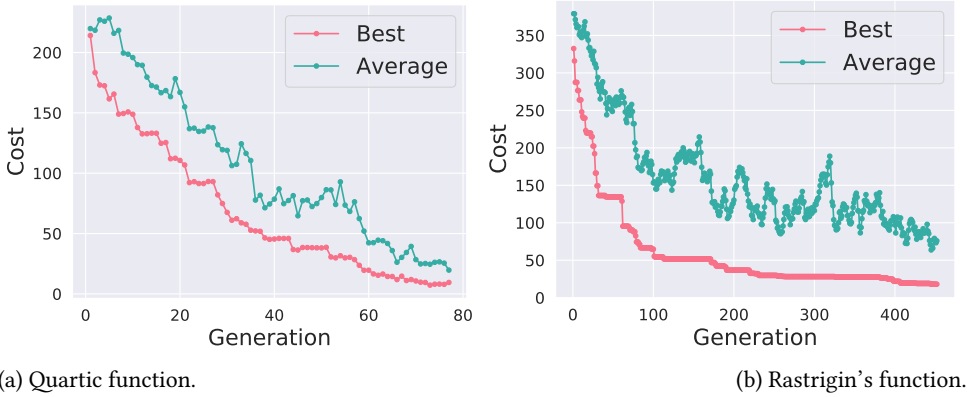


Figure 5.5: Evolution of the cost of best solution (red) and population average (green) for the (a) quartic function and (b) Rastrigin's function over 75 and 400 generations, respectively. The data used in (a) ((b)) was acquired in roughly 1h30 (26h) on consumer-market hardware (Intel Core i7-8665U and 8 GB RAM). These runtimes can be significantly reduced via parallelization and use of high-performance computing clusters. All costs approach zero, the global minimum of both cost functions, with the average cost being consistently higher than the best cost, as expected. This indicates that our GA implementation is capable of finding good solutions for these functions.

results remains. In order to answer it, we consider the simple scenario of a chain of 3 nodes generating Werner states, and we pose the question of what are the worst parameters that can deliver an end-to-end entangled pair of fidelity 0.6 every second. Similarly to the abstract model presented in earlier sections, the nodes in the chain generate elementary links of fidelity F_{EL} with success probability p_{suc} and depolarizing noise parametrized by s_q is applied after entanglement swaps. This is a problem for which we can analytically find expressions for the end-to-end rate and fidelity, and thus for the ideal value of the cost function. We expect that the structure of this problem is similar to that of the one we want to tackle. By this we mean that we expect its cost landscape to show some of the same features as our target problem, namely multiple minima and noisiness. Therefore, despite being simpler, good performance in this problem should indicate that our approach is valid. For details of how we derived analytical results for this setup we defer the interested reader to Appendix 5.11.

In Figure 5.6, we show the evolution of the cost of the best individual in each generation obtained by applying the GA-based method to the setup we described. Also present in the plot, in a dashed line, is the optimum cost.

The cost function drops to the global optimum at around the 30 generation mark, indicating that the algorithm is capable of finding the worst set of repeater parameters satisfying the benchmarks we set. This is then a good indicator that our methodology is well-suited to the optimization of entanglement generation in repeater chains.

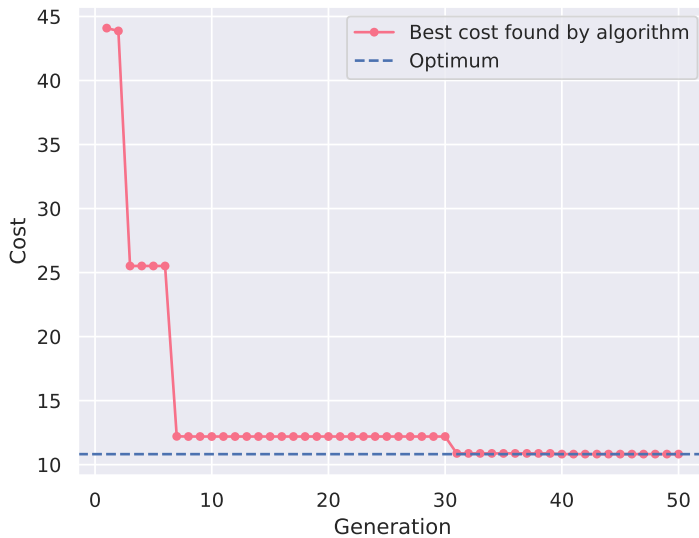


Figure 5.6: Evolution of the lowest value of the cost function over 50 generations. After little more than 30 generations the algorithm finds the parameter set that optimizes the cost function. This optimum is marked in the figure by a blue dashed line.

5.4 Evaluation: use cases

Having validated our methodology, we applied it to two use cases demonstrating its power and potential usefulness. In the past decade, NV centers have been demonstrated to be capable of generating remote entanglement between matter memories with long coherence times [23–25], establishing them as promising candidates for the realization of scalable quantum repeaters [26]. A better understanding of hardware requirements would then be useful in illuminating the path towards scalable NV-based quantum repeaters. We thus used the abstract model of NV-type states that we introduced in Section 5.2.3 in the simulations of all use cases. We furthermore chose to consider, for simplicity, SWAP-ASAP protocols with no memory cut-offs. More precisely, we simulate the protocol introduced in Appendix E 2 of [8], which proceeds as follows: we assign indices to each node going from left to right in the chain and starting with 1. Even-numbered nodes are called initiators, whereas odd-numbered nodes are called responders. As the name implies, initiators are responsible for initiating the process of entanglement generation, which they do by sending a request for entanglement generation to their left-hand neighbors and waiting for a response. Once the responders respond, the process of entanglement generation begins. This is simulated by sampling the time taken to generate entanglement according to the success probability parameter and the cycle time, which determines how long a single entanglement generation attempt takes. Once entanglement is successfully generated, the initiator proceeds to attempt to generate entanglement with the right-hand neighbor and the process unfolds in the same way. Whenever a node holds two entangled qubits in hand, it performs an entanglement swap by measuring them in the Bell basis. The simulation stops once the end nodes of the chain share an entangled pair.

Another roadblock in the way of the quantum internet is that even when quantum repeater technology is at deployment stage, it is expected that it will be very costly. One way of rendering the implementation of quantum networks more cost-effective is to take advantage of preexisting infrastructure by using previously deployed optical fiber networks [27]. With this in mind, we used real-life fiber data of the Netherlands. This was made available to us by SURF, a classical network provider for Dutch education and research institutions. We considered a repeater chain with nodes in Delft, The Hague, Leiden and Amsterdam, as depicted in Figure 5.7, as this is an example of a possible near-term quantum network in the Netherlands. We use real fiber length and attenuation in our simulations. We chose Delft and Amsterdam as the end nodes of the chain as out of these four cities they are the most distant pair. The baseline values used for computing the value of the cost function for each set of parameters were obtained from actual state-of-the-art experimental results using NV centers. The process through which we converted these experimental results to our abstract model parameters is described in detail in Appendix 5.12.1. We set as performance targets end-to-end fidelity $F_{min} = 0.7$ and end-to-end entanglement generation rate $R_{min} = 1$ Hz. The value of F_{min} was chosen to ensure that we remain in the regime where the agreement between the abstract model and the detailed NV model is good (see Appendix 5.10 for details). Besides this practical argument, there is no strong reason to pick a particular number for the fidelity or the rate. These numbers are simply examples, meant to show how our methodology can find the minimal hardware requirements satisfying them.




Figure 5.7: Visualization of the quantum network we will consider. The end nodes, represented by circles, are placed in Delft and Amsterdam. The repeater nodes, represented by squares, are placed in The Hague and Leiden. The placement of the nodes roughly approximates their actual geographical location and the length of the fibers connecting them is included for reference.

In order to study the effects of internode distance, chain length and number of repeaters we further applied our methodology to chains of equally spaced nodes with varying numbers of repeaters. In one case, we kept the internode distance fixed, and in the other we kept the total length fixed as we varied the number of repeaters. More concretely, we considered (i) a chain of equally spaced nodes spanning 800 km and (ii) a chain with

an internode distance of 100 km. For each of these, we considered the cases of 3, 5, 10 and 12 repeater nodes. The baseline parameter values are computed in the same manner as in the previous use case, so we again defer to Appendix 5.12.1 for details. We also consider the same target performance metrics as in the previous use case.

5.4.1 Results

Real network

We will now show the main results obtained by applying our methodology to the network introduced in Figure 5.7.

In Figure 5.8 we show the best and average values of the total cost function (Equation (5.6)) as a function of the optimization step. Contrasting with Figure 5.5, we see that (i) the average value of the cost function remains significantly higher than the best value and that (ii) the best value per generation oscillates. The first observation is explained by the combination of the inherent randomness of the GA and the fact that we used step functions for the cost. A GA generates new candidate solutions through a process of mutation and recombination, as detailed in Appendix 5.9. While these processes allow for a thorough exploration of the parameter space, they may also produce solutions that fall outside the defined target metrics. The step functions in the cost ensure that such solutions will be heavily penalized, explaining the high average values of the cost function in Figure 5.8. The second observation is also explained by a combination of two factors, namely the already mentioned step functions in the cost and the non-deterministic nature of our simulations. Since across different simulations for the same set of parameters there are fluctuations in the values of the end-to-end metrics, it might happen that these sometimes dip below the predefined targets. Due to the step function, the cost associated to this particular set of parameters will become much higher, meaning that it will no longer be the best solution. This effect can be minimized by running our simulations multiple times for each set of parameters, as discussed in Section 5.2.6.

In Table 5.1 we show the parameters of the best solution found using our methodology. For comparison purposes, we also show the baseline values we considered. The biggest

	F_{EL}	p_{suc}	s_q	T_1	T_2
Baseline	0.9698	0.004600	0.8590	10 h	4.9 ms
Solution	0.9806	0.09770	0.9414	10.23 h	22.79 ms

Table 5.1: Experimentally-derived baseline parameter values and values of the best solution found using our methodology for the use case discussed in Section 5.4.1. The biggest relative increases happen for T_2 and p_{suc} , suggesting that improving these parameters is key for achieving scalable NV-based repeaters.

relative increases are in p_{suc} and T_2 , suggesting that induced dephasing noise is the biggest hurdle in the way of NV-based repeater technology. On the other end of the spectrum, the solution's T_1 value is barely higher than that of the baseline, indicating that T_1 coherence times in NV centers are already long enough.

Equally-spaced nodes

We now show the main results obtained by applying our methodology to repeater chains of equally spaced nodes with different numbers of repeaters. To study how the overall




Figure 5.8: Evolution of the (a) best and (b) average values of the total cost function (shown on the y axes) for the use case discussed in Section 5.4.1. The best value converges to roughly 13, whereas the average oscillates around 50000 (note the different scales). The high values of the average cost throughout the optimization process are due to the mutation process, which sometimes produces solutions that do not fulfill the target metrics. Our simulation was run 100 times for each individual.

length of a chain and the internode distance affect the solutions found, we considered two cases: (a) fixed chain length (FCL) and (b) fixed internode distance (FID). For both FCL and FID we applied our methodology to chains of 3, 5, 10 and 12 repeater nodes. We note that each data point in the plots shown in this section corresponds to the best solution found after 200 generations, with 150 population individuals per generation and 100 simulation runs per individual. Running our optimization procedure once with these parameters takes roughly 46 hours locally using consumer-market hardware (Intel Core i7-8665U and 8 GB RAM), underlining the need for access to HPC centers. In fact, by using such a center, the computation time can be reduced to 2 hours (using 2 nodes of the HPC center, each endowed with 64 GB of memory and 24 cores with CPU E5-2690). We note that the vast majority of this time is taken by quantum repeater simulations, with the time needed by the GA being negligible in comparison. We further note that, as shown in Figure 9 of [8], the runtime of repeater chain simulations using NetSquid grows linearly with the number of nodes in the chain. This implies that our method remains applicable for chains that are significantly longer than the ones considered in this chapter.

In Figure 5.9 we show how the total cost of the best solution found varies with the number of repeaters in both cases. We observe a linear growth of the FID cost with the number of repeaters, which is not surprising: fixing the internode distance but increasing the number of repeater nodes corresponds to increasing the total length covered. In fact, the leftmost data point in Figure 5.9 corresponds to a chain spanning 400 km, whereas the rightmost is associated to a chain spanning 1300 km. We would expect connecting end nodes that are further apart to be a greater challenge due to the exponential growth in photon losses, which necessitates repeater parameters of higher quality. This does not

apply to the FCL use case. All the data points in the associated curve correspond to a repeater chain that spans 800 km and we observe in Figure 5.9 that the cost is slightly higher for the 3-repeater setup. It was not *a priori* obvious that this would be the case. A smaller number of repeaters implies that the swap quality and fidelity of the elementary link do not need to be as good, as there will be fewer swaps and hence less fidelity loss. On the other hand the elementary links are longer than in a setup with many repeaters, so the associated baseline values are worse (see Appendix 5.12.1 for details). Any improvement then requires a higher parameter cost, as per Equation 5.5.




Figure 5.9: Total cost, as defined in Equation 5.6, of the best solutions found by our GA for setups with varying number of repeaters. The cost grows linearly for FID. There is no discernible pattern for FCL. Each data point corresponds to the best solution found after 200 generations, with 150 population individuals per generation and 100 simulation runs per individual.

To further explore how the solutions found vary, we plot in Figure 5.10 the end-to-end fidelity and entanglement generation rate of these solutions against the number of repeaters in the chain. We see that, for both use cases and all numbers of repeaters, the end-to-end fidelity is very close to 0.7. On the other hand, the rate decreases from around 80 Hz to 30 Hz as the number of repeater nodes increases from 3 to 12 at FID and it increases slightly from 40 Hz to 50 Hz with the number of repeaters at FCL. While the fidelities obtained are what we expected, since the limit we imposed via the cost function was 0.7, the same is not true for the rates. The penalty term we added to the cost function only comes into effect if the rate drops below 1 Hz, so there is no benefit in terms of the cost to have a solution that results in a rate of e.g. 50 Hz versus one of 1 Hz. We would thus expect the best solutions to have rates close to 1 Hz, which was not the case.

In order to explain this, we note that T_2 and p_{suc} are inextricably linked. T_2 reflects the intensity of the induced dephasing effect, (see Appendix 5.3) with a higher value of T_2 corresponding to a weaker induced dephasing effect, and vice-versa. This type of noise




Figure 5.10: Comparison of the metrics characterizing the best solutions found by our GA for each of the different setups. The end-to-end fidelity is very close to the goal of 0.7 we defined, for both FID and FCL. On the other hand, the end-to-end entanglement rate is well above the 1 Hz goal for both cases. For fixed internode distance, it decreases from roughly 80 Hz in the 3 repeater node setup to about 30 Hz in the 12 repeater node setup. For fixed chain length, it increases slightly from 40 Hz in the 3 repeater node setup to 50 Hz in the 12 repeaters setup. Each data point corresponds to 100 runs of the simulation. The error bars are smaller than the markers.

is applied every time entanglement generation is attempted. Therefore, its intensity heavily depends on p_{suc} : a lower success probability implies more entanglement generation attempts and thus more dephasing. One would naively think that the GA would always converge towards a solution with lower rate (R) up until the limit of 1 Hz we defined, as that would allow for lower values of p_{suc} and hence a lower value of the parameter cost. However, due to the connection between p_{suc} and T_2 , a lower value of the former necessitates a higher value of the latter. This then implies that solutions whose R is closer to the established requirement of 1 Hz, with their lower values of p_{suc} , might actually have higher costs than solutions with higher R , accounting for why the ideal solutions have such high rates.

To conclude our analysis of the solutions found with our optimization procedure, we present in Figure 5.11 the values of each of the parameters in the solutions found for each setup. Starting with the top row, we note that the relative variations of T_1 for different setups are small when compared to the ones of T_2 . Similarly to what we saw in the use case of Section 5.4.1, this indicates that T_1 is not a crucial parameter to improve for NV center-based repeaters. We note also that for FID, T_2 grows with the number of repeaters, whereas it remains roughly constant for FCL. This is again explained by the fact that in the first case the total distance covered increases with the number of repeaters, so one expects that longer coherence times will be required. Regarding p_{suc} , we observe that it tends to be higher for chains with more repeaters, reflecting the fact that in order to achieve similar end-to-end rates across longer chains, one cannot afford to spend as much time generating elementary links as in shorter chains.

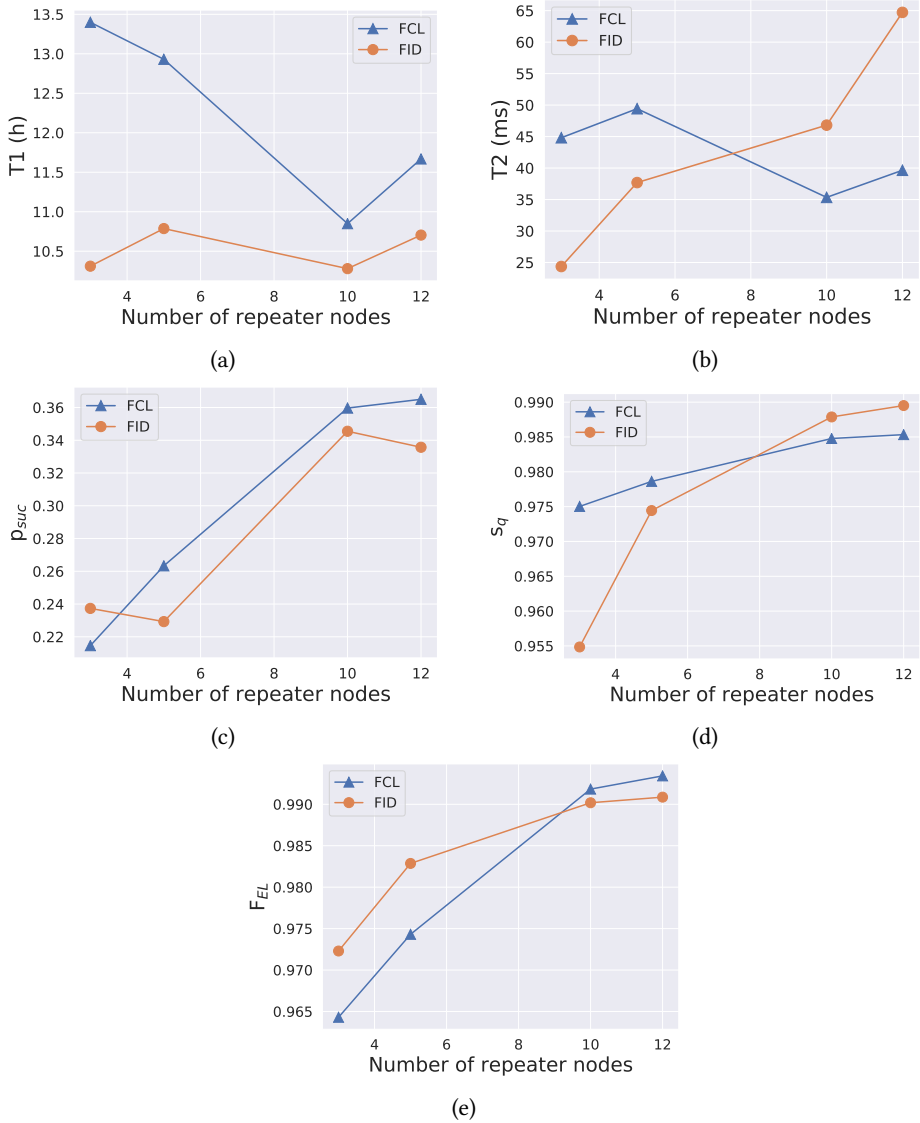


Figure 5.11: Parameters of the best solutions found for FCL and FID with different numbers of repeaters. Each data point corresponds to the best solution found after 200 generations, with 150 population individuals per generation and 100 simulation runs per individual. For a detailed discussion of these results, see the text in Section 5.4.1.

We move now to the bottom row, whose plots concern F_{EL} and s_q . Both increase with the number of repeaters, approaching 1. This was to be expected, as a higher number of repeaters implies more entanglement swaps and hence more decay in fidelity. Therefore, to reach the same end-to-end fidelity one needs better elementary links and swaps. We

further note that for few repeaters, F_{EL} is higher and s_q is lower at FID than at FCL. The opposite is true for many repeaters. We believe this may be explained by the length of the elementary links in the FCL case. For few repeaters, the FCL elementary links are longer than the FID elementary links (133 – 200 km vs 100 km), with the situation being reversed for many repeaters (73 – 89 km vs 100 km). A longer elementary link translates into a worse baseline value of F_{EL} , as detailed in Appendix 5.12.1, and thus more expensive improvements. On the other hand, the baseline value of s_q is the same irrespective of the elementary link length, and thus so is the cost of improving it. Therefore, for few repeaters the less costly solution at FCL has a lower elementary link fidelity and higher swap quality than the the less costly solution at FID. The opposite is true for many repeaters, explaining the observed behaviour.

5.5 Conclusions

We have introduced a methodology for the optimization of entanglement generation and distribution in repeater chains using GAs. In contrast with previous work in this area [1–5], our methodology is systematic, modular and broadly applicable. We validated it by benchmarking our GAs on functions commonly used for this purpose and by applying it to a repeater chain generating Werner states. We can derive analytical results for such a chain and thus gauge how well our methodology performs. Having validated our methodology, we applied it to three use cases. First, we considered a repeater chain built using real-life fiber data, thus demonstrating that our methodology can go beyond simple network topologies. The other two use cases consisted of chains of equally spaced nodes for which we varied the number of repeaters. In one we kept the internode distance constant, and in the other we fixed the total chain length. By applying our methodology to these use cases we found what are the worst parameters achieving end-to-end fidelity and rate of at least 0.7 and 1 Hz, respectively, in different scenarios. Even though this was the question we focused on answering in this chapter, we must note that our methodology is more general and can be applied to a variety of problems, given that they can be restated as optimization problems and that an appropriate cost function is designed.

On a similar note, we must again stress that even though we have here focused on a simplified five-parameter repeater model, in no way is our methodology restricted to such a model. In fact, one interesting application of our methodology would be to consider a more realistic hardware model, such as the one proposed in [28] for NV-center based repeaters. Such models are described by a very large number of parameters, on the order of 30 in this case, which means that the initial search space is too large for a direct application of our methodology. To practically apply our methodology to such a large parameter space, one could opt for a two-stage optimization process. The first stage would be similar to what was shown in this chapter, i.e. applying the methodology to a simpler model that can be mapped to the more accurate one. This step would allow us to both reduce the search space by finding minimal requirements on parameters and to identify which of these parameters have a bigger impact on the target metrics. With this knowledge in hand, we could apply the methodology to a select subset of parameters in the more detailed model, performing the optimization procedure in a reduced, more feasible search space. The outcome of this two-step procedure would then be a realistic picture of what kind of hardware improvements are required to achieve long-range entanglement, constituting a

useful guide for experimental groups working on repeater technology. This establishes the methodology we have proposed as an invaluable tool for the development of a blueprint for the quantum internet.

We have in fact applied this methodology to various other scenarios. In Chapter 6 we investigate minimal hardware requirements to perform a simple form of blind quantum computing for a single-repeater setup on a real-world fiber grid with more accurate hardware models of color centers and trapped ions. In Chapter 7, we study chains of up to seven repeaters on a real-world fiber grid, investigating hardware requirements for blind quantum computing and quantum key distribution. Finally, in Chapter 8 we extend the results shown here by performing hardware and protocol optimization in parallel. In particular, we optimize over protocols for entanglement generation, purification and swapping policies.

5.6 Data availability

The data presented in this work have been made available at <https://doi.org/10.4121/21294714.v1> [29].

5

5.7 Code availability

The code that was used to perform the simulations and generate the plots in this paper has been made available at <https://gitlab.com/FranciscoHS/NetSquid-SimplifiedRepChain> [30].

5.8 smart-stopos

In Figure 5.12, we present a detailed overview of the *smart-stopos* workflow. The user must provide a script, entitled `program.py` in Figure 5.12, that runs the simulation and an `input_file.ini` that contains information about the optimization procedure, such as the number of iterations and parameter specifications. Given these inputs, *smart-stopos* generates sets of parameters for which the simulation will be run according to the specifications given in `input_file.ini`. The outputs of the simulation are then used to generate a new set of parameters for the next iteration. This generation is done in an algorithm-dependent way. We used GAs in this work but in principle any other algorithm could be plugged in, provided that it can be run with only simulation inputs and outputs.

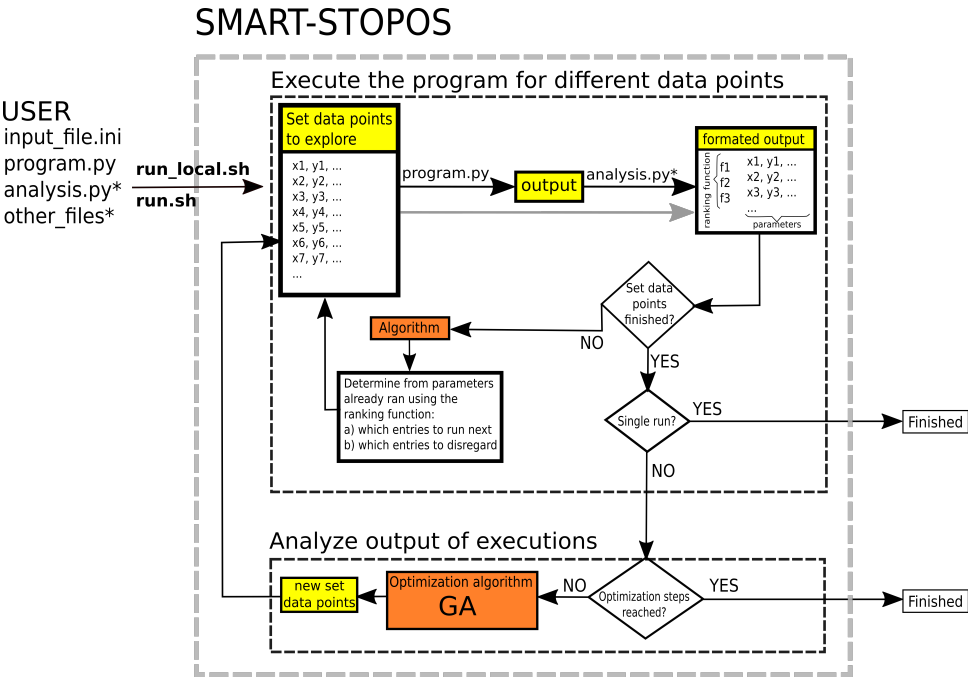


Figure 5.12: Diagram of the *smart-stopos* workflow for parameter optimization. `input_file.ini` is used to define optimization parameters such as algorithm to be used, parameters to optimize and allowed range of values. The execution of a simulation (`program.py`) and the optional post-processing of resulting data (`analysis.py`) can be done locally (`run_local.sh`) or using HPC facilities (`run.sh`). The `stopos` [20] job manager tool is required for running on HPC facilities. Files marked with * are optional.

5.9 Genetic algorithms

In this appendix we give a detailed view of the GA implementation we used for the simulations described in this work.

We started by transforming all parameters to be in the $[0, 1]$ range. This is trivial for the elementary link fidelity, success probability and swap quality. For T_1 and T_2 , which usually live in the $[0, \infty]$ range, we performed the following transformation:

$$T' = \begin{cases} \frac{1}{T+1} & \text{if } T > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{5.9}$$

which results in $T' \in [0, 1]$, as required. A chromosome, i.e. a set of parameters constituting a candidate solution, is thus a set of 5 real numbers in the $[0, 1]$ interval.

We used populations of 150 individuals, as the literature suggests that numbers of this order of magnitude are enough to get adequate parameter space exploration while still being computationally feasible [19].

After the cost function is computed for all members of the population, we select 10 of them, 20% of the total population, according to the Roulette Wheel method [19]. Again, the literature indicates that the percentage of selected individuals should be of this order

of magnitude and we empirically verified that this value produced the best results for our particular use case. One of the major challenges in GA-based optimization is to balance exploration of the search space with exploitation of known minima. If the algorithm performs selection in a purely random manner, it is no different than random search. On the other hand, if it simply selects the best individuals in a given generation, the population will tend to get stuck in local minima and be vulnerable to premature convergence. The Roulette Wheel selection method is a well-known approach to this problem, balancing exploration and exploitation by assigning selection probabilities to individuals biased, but not completely determined, by their fitness value. Applying this method to a maximization problem, the probability p_i of individual i being selected is given by:

$$p_i = \frac{f_i}{\sum_j f_j}, \quad (5.10)$$

with f_j being the value of the fitness function for individual j . The probability of selection is then proportional to how big of a share of the total fitness the individual's fitness represents, i.e. how good it is in comparison to its peers. Our problem is, however, one of minimization, not maximization. Therefore, we adapted this method by simply inverting the values of the fitness function.

Crossover is subsequently applied on the 10 selected members of the population, known as parents. This is done by randomly choosing two of the parents, sampling a crossover point, and mixing the two accordingly. To give a concrete example, if the chromosomes of the two parents are given by $[a_1, a_2, a_3, a_4, a_5]$ and $[b_1, b_2, b_3, b_4, b_5]$ and the crossover point was 2, the resulting child would have chromosome $[a_1, a_2, b_3, b_4, b_5]$. The number of children generated in this way is given by the crossover parameter, a hyperparameter of the algorithm defining how often crossover happens, times the desired population size.

The parents plus the children resulting from the crossover process are then mutated. In this process, all chromosomes of a given member of the population are randomly changed by some value that keeps them inside their range. For the mutation probability of a given parent, we implemented the adaptive scheme introduced in [31], which was shown to reduce the likelihood of corrupting a high-quality solution and enhance the exploratory properties of the algorithm. In this scheme, the probability of parent k being mutated is given by:

$$p_m = \begin{cases} 0.5 & \text{if } c_k > \bar{c} \\ 0.5 \frac{c_k - c_{\min}}{\bar{c} - c_{\min}} & \text{otherwise,} \end{cases} \quad (5.11)$$

where c_k is the value of the cost function for parent k , \bar{c} is the value of the cost function averaged over the previous generation's population and c_{\min} is its minimum value. For the children generated in the crossover process, for which there is no cost value yet, the mutation probability is a hyperparameter of the algorithm. Previous work suggests that a high crossover parameter and low mutation probability produce good results [19], so we used a crossover parameter of 0.7 and a mutation probability of 0.02 to obtain the results showed in this work.

Since generation of new individuals is to some extent probabilistic, the size of a generation can vary. To keep our population size fixed, we either randomly remove elements or add some of the best members of the previous generation. We also implement a form

of elitism, meaning that the best element of the previous generation is always preserved in the following generation, in order to prevent the algorithm wasting time searching for solutions it has already found [32].

We have empirically determined that 200 generations are usually enough to achieve satisfying solutions while still being computationally feasible on a cluster.

5.10 Abstract model validation

In this appendix we show how we validated the abstract model against a physically-accurate NV model.

5.10.1 Matching to nitrogen-vacancy center model

In order to ensure that the simulations of the abstract model can contribute to our understanding of actual physical implementations of quantum repeaters, we must verify that this abstract model captures the relevant physics to a reasonable extent. To do so, we will compare the results of simulations of a repeater chain in the abstract model with those of a repeater chain running a physically accurate model. For this purpose, any model of a physical system being studied as a possible platform for quantum repeaters would do. We will thus focus on one such system, namely NV centers, modelled as described in [33]. This is a very detailed model that accurately captures the physics of NV centers, including for instance modelling the photon emission, capture and detection processes as well as differentiating between communication and memory qubits, with all the restrictions that entails. In contrast, the simplified model we consider abstracts away all of the subtleties of photon emission and detection into an overarching success probability and treats all qubits as equal. Another key difference is that in the NV model the parameters are not mutually independent e.g. there is a relation of inverse proportionality between the fidelity of the generated entangled states and the rate at which they are generated due to the fact that both of these parameters depend on the bright state population. On the other hand, in the abstract model we make the simplifying assumption that all parameters are independent from one another. We must however emphasize that this does not reflect a limitation of our method. Taking the constraints arising from interparameter dependence into account would be possible, but we chose not to consider any such constraints in this preliminary study.

More concretely, we will perform the validation of the abstract model by taking a set of parameters describing an NV center in the model, converting it to the five parameter set that defines our model, running both simulations, and checking how the end-to-end fidelity and entanglement generation rate compare.

We start by proposing a mapping from the NV model in [33] to the five-parameter abstract model we introduced in 5.2.3. We assume that elementary link states generated in the abstract repeater chain are of the form:

$$|\phi\rangle\langle\phi| = F_{EL}|\psi\rangle\langle\psi| + (1 - F_{EL})|\uparrow\uparrow\rangle\langle\uparrow\uparrow|, \quad (5.12)$$

where $|\psi\rangle\langle\psi|$ is the ideal Bell state, F_{EL} is the elementary link fidelity and $|\uparrow\uparrow\rangle\langle\uparrow\uparrow|$ is given by:

$$|\uparrow\uparrow\rangle\langle\uparrow\uparrow| = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The overlap between $|\psi\rangle$ and $|\uparrow\uparrow\rangle$ is 0, so F_{EL} is in fact the elementary link fidelity, the sole parameter defining elementary link states. To map states from one model to another we compute the fidelity of the NV state described in the appendix of [33] and use the result to define the abstract model state as in Equation (5.12). The probability of successfully generating these elementary links is obtained in an identical manner.

We take into account any errors that might occur in an entanglement swap, which include gate errors, measurement errors and initialization errors by modelling them all as depolarizing channels, with parameters $\{p_i\}$ and multiplying them to obtain a single parameter, s_q , as shown in Equation (5.13).

$$s_q = \prod_i (1 - p_i) \quad (5.13)$$

$1 - s_q$ is then used to parameterize a depolarizing channel that is applied after an ideal Bell state measurement. The action of this channel Φ on a given state ρ as a function of s_q is given by

$$\Phi(\rho, s_q) = \left(\frac{1 + 3s_q}{4}\right)\rho + \frac{1 - s_q}{4}(X\rho X + Y\rho Y + Z\rho Z). \quad (5.14)$$

This implies that s_q is a measure of the quality of an entanglement swap, and it is thus named swap quality.

The two remaining parameters in the abstract model are T_1 and T_2 . An NV center's qubits can be either electrons, used as communication qubits, or carbons, used as memory qubits, each of them having different coherence times. This subtlety is lost when going to the abstract model, in which all qubits are created equal. We expect that decoherence will be more relevant in the memory qubits than in the communication qubits, so we ignore it for the latter. Besides this, one of the major sources of noise in NV centers is induced dephasing, the dephasing applied to the memory qubits whenever the communication qubit attempts to generate entanglement [34]. This noise source can also be accurately modelled by a T_1 , T_2 noise model. In such a model, one applies dephasing noise with probability given by

$$p = \frac{1 - e^{-t(1/T_2 - 1/2T_1)}}{2}, \quad (5.15)$$

with t being the relevant time period. This is formalized by means of a dephasing channel Φ_d whose action on a given state ρ is given by

$$\Phi_d(\rho, p) = (1 - p)\rho + pZ\rho Z. \quad (5.16)$$

On the other hand, the noise introduced in a NV center's carbon atoms over n entanglement generation attempts can be modelled by a dephasing noise process of probability

$$p_n = \frac{1 + (2(1 - p_1) - 1)^n}{2}, \quad (5.17)$$

with p_1 being the probability of a single attempt inducing dephasing noise, which can be experimentally determined [34]. If we assume that a node is always trying to generate entanglement through its electron, we can write n as a function of time:

$$n = \frac{t}{T_{cycle}}, \quad (5.18)$$

with T_{cycle} being the time it takes the NV to go through one entanglement generation attempt.

Matching the probability in Equation (5.15) to the one in Equation (5.17) and solving for T_2 , we find:

$$T_2 = \frac{1}{1/2T_1 - \log(1 - 2p_1)/T_{cycle}}. \quad (5.19)$$

This allows us to account for the effect of induced dephasing in our simulations by modelling it as a T_2 noise process. We note that, in order to more closely capture induced dephasing, this noise should only be simulated when nodes are attempting entanglement generation.

In summary, we have two important sources of noise that can be modelled by T_1 , T_2 processes: induced dephasing and memory decoherence. Since we want to restrict our model to 5 parameters, we must restrict ourselves to account for one of the two. In order to make an informed decision regarding which noise source to model, we run repeater chain simulations using the abstract model and the NV model introduced in [33]. For simplicity, we ignore distillation and consider a SWAP-ASAP protocol where the nodes can only attempt entanglement generation, wait or perform an entanglement swap. In order to obtain a better agreement between the entanglement generation rates of both models, we impose that nodes in the abstract model simulation can only generate entanglement with one neighbour at a time, as is the case for NV centers.

5

5.10.2 Comparison of nitrogen-vacancy center and abstract models

We will look into how the internode distance affects the metrics we are interested in, namely end-to-end fidelity and entanglement generation rate, in the two models. To do so, we will focus on chains of equally spaced nodes, for which varying the internode distance is equivalent to varying the total length of the repeater chain.

In Figure 5.13, we plot the end-to-end fidelity of the states generated by a chain of five equally spaced nodes as the chain's length is varied in the NV model and in both abstract model mappings. The three fidelity curves are very similar for shorter chains, roughly overlapping in chains of up to 200 km. At this point the curve for the NV model starts to diverge, dropping abruptly.

Overall, the difference between the two mappings is small. They both show very good agreement at short chain lengths, and they both perform poorly as the distances grows. This indicates that for longer distances or, alternatively, for lower fidelities, ignoring either of the noise sources results in poor agreement with the NV model. In this work we will focus on scenarios where the obtained fidelities are high, above 0.7, so that the agreement is good. We will consider the induced dephasing mapping.

We turn our attentions now to the other metric of interest, the end-to-end entanglement generation rate. In Figure 5.14, we plot the end-to-end rate against the total chain




Figure 5.13: Variation of the end-to-end fidelity of states generated by a chain of five equally spaced nodes as the chain's length is varied in the NV model (green, triangles) and in both abstract model mappings: memory decoherence (blue, circles) and induced dephasing (red, inverted triangles). The curves overlap for short chains, but as the internode distance grows the fidelity of the NV chain falls faster. The results of the two mappings are virtually identical. The error bars are smaller than the markers.

5

length for the same setup in both models. The behaviour of the two curves is similar, although the rates in the abstract model are significantly higher. We believe that this is due to the fact that, since NV centers have only one communication qubit, they must swap established entanglement from it to a memory qubit as soon as it is generated. This doesn't happen in the abstract model, and thus there is no time spent on swapping the entangled states around qubits, allowing for a higher entanglement generation rate. The difference between the two curves becomes smaller as the distance increases, which could be explained by the fact that at long distances, the majority of the time is spent on generating elementary links, as success probabilities become low. The duration of local node operations become negligible in comparison, and the time taken by internal swaps is not as important in this regime.

In order to verify this, we reran the NV simulation with the internal swap being performed instantly. The results are shown in Figure 5.15. The curves overlap over all distances the simulation covered, corroborating our hypothesis.

We conclude that the entanglement generation rates attained by the two models are similar across the board, with the biggest difference, which happens at short internode distances, being a factor of roughly 1.8. At longer distances, the rates are the same up to statistical fluctuations.




Figure 5.14: Variation of the end-to-end entanglement generation rate in a chain of five equally spaced nodes as the chain's length is varied in the NV model (triangles) and in the abstract model (circles). At short lengths, the rates achieved are higher in the abstract model by a factor of almost 2. As the distance increases, the two curves overlap.




Figure 5.15: Variation of the end-to-end entanglement generation rate in a chain of five equally spaced nodes as the chain's length is varied in the NV model (triangles) and in the abstract model (circles) with instantaneous SWAP gates. The two curves are very close for all simulated chain lengths. The error bars are smaller than the markers.

5.11 Werner chains

In this Appendix we give details about our approach for validating our GA-based optimization approach by applying it to a repeater chain generating Werner states.

The crux of this validation procedure is that we are able to find the optimum value of the cost function by a method other than the GA-based one we proposed. In order to do so, we require closed-form expressions for end-to-end fidelity and entanglement generation rate as functions of the input parameters, elementary link fidelity, success probability and swap quality.

Consider first, for simplicity, a three-node chain. The nodes establish elementary links whose states are of the form

$$\rho(x) = x|\psi^+\rangle\langle\psi^+| + (1-x)\frac{\mathbb{I}}{4}, \quad (5.20)$$

where $|\psi^+\rangle = 1/\sqrt{2}(|01\rangle + |10\rangle)$ is the ideal Bell state and \mathbb{I} is the identity. x is the Werner parameter and is related to the fidelity f of the Werner state with the ideal Bell state by $f = (1+3x)/4$. Performing an ideal BSM on two of these states, both of parameter x , results in a Werner state of parameter x^2 , i.e. the post-BSM state ρ_{BSM} is given by

$$\rho_{BSM} = x^2|\psi^+\rangle\langle\psi^+| + (1-x^2)\frac{\mathbb{I}}{4}. \quad (5.21)$$

To simulate a noisy BSM, we then apply noise via two single-qubit depolarizing channels, one on each of the two qubits involved in the BSM. Both of these channels are parametrized by the swap quality s_q , as defined in Equation (5.14). The resulting Werner state has fidelity F with the ideal Bell state:

$$F(f, s_q) = \frac{1}{4} + s_q \left(\frac{1}{2} + \frac{s_q}{4} \right) \left(\frac{4f-1}{3} \right)^2. \quad (5.22)$$

Iterating this process, one arrives at the following expression for the end-to-end fidelity

$$F(N, f, s_q) = \frac{1}{4} + s_q^N \left(\frac{1}{2} + \frac{s_q^N}{4} \right) \left(\frac{4f-1}{3} \right)^{N+1}, \quad (5.23)$$

where N is the number of repeater nodes in the chain. As a sanity check, we ran simulations of a 10-node chain for a fixed f while varying s_q and compared the obtained end-to-end fidelity with the values obtained with Equation (5.23). These results are shown in Figure 5.16.

An attentive reader might notice that Equation (5.23) slightly differs from the well-known result first derived in [35] in how it accounts for the effect of imperfect operations in the end-to-end fidelity. This is due to the fact that we have here parametrized the depolarizing noise in a slightly different manner, through two single-qubit channels.

We shift now our focus to the computation of the end-to-end entanglement generation rate across a 3-node repeater chain. We note that this quantity is simply the inverse of the waiting time, which we denote by T . Let us start with the generation of elementary links. Since we model elementary link generation attempts as processes succeeding with a fixed




Figure 5.16: End-to-end fidelity of states generated by a chain of ten nodes as the swap quality is varied, for an elementary link fidelity of 0.99. The analytical and simulation curves perfectly overlap.

probability p_{suc} , T_0 is a discrete random variable following a geometric distribution. Its expected value is then given by:

$$\mathbb{E}(T_0) = \frac{1}{p_{suc}} T_{cycle}, \quad (5.24)$$

where \mathbb{E} denotes the expected value and T_{cycle} is the cycle time, i.e. the time a single entanglement generation attempt takes. We consider a sequential repeater chain, i.e. one in which nodes can only attempt entanglement generation with one of their neighbours at a time. Therefore, the end-to-end waiting time is given by:

$$\mathbb{E}(T) = 2\mathbb{E}(T_0) + T_{SWAP}, \quad (5.25)$$

where T_{SWAP} is the time an entanglement swap takes. This holds because the repeater node has to generate elementary links with both its neighbours, and it can only start generating the second once it has finished generating the first. Furthermore, after having generated these links, it must swap them. We then define the entanglement rate R as the inverse of the expected waiting time:

$$R = \frac{1}{\mathbb{E}(T)}. \quad (5.26)$$

With Equations (5.23) and (5.26) in hand, we can compute the end-to-end fidelity and entanglement generation rate using only the input parameters f , s_q and p_{suc} and the simulation parameters T_{cycle} and T_{SWAP} . This implies that we can also directly compute the

cost function, as we have analytical expressions for every term appearing in the cost function defined Equation (5.6). We then used the Basin-Hopping algorithm [36] to find the global minimum of this cost function for a target fidelity $f_{min} = 0.6$ and a target entanglement generation rate $r_{min} = 1$ Hz over a chain of equally spaced nodes. We took as baseline values $f_b = s_{qb} = 0.5$ and $p_{suc_b} = 10^{-10}$. The Basin-Hopping algorithm is available in the SciPy library.

5.12 Computing baseline values in the abstract model

5.12.1 Uniform spacing

In order to use a realistic and up to date set of baseline values, we considered the latest results achieved in Ronald Hanson's lab at QuTech, in Delft [37]. The values for T_1 and T_2 can be directly computed from experimental values. The same is true for s_q , which can be derived from entanglement swap experiments. This does not hold for the elementary link-related parameters, namely the fidelity F_{EL} and success probability p_{suc} . Their values are heavily distance-dependent, and to date entanglement generation experiments using NV centers have only been realized at distances on the single kilometer scale [23]. We therefore use instead the model proposed in [33] with the experimental values we obtained from the Hanson group as inputs to compute the baseline values for F_{EL} and p_{suc} for the elementary link lengths we consider. In Table 5.2 we list the values used as inputs to the NV model to compute the baseline abstract parameter values. Explaining the physical meaning of each of these parameters would require a detailed exposition of the NV model, which is beyond the scope of this work. This can instead be found in [28, 33]. We note that although these parameter values have all been measured in actual laboratory experiments, they are not absolute truths. Different setups might achieve slightly different performances, and even in the same NV center not all nuclear spins are identical nor do they couple in exactly the same way to the electron spin. These nonetheless provide a valuable picture of the current state of the art.

5

Parameter	Value
visibility	0.90
σ phase drift	0.35 rad
$p_{\text{double excitation}}$	0.06
$p_{\text{electron measure error}}$	0.025
$p_{\text{electron 1 qubit error}}$	0.
$F_{\text{carbon Z rot}}$	0.999
F_{EC}	0.97
T_1 carbon	10 h
p_{det}	0.00013
$p_{\text{dark count}}$	2.5×10^{-6}
$N_{1/e}$	1400
$p_{\text{loss length}}$	0.5 dB/km

Table 5.2: Values considered for NV model parameters. See e.g. [26, 28, 33] for detailed explanations of parameters.

The bright state population α is also a required parameter in the model. We chose not to include it in Table 5.2 as this parameter is not defined by the quality of the hardware but can instead be chosen. It represents the fraction of the NV electron spin that is in the bright state, i.e. the state that emits photons. It therefore has a direct effect on the success probability of establishing elementary links, as a bigger α results in a higher photon emission probability. On the other hand, increasing α also increases the fraction of terms orthogonal to the Bell basis in the entangled state, decreasing the elementary link fidelity. There is thus a trade-off between elementary link fidelity and success probability when varying an NV center's bright state population [33]. However, in our simplified abstract model we ignore any correlations between parameters, so such a trade-off is not present. We therefore chose to ignore the existence of the trade-off in NV centers when computing the baseline value. Our process for computing these values consisted of performing a parameter scan over α with the NV model and choosing the highest achievable elementary link fidelity and success probability. In practice, this means that the baseline values considered for the elementary link fidelity were obtained with very low values of α and, conversely, the baseline values of the elementary link success probability were obtained with the highest values of α . We note that we restricted the parameter scan to the $[0, 0.5]$ interval, because for $\alpha > 0.5$ entanglement is impossible even for perfect parameters.

Taking all of this into account, we show in Table 5.3 the baseline values we obtained for the abstract model parameters. The distances in the table correspond to the elementary link lengths we considered in the two uniform spacing use cases.

	73 km	89 km	100 km	133km	200 km
T_1	10 h				
T_2	4.9 ms				
s_q	0.8459				
F_{EL}	0.95	0.94	0.90	0.80	0.52
p_{suc}	1.3×10^{-4}	7.0×10^{-5}	1.5×10^{-5}	2.2×10^{-6}	9.6×10^{-8}

Table 5.3: Baseline values of the abstract model parameters for the different elementary link lengths considered.

5.12.2 Real network

The way we arrive at the baseline values used in this use case is identical to what was described in the previous section, with the exception of F_{EL} and p_{suc} . We will now explain why these values must be computed in a different manner, as well as the process we employed to do so.

In order to arrive at realistic baseline values for the network we introduced in Figure 5.7 we used real-life fiber data that was made available to us by SURF. Although we cannot share this data, we used both the physical length of the fibers connecting the locations indicated in the Figure 5.7 and their measured attenuation values. These two quantities then have an impact on the baseline values we consider for F_{EL} and p_{suc} , resulting in three different sets of baseline values, one for each of the links in the network. This raises some questions about how the value of the cost function introduced in Equation (5.5) should be computed, as this function takes as input only one set of baseline values and a respective

set of improved values. There are multiple ways to address this. We will now explain the approach we took.

We start by computing four sets of baseline values: one for each of the links in the network plus one at negligible fiber length. By this we mean that the length we use as an input to the model in [33] is such that the impact of losses in the fiber are negligible. The cost associated with a given set of parameters is computed with respect to the set of baseline values at negligible fiber length. One can then think of this set of parameters as the improved parameters at negligible fiber length. In order to obtain the sets of parameters that will be used in our simulation we start by obtaining the improvement factor, defined in Equation (5.4), for each of the parameters. These improvement factors are then applied to the baseline values of each of the links according to Equation (5.1). The resulting three sets of values, one for each of the links, are finally the ones fed into our simulation. We reiterate that this process only applies to F_{EL} and p_{suc} . The baseline values of the remaining parameters, not being dependent on fiber length, are computed in the same way as described in the previous section. In Table 5.4 we present the baseline values we arrived at through the aforementioned process.

5

	p_{suc}	F_{EL}
DH	0.002588	0.9683
HL	0.0009187	0.9643
LA	0.0009082	0.9642
NL	0.004600	0.9698

Table 5.4: Baseline values for the links (DH stands for Delft - The Hague, HL for The Hague - Leiden and LA for Leiden - Amsterdam) and at negligible fiber length (NL).

5.13 Search space reduction using previous runs

We can use previous optimization runs to limit the search space of new runs and hence increase the probability of a good solution being found. As an example of how this can be done, suppose we have performed an optimization run over a repeater chain of 5 uniformly spaced nodes spanning some distance L . This resulted in a solution that achieves an end-to-end entanglement generation rate of $R = 1$ Hz with an elementary link success probability of p_{suc_5} , the subscript being here used to denote the number of nodes in the chain. Say we now want to apply our optimization method to a chain of 7 uniformly spaced nodes spanning the same distance L . As more elementary links need to be established and more entanglement swaps need to be performed, we know with certainty that, in order to achieve the same R a higher elementary link success probability will be needed, i.e. $p_{suc_7} > p_{suc_5}$. We can thus impose a lower bound of p_{suc_5} on the search space, reducing it.

These considerations are easy to make for the case of the elementary link success probability. Since we hold the operation times constant and implement no cut-off, it is the only parameter influencing the end-to-end entanglement generation rate. The same is not true for the other metric of interest, the end-to-end fidelity. As a concrete example, assume that the best solution found for a repeater chain of 5 uniformly spaced nodes had

an elementary link fidelity $F_{EL} = 0.96$ and a swap quality $s_q = 0.98$, resulting in an end-to-end fidelity of 0.75. One could be inclined to, in a future optimization run, upper bound the search space of F_{EL} by 0.96 to help lead the algorithm to a solution with an end-to-end fidelity closer to the target value of 0.7. However, it might be that there is a solution with $F_{EL} > 0.96$ and $s_q < 0.98$ that results in a lower cost function value than any solution with $F_{EL} < 0.96$. Therefore, by imposing this upper bound we could be preventing the algorithm from ever finding the ideal solution.

References

- [1] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, *Machine learning for long-distance quantum communication*, PRX Quantum **1**, 010301 (2020).
- [2] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, *Optimal architectures for long distance quantum communication*, Scientific reports **6**, 20463 (2016).
- [3] L. Jiang, J. M. Taylor, N. Khaneja, and M. D. Lukin, *Optimal approach to quantum communication using dynamic programming*, Proceedings of the National Academy of Sciences **104**, 17291 (2007).
- [4] S. Santra, L. Jiang, and V. S. Malinovsky, *Quantum repeater architecture with hierarchically optimized memory buffer times*, Quantum Science and Technology **4**, 025010 (2019).
- [5] K. Goodenough, D. Elkouss, and S. Wehner, *Optimising repeater schemes for the quantum internet*, arXiv preprint arXiv:2006.12221 (2020).
- [6] S. Krastanov, V. V. Albert, and L. Jiang, *Optimized entanglement purification*, Quantum **3**, 123 (2019).
- [7] *Netsquid*, <https://netsquid.org>.
- [8] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. Oliveira, M. Papendrecht, J. Rabbie, F. Rozpędek, M. Skrzypczyk, L. Wubben, W. de Jong, D. Podareanu, A. T. Knoop, D. Elkouss, and S. Wehner, *Netsquid, a discrete-event simulation platform for quantum networks*, (2020), arXiv:2010.12535 [quant-ph].
- [9] A. Torres-Knoop, T. Coopmans, D. Maier, and F. Silva, *smart-stopos*, <https://github.com/aritoka/smart-stopos> (2020).
- [10] R. F. Werner, *Quantum states with einstein-podolsky-rosen correlations admitting a hidden-variable model*, Physical Review A **40**, 4277 (1989).
- [11] N. Gunantara, *A review of multi-objective optimization: Methods and its applications*, Cogent Engineering **5**, 1502242 (2018).
- [12] J. D. Schaffer, *Some experiments in machine learning using vector evaluated genetic algorithms (artificial intelligence, optimization, adaptation, pattern recognition)*, (1986).

- [13] P. A. Vikhar, *Evolutionary algorithms: A critical review and its future prospects*, in *2016 International conference on global trends in signal processing, information computing and communication (ICGTSPICC)* (IEEE, 2016) pp. 261–265.
- [14] J. H. Holland *et al.*, *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence* (MIT press, 1992).
- [15] H.-G. Beyer and H.-P. Schwefel, *Evolution strategies—a comprehensive introduction*, *Natural computing* **1**, 3 (2002).
- [16] R. Storn and K. Price, *Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces*, *Journal of global optimization* **11**, 341 (1997).
- [17] J. Kennedy and R. Eberhart, *Particle swarm optimization*, in *Proceedings of ICNN'95-International Conference on Neural Networks*, Vol. 4 (IEEE, 1995) pp. 1942–1948.
- [18] Y. Shi and R. Eberhart, *A modified particle swarm optimizer*, in *1998 IEEE international conference on evolutionary computation proceedings. IEEE world congress on computational intelligence (Cat. No. 98TH8360)* (IEEE, 1998) pp. 69–73.
- [19] D. E. Goldberg, *Genetic algorithms* (Pearson Education India, 2006).
- [20] stopos, <https://gitlab.com/surfsara/stopos> (2019).
- [21] J. G. Digalakis and K. G. Margaritis, *On benchmarking functions for genetic algorithms*, *International journal of computer mathematics* **77**, 481 (2001).
- [22] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, 1st ed. (Addison-Wesley Longman Publishing Co., Inc., USA, 1989).
- [23] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. Vermeulen, R. N. Schouten, C. Abellán, *et al.*, *Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres*, *Nature* **526**, 682 (2015).
- [24] P. C. Humphreys, N. Kalb, J. P. Morits, R. N. Schouten, R. F. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network*, *Nature* **558**, 268 (2018).
- [25] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A ten-qubit solid-state spin register with quantum memory up to one minute*, *Physical Review X* **9**, 031045 (2019).
- [26] F. Rozpędek, R. Yehia, K. Goodenough, M. Ruf, P. C. Humphreys, R. Hanson, S. Wehner, and D. Elkouss, *Near-term quantum-repeater experiments with nitrogen-vacancy centers: Overcoming the limitations of direct transmission*, *Physical Review A* **99**, 052330 (2019).
- [27] J. Rabbie, K. Chakraborty, G. Avis, and S. Wehner, *Designing quantum networks using preexisting infrastructure*, arXiv preprint arXiv:2005.14715 (2020).

- [28] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpędek, M. Pompili, A. Stolk, P. Pawelczak, R. Knegjens, J. de Oliveira Filho, *et al.*, *A link layer protocol for quantum networks*, in *Proceedings of the ACM Special Interest Group on Data Communication* (2019) pp. 159–173.
- [29] F. Ferreira da Silva, A. Torres-Knoop, T. Coopmans, D. Maier, and S. Wehner, *Replication data for: Optimizing entanglement generation and distribution using genetic algorithms*, <https://doi.org/10.4121/21294714.v1> (2023).
- [30] F. Ferreira da Silva, A. Torres-Knoop, T. Coopmans, D. Maier, and S. Wehner, *Simulation code for optimizing entanglement generation and distribution using genetic algorithms*, <https://gitlab.com/FranciscoHS/NetSquid-SimplifiedRepChain>.
- [31] M. Srinivas and L. M. Patnaik, *Adaptive probabilities of crossover and mutation in genetic algorithms*, *IEEE Transactions on Systems, Man, and Cybernetics* **24**, 656 (1994).
- [32] S. Luke, *Essentials of Metaheuristics*, 2nd ed. (Lulu, 2013) available for free at <http://cs.gmu.edu/~sean/book/metaheuristics/>.
- [33] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356**, 928 (2017).
- [34] N. Kalb, P. C. Humphreys, J. Slim, and R. Hanson, *Dephasing mechanisms of diamond-based nuclear-spin memories for quantum networks*, *Physical Review A* **97**, 062330 (2018).
- [35] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller, *Quantum repeaters based on entanglement purification*, *Physical Review A* **59**, 169 (1999).
- [36] D. J. Wales and J. P. Doye, *Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms*, *The Journal of Physical Chemistry A* **101**, 5111 (1997).
- [37] S. Hermans, private communication (2020).

6

Requirements for a processing-node quantum repeater on a real-world fiber grid

Guus Avis¹, Francisco Ferreira da Silva¹, Tim Coopmans, Axel Dahlberg, Hana Jirovská, David Maier, Julian Rabbie, Ariana Torres-Knoop and Stephanie Wehner.

6

We numerically study the distribution of entanglement between the Dutch cities of Delft and Eindhoven realized with a processing-node quantum repeater and determine minimal hardware requirements for verifiable blind quantum computation using color centers and trapped ions. Our results are obtained considering restrictions imposed by a real-world fiber grid and using detailed hardware-specific models. By comparing our results to those we would obtain in idealized settings we show that simplifications lead to a distorted picture of hardware demands, particularly on memory coherence and photon collection. We develop general machinery suitable for studying arbitrary processing-node repeater chains using NetSquid, a discrete-event simulator for quantum networks. This enables us to include time-dependent noise models and simulate repeater protocols with cut-offs, including the required classical control communication. We find minimal hardware requirements by solving an optimization problem using genetic algorithms on a high-performance-computing cluster. Our work provides guidance for further experimental progress, and showcases limitations of studying quantum-repeater requirements in idealized situations.

¹These authors contributed equally.

This chapter is based on the publication npj Quantum Inf 9, 100 (2023).

Part of the challenge in building quantum repeaters is that their hardware requirements remain largely unknown. Extensive studies have been conducted to estimate such requirements both analytically (see, e.g., [1–24]), as well as using numerical simulations (see, e.g., [25–33]). While greatly informative in helping us understand minimal hardware requirements needed to bridge long distances, they have mostly been conducted in idealized settings where all repeaters are equally spaced, and one assumes a uniform loss of typically 0.2 dB/km on each fiber segment (exceptions are [20–22]). Furthermore, with few exceptions [8, 21, 22, 24, 31], such studies only provide rough approximations of time-dependent noise, and do not take into account platform-specific physical effects such as noise on the memory qubits during entanglement generation on NV centers [34] or collective Gaussian dephasing in ion traps (see Figure 6.1).

6.1 Results

In this chapter, we present the first study that takes into account time-dependent noise, platform-specific noise sources and classical control communication, as well as constraints imposed by a real-world fiber network, and optimizes over parameters of the repeater protocols used to generate entanglement. Our investigation is conducted using fiber data from SURF, an organization that provides connectivity to educational institutions in the Netherlands. Specifically, we will consider a network path connecting the Dutch cities of Delft and Eindhoven, separated by 226.5 km of optical fiber (see Figure 6.1 (a)). In placing equipment, we restrict ourselves to SURF locations, which leads to the repeater being located closer to Delft than to Eindhoven. Intermediary stations used for heralded entanglement generation (see Figure 6.1 (b)) cannot be placed equidistantly from both nodes either, as is generally assumed in idealized studies. We emphasize that we restrict ourselves to existing infrastructure, and therefore do not investigate the possibility of altering the fiber links. Related work which focuses on determining hardware requirements while taking into account how many repeaters to use and their placement is presented in Chapter 7.

We consider the case where the network path is used to support an advanced quantum application, namely Verifiable Blind Quantum Computation (VBQC) [35], with a client located in Eindhoven and a powerful quantum-computing server located in Delft. We chose VBQC because since their introduction blind-quantum-computing protocols have attracted a lot of interest, being widely cited as one of the principal future applications of quantum networks (see, e.g., [35–42]). While it is true that VBQC is somewhat unique in that it is highly asymmetrical in terms of the resources it requires from client and server, it is representative for many other quantum-networking applications in that it requires multiple live qubits. Additionally, the noise resilience of the specific VBQC protocol we consider [35] makes it particularly suitable to study the performance of such applications in the presence of hardware imperfections. Specifically, we consider the smallest instance of VBQC, where two entangled pairs are generated between the client and the server. Such entanglement is used to send qubits from the client to the server. We show in Section 6.7 that this can be done through remote state preparation [43]. To set the requirements of our quantum-network path, we impose that its hardware must be good enough to execute VBQC with the largest acceptable error rate [35]. This demand can be translated to requirements on the fidelity and rate at which entanglement is produced. Both depend on the lifetime of the server’s memory, as the server needs to be able to wait until both qubit

states have been generated before it can begin processing. Additionally, the requirements on the fidelity and rate can also be understood as the fidelity and rate at which we can deterministically teleport unknown data qubits between the client and the server. Therefore, while our investigation focuses on VBQC, our results can also be interpreted from the perspective of quantum teleportation.

In our study, we obtain the following results, described in more detail below: First, we investigate the *minimal hardware requirements* that are needed to realize target fidelities and rates that allow executing VBQC using our network path. These correspond to the minimal improvements over state-of-the-art hardware parameters that enable meeting the targets. Specifically, we consider parameters measured for networked color centers (specifically, for NV centers in diamond) [44–51] and ion traps [52–59]. We find that considerable improvements are needed even to bridge relatively modest distances, with our study also shining light on which parameters require significantly more improvement than others. To obtain this result, we have built an extensive simulation framework on top of the discrete event simulator NetSquid [60], which includes models of color centers (specifically adapted from NV centers in diamond), ion traps, a general abstract model applicable to all processing nodes, as well as different schemes of entanglement generation. Our framework can be readily re-configured to study other network paths of this form, including the ability to configure other types of processing-node hardware, or entanglement-generation schemes. Being able to simulate the Delft-Eindhoven path, we then perform parameter optimization based on genetic algorithms to search for parameter improvements that minimize a cost function (see Section 6.3 for details) on SURF’s high-performance-computing cluster Snellius.

Second, we examine the *absolute minimal requirements* for all parameters in our models (for color centers and ion traps), if all other parameters are set to their perfect value (except for photon loss in fiber). We observe that the minimal hardware requirements impose higher demands on each individual parameter than the absolute minimal requirements. This highlights potential dangers in trying to maximize individual parameters without taking into account global requirement trade-offs. However, somewhat surprisingly, we find that the absolute minimal requirements are typically of the same order of magnitude as the minimal requirements, and can therefore still be valuable as a first-order approximation. Our results are obtained using the same NetSquid simulation, by incrementally increasing the value of a parameter until the target requirements are met.

Finally, we investigate whether the idealized network paths usually employed in the repeater literature would lead to significantly different minimal hardware improvements. Specifically, in such idealized setups all repeaters and heralding stations are equally spaced, all fibers are taken to have 0.2 dB/km attenuation, and the models employed for the processing-node hardware are largely platform-agnostic. We find that considering real-world network topologies such as the SURF grid imposes significantly more stringent demands.

Let us now be more precise about the setup of our network path, as well as the requirements imposed by VBQC:

6.1.1 Quantum-network path

The network path we consider consists of three processing nodes that are assumed to all have the same hardware. That is, the stated hardware requirements are sufficient for

all nodes and we do not differentiate between the three nodes. On an abstract level, all processing nodes have at least one so-called communication qubit, which can be used to generate entanglement with a photon. The repeater node in the middle (Nieuwegein, Figure 6.1 (a)) has two qubits available (at least one of which is a communication qubit) that it can use to simultaneously hold entanglement with the node in Delft, as well as the one in Eindhoven. Once entanglement has been generated with both Delft and Eindhoven, the repeater node may perform an entanglement swap [61] in order to create end-to-end entanglement between Delft and Eindhoven (see Figure 6.2). On processing nodes, such a swap can be realized deterministically, i.e., with success probability 1, since it can be implemented using quantum gates and measurements on the processor. We note that even when the gates and measurements are noisy the swap remains deterministic, although it will induce noise on the resulting entangled state.

For all types of processing nodes, we here assume the repeater to act sequentially [21] due to hardware restrictions. That is, it can only generate entanglement with one of the other two nodes at a time. To minimize the memory requirements at the repeater node (Nieuwegein), we will always first produce entanglement with the farthest node (Eindhoven). Once this entanglement has been produced, the repeater generates entanglement with the closest node (Delft). To combat the effect of memory decoherence, entangled qubits are discarded after a cut-off time [21]. This means that if entanglement between Delft and Nieuwegein is not produced within a specific time window following the successful generation of entanglement between Nieuwegein and Eindhoven, all entanglement is discarded and we restart the protocol by regenerating entanglement between Nieuwegein and Eindhoven. Classical communication is used to initiate entanglement generation between nodes and notify all nodes when swaps or discards are performed.

We consider three types of processing nodes (see Figure 6.1 (c) and (d)): (1) color centers, specifically modeled on NV centers in diamond, (2) ion traps and (3) a general abstract model applicable to all processing nodes. Let us now provide more specific details on each of these models required for the parameter analysis below.

(1) NV centers are a prominent example of color centers for which significant data is available from quantum-networking experiments [44–49]. Here, the color center’s optically-active electronic spin is employed as a communication qubit. The second qubit is given by the long-lived spin state of a Carbon-13 atom, which is coupled to the communication qubit and used as a memory qubit. Our color-center model accounts for the following:

- Restricted topology, with one optically-active communication qubit and one memory qubit (note however that larger registers have been realized, for example in [51]);
- Restricted gate set, with arbitrary rotations on the communication qubit, Z-rotations on the memory qubit and a controlled rotation gate between the two qubits;
- Depolarizing noise in all gates, bit-flip noise in measurement;
- Qubit decoherence in memory modeled through amplitude damping and dephasing channels with decay times T_1 and T_2 (we consider the experimentally-realized times of $T_1 = 1$ hour (10 hours) and $T_2 = 0.5$ s (1 s) for the communication (memory) qubit [49–51]);

- Induced dephasing noise on the memory qubit whenever entanglement generation using the communication qubit is attempted [34, 48].

The efficiency of the photonic interface in NV centers is limited to 3% due to the zero-phonon line (ZPL). It is likely that executing VBQC using the path we investigate will require overall photon detection probabilities higher than 3%. Little data is presently available for other color centers (SiV, SnV). We hence focus on the NV model, but do allow a higher emission probability, which could be achieved either by using a color center with a more favorable ZPL (65-90% for SiV [62], 57% for SnV [62]), or by placing the NV in a cavity [63]. More details about our color-center model, and a validation of the model against experimental data for NV centers, can be found in Section 6.6.

(2) Trapped ions are charged atoms suspended in an electromagnetic trap, the energy levels of which can be used as qubits. Our trapped-ion model accounts for the following:

- Two identical, optically active ions in a trap;
- Restricted gate set as described in [64], with arbitrary single-qubit Z rotations, arbitrary collective rotations around axes in the XY plane, and an entangling Mølmer-Sørensen gate [65];
- Depolarizing noise in all gates, bit-flip noise in measurement;
- Qubit decoherence modeled as collective Gaussian dephasing, with a characteristic coherence time [31];
- Off-resonant scattering that adds a random delay to the emission time of photons, which is counteracted using a tunable coincidence time window (as captured by a toy model introduced in Section 6.9).

More details about our trapped-ion model, and a validation of the model against experimental data, can be found in Section 6.6.

(3) We further investigate an abstract, platform-agnostic processing-node model. This model accounts for depolarizing noise in all gates and in photon emission, as well as amplitude-damping and phase-damping noise in the memory. It does not account for any platform-specific restrictions on topology, gate set or noise sources. Later on, we show that using the abstract model instead of hardware-specific models leads to an inaccurate picture of minimal hardware requirements. Even so, the abstract model can be valuable to study systems for which hardware-specific models are as of yet unavailable. Additionally, we note that the smaller number of hardware parameters in the abstract model as compared to the hardware-specific models means that the parameter space can be explored more efficiently, making it easier to, e.g., find minimal hardware parameters.

To entangle two processing nodes, one can use different schemes for entanglement generation, and we here consider the so-called single-click [66] and double-click schemes [67]. Both of these start with two distant nodes generating matter-photon entanglement and sending the photon to a heralding station. In the single (double)-click protocol, matter-matter entanglement is heralded by the detection of one (two) photons after interference. The trapped-ion nodes we investigate perform only double-click entanglement generation as single-click entanglement generation has not been realized for the type of trapped-ion

devices we consider, i.e., trapped ions in a cavity. The color-center nodes and abstract nodes perform both single and double click. Our entanglement-generation models account for the following physical effects:

- Emission of the photon in the correct mode, modeled through a loss channel;
- Imperfect photon emission modeled through a depolarizing channel;
- Capture of the photon into the fiber, modeled through a loss channel;
- Photon frequency conversion, modeled through a loss channel (as a first-order approximation, we assume this is a noiseless process);
- Photon attenuation in fiber, modeled through a loss channel;
- Photon delay in fiber;
- Photon detection at the detector, modeled through a loss channel;
- Detector dark counts;
- Photon arrival at the detector at different times;
- Imperfect photon indistinguishability.

6

While photon attenuation losses depend on the characteristics (such as the length) of the fiber that is used to deploy a quantum network, the other losses depend only on the quantum hardware that is used. For convenience we collect all the hardware-related losses into a single parameter, called the photon detection probability excluding attenuation losses.

The hardware parameters used in our models are based on quantum-networking experiments with NV centers (single-click [46–49] and double-click [44, 45]), and trapped ions (double-click [53]).

6.1.2 Blind quantum computation

Having discussed our modeling of the path between Delft and Eindhoven, we turn to the end nodes.

Both end nodes are processing nodes. The end node in Eindhoven takes the role of client in the VBQC protocol. In Delft, there is not only an end node, but also a powerful quantum-computing server. After entanglement is established by the end node in Delft it transfers its half of the entangled state to this server. The client in Eindhoven simply measures its half of the entangled state. The Delft scenario is similar to the setting investigated in [68], where the authors consider an architecture in which a node contains two NV centers, one of them used for networking and the other for computing. Here, we make some simplifying assumptions that allow us to focus on the network path: we take the state transfer process to be instantaneous and noiseless, and assume that the computing node is always available to receive the state. Further, we assume that the quantum gates performed by the server are noiseless and instantaneous, and that their qubits are subject to depolarizing noise with memory coherence time $T = 100$ s. Because of these assumptions, the requirements we find are limited primarily by imperfections in the network path itself rather than in the computing node.

We investigate hardware requirements on three processing nodes (two end nodes and one repeater node) so that a client in Eindhoven can perform 2-qubit VBQC, a particular case of the protocol described in [35], using the Delft server. In this protocol, the client prepares qubits at the server, which are then used to perform either computation or test rounds. In test rounds, the results of the computation returned by the server are compared to expected results. The protocol is only robust to noise if the noise does not cause too large an error rate. The protocol is shown in [35] to remain correct if the *maximal* probability of error in a test round can be upper-bounded by 25%. We prove in Section 6.7 that the protocol is still correct if the *average* probability of error in a test round can be upper-bounded by 25%. We further prove in the same section that if the entangled pairs distributed by the network path can be used to perform quantum teleportation at a given rate and quality, the protocol can be executed successfully. Namely, this is true if the average fidelity at which unknown pure quantum states can be teleported using the entangled pairs distributed by the network path (F_{tel}) and the entangling rate R satisfy a specific bound. We note that this bound takes into account potential jitter in the delivery of entanglement (i.e., the fact that the time required to generate entanglement, and hence the time entangled states need to be stored in memory, can fluctuate around its expected value). We consider two distinct pairs of F_{tel} and R that satisfy this bound as our target metrics, namely:

- Target 1: $F_{\text{tel}} = 0.8717$, $R = 0.1$ Hz,
- Target 2: $F_{\text{tel}} = 0.8571$, $R = 0.5$ Hz.

The choice of these specific values was motivated by the fact that there is no fidelity $F_{\text{tel}} \leq 1$ for $R \approx 0.014$ Hz such that the VBQC condition is satisfied, therefore all target rates should satisfy $R > 0.014$ Hz, preferably with some margin to avoid trivial solutions. Additionally, Target 1 is achievable using either the single-click or double-click protocol and using either one or zero repeaters on the fiber path under consideration, given sufficient hardware improvements. In contrast, Target 2 is achievable only using the single-click protocol and one repeater (see also Sections 6.14.3 and 6.14.4). This suggests that the difference between the two targets is large enough to lead to significantly different results.

The derivation of this bound assumes that the client prepares qubits at the server by first generating them locally and then transmitting them to the server using quantum teleportation. We note that alternatively the remote-state-preparation protocol [43] can be used, which will likely be more feasible in a real experiment as it requires fewer quantum operations by the client. In Section 6.7 we describe a way how the VBQC protocol [35] can be performed using remote state preparation. Note however that we have not investigated the security of the protocol in this case. We show that under the assumption that local operations are noiseless, quantum teleportation and remote state preparation lead to the exact same requirements on the network path. Thus, in case the target is met, VBQC can be successfully executed using either quantum teleportation or remote state preparation. Lastly, we note that there is a linear relation between the average teleportation fidelity F_{tel} and the fidelity of the entangled pair [69].

6.1.3 Minimal hardware requirements

Here, we aim to find the smallest improvements over current hardware to generate entanglement enabling VBQC. These are shown at the bottom of Figure 6.3 for color centers (left) and trapped ions (right). In the table at the top of Figure 6.3 we show a selection of the actual values for the minimal hardware requirements (the set of parameters representing the smallest improvement over state-of-the-art parameters, see Section 6.3 for details on how we determine this), as well as the absolute minimal requirements (the minimal value for each parameter assuming that every other parameter except for photon loss in fiber is perfect). All the parameters are explained in Section 6.3, and their state-of-the-art values that we consider are given in Table 6.1.

The minimal color-center hardware requirements for Target 1 (blue line in Figure 6.3, bottom left) correspond to the usage of a double-click protocol, as we found that this allows for laxer requirements than using a single-click protocol. On the other hand, the minimal requirements for Target 2 (orange line in Figure 6.3, bottom left) correspond to the usage of a single-click entanglement-generation protocol. This is because achieving Target 2 in the setup we studied is not possible at all with a double-click protocol even if every parameter except for photon loss in fiber is perfect. Therefore, and since we do not model single-click entanglement generation with trapped ions, the bottom-right plot of Figure 6.3 depicts only the requirements for trapped ions to achieve Target 1.


We thus find that in the setup we investigated performance targets with relatively higher fidelity and lower rate are better met by using a double-click protocol. On the other hand, higher rates can only be achieved with single-click protocols. This was to be expected, as (a) states generated with single-click protocols are inherently imperfect, even with perfect hardware and (b) the entanglement-generation rate of double-click protocols scales poorly with both the distance and the detection probability due to the fact that two photons must be detected to herald success.

6.1.4 Absolute minimal requirements


We now aim to find the minimal parameter values that enable meeting the targets, if the only other imperfection were photon loss in fiber. These are the absolute minimal requirements, presented in the table at the top of Figure 6.3. We observe that while there is a gap between them and the minimal hardware requirements, it is perhaps surprisingly small. For example, the minimal photon detection probability excluding attenuation losses required to achieve Target 1 with color centers is roughly 1.5 times larger than the corresponding absolute minimal requirement. However, both requirements represent a three order of magnitude increase with respect to the state-of-the-art, which makes a factor of 1.5 seem small in comparison.

We remark on the feasibility of achieving the minimal hardware requirements for color centers. NV centers, on which we have based the state-of-the-art parameters used in this work, are the color center that has been most extensively used in quantum-networking experiments (see [62] for a review). As discussed in Section 6.1, the efficiency of the photonic interface in this system is limited to 3% due to the zero-phonon line. Both targets we investigated place an absolute minimal requirement on the photon detection probability excluding attenuation losses above this value. Improving the photonic interface of NV centers beyond the limit imposed by the zero-phonon line is only possible through integra-

tion of the NV center into a resonant cavity [63]. Alternatively, other color centers with a more efficient photonic interface could be considered as alternatives for long-distance quantum communication [62].



(a)



(b)

(c)

(d)

Figure 6.1: **(a)** Satellite photo of the Netherlands overlaid with a depiction of the hypothetical one-repeater setup connecting the Dutch cities of Delft and Eindhoven that we investigate. The white circles represent processing nodes. They are connected to each other and to heralding stations through fiber drawn in white. The black dots within the processing nodes represent qubits (the distinction between communication and memory qubits is not represented here). The placement of nodes and heralding stations is constrained by the fiber network, and their position on the figure roughly approximates their actual geographic location. All distances are given in kilometers, with a total fiber distance between Delft and Eindhoven of 226.5 km. **(b)** Heralding station. Photons emitted by a processing node travel through the optical fiber and are interfered at a beam splitter. Photon detection heralds entanglement between processing nodes. This process is affected by the overall probability that emitted photons are detected, the coincidence probability, i.e., the probability that photons arrive in the same time window, the imperfect indistinguishability of the photons as measured by the visibility and dark counts in the detector. **(c)** Color center in diamond, one of the processing nodes we investigate. We consider an optically-active electronic spin used as a communication qubit, and a carbon spin used as a memory qubit. Decoherence in both qubits is modeled through amplitude damping and phase damping channels with characteristic times T_1 and T_2 , respectively. These are different for the two qubits. The existence of an always-on interaction between the qubits allows for the execution of two-qubit gates, but also means that entangling attempts with the communication qubit induce noise on the memory qubit. **(d)** Ion trap, the other processing node we investigate. We consider two optically active ions trapped in an electromagnetic field generated by electrodes, whose energy levels are used as qubits. The ions interact through their collective motional modes, which enables the implementation of two-qubit gates. They are subject to collective Gaussian dephasing noise characterized by a coherence time.




Figure 6.2: Protocol executed in the setup we investigate. 1. No entanglement is shared *a priori*. E.N. stands for End Node, R.N. stands for Repeater Node and H.S. stands for Herald Station. 2. Entanglement generation attempts begin along the longer link, which connects the repeater node to the Eindhoven node. 3. After entanglement has been established along the longer link, attempts for entanglement generation along the shorter link start. In case this takes longer than a given cut-off time, the previously generated entanglement is discarded and we go back to 2. 4. After entanglement is generated on both links, the repeater node performs an entanglement swap, creating an end-to-end entangled state. 5. The Delft node maps its half of the state to a powerful quantum-computing server, while the Eindhoven node measures its half.

Setup		p_{det}	T_{coh} (s)	
CC	Baseline	0.00051	0.5	
	R = 0.1 Hz	Minimal hardware requirement	0.71	7.2
		Absolute minimal requirement	0.48	1.0
	R = 0.5 Hz	Minimal hardware requirement	0.88	7.5
		Absolute minimal requirement	0.22	1.0
TI	Baseline	0.12	0.085	
	R = 0.1 Hz	Minimal hardware requirement	0.96	0.67
		Absolute minimal requirement	0.59	0.42




Figure 6.3: **Top:** Parameter values required to connect the Dutch cities of Delft and Eindhoven using color-center (CC) and trapped-ion (TI) repeaters for an entanglement-generation rate of 0.1 Hz and an average teleportation fidelity of 0.8717 (Target 1) and a rate of 0.5 Hz and average teleportation fidelity of 0.8571 (Target 2). The baseline parameter values have been demonstrated in state-of-the-art experiments. The absolute minimal requirements are the required parameter values assuming that there are no other sources of noise or loss with the exception of fiber attenuation. The coherence-time values in the table are the communication-qubit dephasing time for CC and the collective dephasing time for TI (see Section 6.3 for an explanation of these parameters). The TI requirements are for running a double-click entanglement-generation protocol. The CC requirements are for running a double-click protocol for Target 1, and a single-click protocol for Target 2. We note that all the minimal requirements found have a photon detection probability excluding attenuation losses above 30%, the current state-of-the-art value for frequency conversion [56]. **Bottom:** Directions along which hardware must be improved to connect the Dutch cities of Delft and Eindhoven using a CC (**left**) and TI (**right**) repeater. The further away the line is from the center towards a given parameter, the larger improvement that parameter requires. Improvement is measured in terms of the “improvement factor”, which tends to infinity as a parameter tends to its perfect value (see Section 6.3 for the definition). In both plots a logarithmic scale is used. The origin of the plots corresponds to an improvement factor of 1, i.e., no improvement with respect to the state of the art. On the **bottom left** (CC), the blue (orange) line corresponds to the minimal requirements for Target 1 (Target 2). Improvement is depicted for the following parameters, clockwise from the top: photon detection probability excluding attenuation losses in fiber, dephasing time of the communication qubit, dephasing time of the memory qubit, noise in the two-qubit gate, visibility of photon interference and dephasing noise induced on memory qubits when entanglement generation is attempted. On the **bottom right** (TI), the line corresponds to the minimal requirements for Target 1. Improvement is depicted for the following parameters, clockwise from the top: photon detection probability excluding attenuation losses in fiber, qubit collective dephasing coherence time, spin-photon emission fidelity, visibility of photon interference and probability that two emitted photons coincide at the detection station. All parameters are explained in Section 6.3, and their state-of-the-art values that are being improved upon are given in Table 6.1.

6.2 Discussion

6.2.1 Hardware requirements in simplified settings

Since we made use of real-life fiber data and elaborate, platform-specific hardware models, the results above would be difficult to obtain analytically. For instance, collective Gaussian dephasing in ion traps could be challenging to analyze. Analytical results are however attractive, as they provide a more intuitive picture of the problem at hand. In order to find them, an approach commonly taken in the literature is to simplify the setup under study so that it becomes analytically tractable. A usual simplification is to assume what we name the *standard* scenario, in which nodes and heralding stations are equally spaced, and where the fiber attenuation is 0.2 dB/km throughout. Another common simplification is to consider simplified physical models for the nodes and the entangled states they generate (see, among others, [19, 70–72]). In order to investigate how hardware requirements change if such simplifications are used, we now apply our methodology to these two simplified situations and compare the resulting hardware requirements with the ones for our setup. We hope to understand whether considering these setups leads to similar results, indicating that the simplifying approach is a good one, or if doing so paints an unrealistic picture of the hardware requirements, which would favor our approach.

a. Effect of Existing Fiber Networks on Hardware Requirements We investigate how the hardware requirements in the standard scenario differ from the fiber-network-based setup. We thus present in Figure 6.4 a comparison of the hardware requirements for color centers in the two situations. In both cases, we consider double-click entanglement generation, targeting an entanglement-generation rate of 0.1 Hz and an average teleportation fidelity of 0.8717. Significant improvements over the state-of-the-art are required in both scenarios, but the magnitude of these improvements would be understated in case one were to consider the standard scenario and ignore existing fiber infrastructure. For example, doing so would lead to underestimating the required coherence time of the memory qubits by a factor of four. More broadly, we see that the improvement required is larger in the fiber-network scenario for (i) the photon detection probability excluding attenuation losses and (ii) memory parameters (coherence times and tolerance to entanglement-generation attempts). Both of these results can be explained by the fact that when a real-world fiber network is considered there is more attenuation and the nodes are not evenly spaced. As a consequence, better photonic interfaces are required to achieve similar rates, and states likely spend a longer time in memory, necessitating longer coherence times. This emphasizes the need for considering limitations imposed by existing fiber infrastructure when estimating requirements on repeater hardware.

b. Effect of Platform-Specific Modeling on Hardware Requirements Finally, we look into how the hardware requirements are affected if the processing nodes are modeled in a simplified, platform-agnostic way. We thus compare the hardware requirements for color-center and trapped-ion repeaters with those for a platform-agnostic abstract model for a quantum repeater. This is a simple processing-node model that accounts for generic noise sources such as memory decoherence and imperfect photon indistinguishability, but does not take platform-specific considerations such as restricted topologies into account. For more details on the platform-agnostic abstract model, see Section 6.6.7. We consider double-click entanglement generation in the fiber-network-based setup, targeting an entanglement-generation rate of 0.1 Hz and an average teleportation fidelity of






Figure 6.4: Hardware requirements for connecting the Dutch cities of Delft and Eindhoven using a color center repeater performing double-click entanglement generation on an actual fiber network (blue) and assuming the standard scenario (orange, dashed). Requirements are for achieving an entanglement-generation rate of 0.1 Hz and an average teleportation fidelity of 0.8717. Parameters shown are, from top to bottom: visibility of photon interference, dephasing noise induced on memory qubits when entanglement generation is attempted, dephasing time of communication qubit, dephasing time of memory qubit, photon detection probability excluding attenuation losses in fiber and two-qubit gate fidelity.

0.8717.

To perform the comparison, we proceed as follows: (i) map the state-of-the-art hardware parameters to abstract-model parameters, (ii) run the optimization process for the platform-specific model and the abstract model in order to find the minimal hardware requirements for both, (iii) map the obtained platform-specific hardware requirements to the abstract model and (iv) compare them to the hardware requirements obtained by running the optimization process for the abstract model. The results of this comparison can be seen in Figure 6.5. The hardware requirements are significantly different for the abstract model and for the trapped-ion and color-center models. This can be explained by the greater simplicity of the abstract model. Take coherence time as an example. The communication and memory qubits of color centers decohere at different rates, a complexity which is not present in the abstract model. Therefore, improving the coherence time in the abstract model has a bigger impact than improving a given coherence time in the color center model. This means that in the abstract model it is comparatively cheaper to achieve the same performance by improving the coherence time rather than other parameters. The fact that memory noise in trapped ions is modeled differently than in the abstract model (the trapped-ion memory noise is Gaussian, arising from a collective dephasing process. See Equations 6.7 and 6.9) could also explain the difference in the requirements for the coherence times seen in that case.



(a) Color center.



(b) Ion trap.

Figure 6.5: Comparison of hardware requirements for connecting the Dutch cities of Delft and Eindhoven using a repeater performing double-click entanglement generation considering a simple abstract model and more detailed color center (left) and ion trap (right) models. Requirements are for achieving an entanglement-generation rate of 0.1 Hz and an average teleportation fidelity of 0.8717. Parameters shown are, from top to bottom: spin-photon emission fidelity (trapped ion only), visibility of photon interference, photon detection probability excluding attenuation losses in fiber, fidelity of entanglement swap and qubit coherence time.

6.2.2 Entanglement without a repeater

We note that one of the set of targets we investigated, namely an entanglement-generation rate of 0.1 Hz and an average teleportation fidelity of 0.8717, could also be achieved in the setup we investigated without using a repeater node if a single-click entanglement-generation protocol were employed. Furthermore, the hardware improvements required would be more modest in this case than if a repeater were used. For more details on this, see Section 6.14.3.

6.2.3 Outlook

In order to design and realize real-world quantum networks, it is important to determine minimal hardware requirements in more complex scenarios such as heterogeneous networks with multiple repeaters and end nodes. The method presented in this work is well suited for this. Furthermore, it would be valuable to investigate what limitations the assumptions we have made in our modeling place on our results. For example, we did not consider the effects of fiber dispersion. These effects could hamper entanglement generation and hence affect the minimal hardware requirements. Even though preliminary investigations suggest that these effects might be small, quantifying them would represent a step forward in determining realistic minimal repeater-hardware requirements. Another interesting open question is what effect the use of entanglement-distillation protocols (see [73] for a review) would have on the minimal hardware requirements.

6

6.3 Methods

In this section we elaborate on our approach for determining the minimal and absolute minimal hardware requirements for processing-node repeaters to generate entangled states enabling VBQC.

6.3.1 Conditions on network path to enable VBQC

In our setup, a client wishes to perform 2-qubit VBQC, a particular case of the protocol described in [35], on a powerful remote server whose qubits are assumed to suffer from depolarizing noise with coherence time $T = 100$ s. We further assume that the computation itself is perfect, with the only imperfections arising from the network path used to remotely prepare the qubits. This protocol is shown in [35] to be robust to noise, remaining correct if the *maximal* probability of error in a test round can be upper-bounded by 25%. We argue in Section 6.7 that the protocol is still correct if the *average* probability of error in a test round can be upper-bounded by 25%, as long as we assume that the error probabilities are independent and identically distributed across different rounds of the protocol. This is the case for the setup studied here, as the state of the network is fully reset after entanglement swapping takes place at the repeater node. This condition, together with the assumption on the server's coherence time, can be used to derive bounds on the required average teleportation fidelity and entanglement-generation rate, as shown in Section 6.7.

6.3.2 Average teleportation fidelity

We use the average teleportation fidelity F_{tel} that can be obtained with the teleportation channel Λ_σ arising from the end-to-end entangled state σ generated by the network we

investigate as a target metric:

$$F_{\text{tel}}(\sigma) \equiv \int_{\psi} \langle \psi | \Lambda_{\sigma}(|\psi\rangle\langle\psi|) | \psi \rangle d\psi, \quad (6.1)$$

where the integral is taken over the Haar measure. See Section 6.7.1 for more details.

6.3.3 Hardware improvement for VBQC as an optimization problem

We want to find the minimal hardware requirements that achieve a given average teleportation fidelity F_{target} and entanglement-generation rate R_{target} . We restate this as a constrained optimization problem: we wish to minimize the hardware improvement, while ensuring that the performance constraints are met. These constraints are relaxed through scalarization, resulting in a single-objective problem in which we aim to minimize the sum of the hardware improvement and two penalty terms, one for the rate target and one for the teleportation fidelity target. The resulting cost function is given by

$$\begin{aligned} C = & w_1 \left(1 + (F_{\text{target}} - F_{\text{tel}})^2 \right) \Theta(F_{\text{target}} - F_{\text{tel}}) \\ & + w_2 \left(1 + (R_{\text{target}} - R)^2 \right) \Theta(R_{\text{target}} - R) \\ & + w_3 H_C(x_1, \dots, x_N), \end{aligned} \quad (6.2)$$

where H_C is the hardware cost associated with parameter set $\{x_1, \dots, x_N\}$, w_i are the weights of the objectives, Θ is the Heaviside function and F_{tel} and R are the average teleportation fidelity and entanglement-generation rate achieved by the parameter set, respectively. The hardware cost function H_C maps sets of hardware parameters to a cost that represents how large of an improvement over the state of the art the set requires. To compute this consistently across different parameters we use no-imperfection probabilities, as done in [60] (where they are called no-error probabilities). A parameter is improved by a factor k , called the *improvement factor*, if its corresponding no-imperfection probability p_{ni} becomes $\sqrt[k]{p_{\text{ni}}}$. For example, if the error probability of a gate is 40%, its probability of no-imperfection is 0.6. After improving it by a factor of 4 the no-imperfection probability becomes $\sqrt[4]{0.6} \approx 0.88$, corresponding to an error probability of approximately 12%. The hardware cost associated with a set of hardware parameters is the sum of the respective improvement factors, i.e.,

$$H_C(x_1, \dots, x_N) = \sum_{i=1}^N \frac{\ln\{p_{\text{ni}}(b_i)\}}{\ln\{p_{\text{ni}}(x_i)\}}, \quad (6.3)$$

where $p_{\text{ni}}(x_i)$ is the no-imperfection probability corresponding to the value x_i of parameter i and $p_{\text{ni}}(b_i)$ is the no-imperfection probability corresponding to the baseline value b_i of parameter i . We have here for concreteness used natural logarithms, but the hardware cost is invariant to changes in the logarithms' bases. We note that these improvement factors are the quantities shown in Figure 6.3. The weights w_i are chosen such that the first two terms are larger than the last one for near-term parameters, guaranteeing that the set of parameters minimizing C meets performance targets. We are then effectively restricted to the region of parameter space in which the performance constraints are satisfied, as all points corresponding to near-term parameters in this region have a lower cost

than points outside it. The problem then becomes one of minimizing the hardware cost in this region. We have verified that the expected values of the average teleportation fidelity and entanglement-generation rate of the parameter sets found meet the constraints, thus enabling VBQC conditional on our assumptions. Our method guarantees that the set of parameters found is 'minimal' in the sense that making any of the parameters worse would result in the target not being met. However, we note that there exist many such solutions, and if specific knowledge is available about how hard it is to improve particular parameters, the cost function could be adapted to pick out minimal parameter sets that may be easier to attain. An example of this is the efficiency of the NV center's photonic interface, which is limited to 3% due to the ZPL. Going beyond this limit requires integration into a cavity, which carries with it a host of challenges [62, 63]. One could then modify the cost function to make improving the efficiency of the photonic interface beyond 3% more expensive than improving other parameters. However, as it is challenging to accurately estimate the hardness associated with specific improvements and, furthermore, the hardness may depend on the specific expertise available within a given research group, we have refrained from making such estimates.

6.3.4 Optimization parameters

Using the methodology described later on in this section, we perform an optimization over both protocol and hardware parameters. First we enumerate the protocol parameters:

- Cut-off time, the time after which a stored qubit is discarded;
- Bright-state parameter (single-click entanglement generation only), the fraction of a matter qubit's superposition state that is optically active;
- Coincidence time window (double-click entanglement generation with ion traps only), the maximum amount of time between the detection of two photons for which a success is heralded. We model the effect of the coincidence time window using a toy model, see Section 6.9.

Second, we enumerate the hardware parameters:

- The Hong-Ou-Mandel visibility [74] is a measure for the indistinguishability of interfering photons and is defined by [75]

$$1 - \frac{C_{\min}}{C_{\max}}. \quad (6.4)$$

Here C_{\min} is the probability (coincidence count rate) that two photons that are interfered on a 50:50 beamsplitter are detected at two different detectors when the indistinguishability is optimized (as is the case when using interference to generate entanglement), while C_{\max} is the same probability when the photons are made distinguishable.

- The probability of double excitation is the probability that two photons are emitted instead of one in entanglement generation with color centers;

- The induced memory qubit noise is the dephasing suffered by the memory qubit when the communication qubit is used to attempt entanglement generation. The number given for this parameter in Table 6.1 corresponds to the number of electron spin pumping cycles after which the Bloch vector length of the memory qubit in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ in the $X - Y$ plane of the Bloch sphere has shrunk to $1/e$ when the communication qubit has bright-state parameter 0.5 [34];
- The interferometric phase uncertainty is the uncertainty in the phase acquired by the two interfering photons when they travel through the fiber in single-click entanglement generation with color centers;
- The photon detection probability excluding attenuation losses is the probability that a photon is detected given that emission was attempted, and assuming that the fiber length is negligible, i.e., considering every form of photon loss (including coupling to fiber) except the length-dependent attenuation loss in fiber;
- Every gate is parameterized by a depolarizing-channel fidelity;
- For color centers, T_1 and T_2 are the characteristic times of the time-dependent amplitude damping and phase damping channels affecting the qubits, and are different for the communication and memory qubits. The effect of the amplitude (phase) damping channel after time t is given by equation (6.5) ((6.6))

$$\begin{aligned} \rho \rightarrow & \left(|0\rangle\langle 0| + \sqrt{e^{-t/T_1}} |1\rangle\langle 1| \right) \rho \\ & \left(|0\rangle\langle 0| + \sqrt{e^{-t/T_1}} |1\rangle\langle 1| \right)^\dagger \\ & + \sqrt{1 - e^{-t/T_1}} |0\rangle\langle 1| \rho \left(\sqrt{1 - e^{-t/T_1}} |0\rangle\langle 1| \right)^\dagger \end{aligned} \quad (6.5)$$

$$\begin{aligned} \rho \rightarrow & \left(1 - \frac{1}{2} \left(1 - e^{-t/T_2} e^{-t/(2T_1)} \right) \right) \rho \\ & + \frac{1}{2} \left(1 - e^{-t/T_2} e^{-t/(2T_1)} \right) Z \rho Z; \end{aligned} \quad (6.6)$$

- For ion traps, the coherence time characterizes the time-dependent collective Gaussian dephasing process that the qubits undergo, which is given by [31]:

$$\rho \rightarrow \int_{-\infty}^{\infty} K_r \rho K_r^\dagger p(r) dr, \quad (6.7)$$

where

$$K_r = \exp\left(-ir \frac{t}{\tau} \sum_{j=1}^n Z_j\right), \quad (6.8)$$

Z_j denotes a Pauli Z acting on qubit j , n is the total number of ions in the trap, τ the coherence time and t the storage time, and

$$p(r) = \frac{1}{\sqrt{2\pi}} e^{-r^2/2}; \quad (6.9)$$

- The noise on matter-photon emission is parameterized by a depolarizing-channel fidelity (i.e., the matter-photon state directly after emission is a mixture between a maximally entangled state and a maximally mixed state);
- The dark-count probability is the probability that a detection event is registered at a detector without a photon arriving.

The state-of-the-art values we use for the hardware parameters are shown in Table 6.1. For more details on how the effects of the different hardware parameters are included in our models, see Section 6.6. We note that some of the hardware parameters we consider in fact conceal trade-offs. For example, the probability of getting a double excitation when using color centers to emit photons can to an extent be tuned. In this case, a lower probability of double excitation would come at the cost of getting fewer events. However, optimizing over all such trade-offs is beyond the scope of this work.

	Color center	Ion trap
Visibility	0.9 [49]	0.89 [53]
Probability of double excitation	0.06 [49]	-
Induced memory qubit noise (entanglement attempts until dephasing)	5300 [49]	-
Interferometric phase uncertainty (rad)	0.21 [49]	-
Photon detection probability excluding attenuation losses	5.1×10^{-4} [49]	0.111 [54–56]
Two-qubit gate fidelity	0.97 [46]	0.95 [52]
Two-qubit gate duration	500 μ s [46]	107 μ s [52]
Communication T1	1 h [50]	-
Communication T2	0.5 s [49]	-
Memory T1	10 h [51]	-
Memory T2	1 s [51]	-
Coherence time	-	85 ms [52]
Matter-photon emission fidelity	1 [76]	0.99 [77]
Matter-photon emission duration	3.8 μ s [48]	50 μ s [52, 55]
Dark count probability	1.5×10^{-7} [49]	1.4×10^{-5} [54]

Table 6.1: State-of-the-art color center and trapped-ion hardware parameters. For the trapped-ion parameters, a detection time window of 17.5 μ s and a coincidence time window of 0.5 μ s are assumed (see Section 6.6 for more details). All fidelities are depolarizing-channel fidelities. A dash (“-”) indicates that a value would not be well defined (for instance, there is no T1 or T2 time defined for trapped ions, while there is no coherence time defined for color centers). We note that not all of these parameter values have been realized in a single experiment.

6.3.5 Evaluating hardware quality

In order to minimize the cost function C , we require an efficient way of evaluating the performance attained by each parameter set. We do this through simulation of end-to-end entanglement generation using NetSquid. The full density matrix of the states generated, as well as how long their generation took in simulation time are recorded and used to compute the average teleportation fidelity and rate of entanglement generation. Since

entanglement generation is a stochastic process, multiple simulation runs are performed in order to collect representative statistics.

6.3.6 Framework for simulating quantum repeaters

In our NetSquid simulation framework, we have implemented hardware models for color centers, trapped ions and a platform-agnostic abstract model. This includes the implementation of different circuits for entanglement swapping and moving states for each platform, conditioned on their respective topologies and gate sets. Additionally, we have implemented both single and double-click entanglement-generation protocols. In order to combine these different building blocks that are required to simulate end-to-end entanglement distribution, we define services that each have a well-defined input and output but can have different implementations. For example, the entanglement-generation service can either use the single-click or double-click protocol, and entanglement swapping can be executed on either color center or trapped-ion hardware. End-to-end entanglement generation is then orchestrated using a link-layer protocol (inspired on the one proposed in [78]) that makes calls to the different services, agnostically of how the services are implemented. This allows us to use the same protocol for each different configuration of the simulation. Switching between configurations in our simulation framework then only requires editing a human-readable configuration file. The modularity of the simulation framework would make it simple to investigate further hardware platforms and protocols.

The link-layer protocol is itself an implementation of the link-layer service defined in [78]. From a user perspective, this simplifies using the simulation as all that needs to be done to generate entanglement is make a call to the well-defined link-layer service, without any knowledge of the protocol that implements the service. In this work, the link-layer protocol is the one for a single sequential repeater illustrated in Figure 6.2. However, the protocols included in our simulation code are able to simulate entanglement generation on chains of an arbitrary number of (sequential) repeaters that use classical communication to negotiate when to generate entanglement and that implement local cut-off times.

6.3.7 Finding minimal hardware improvements

In order to find the sets of parameters minimizing the cost function C , we employ the optimization methodology introduced in [79], which integrates genetic algorithms and NetSquid simulations. A genetic algorithm is an iterative optimization method, which initiates by randomly generating a population consisting of many sets of parameters, also known as individuals. These are then evaluated using the NetSquid simulation and the cost function, and a new population is bred through mutation and crossover of individuals in the previous population. The process then iterates, with better-performing individuals being more likely to propagate to further iterations. For further details on the optimization methodology employed, see Section 6.11 and [79].

This methodology is computationally intensive, so we execute it on the Snellius supercomputer. We use one node of the Snellius supercomputer, which contains 128 2.6 GHz cores and a total of 256 GiB of memory. Based on previously observed data reported in [79], we employ a population size of 150 evolving for 200 generations. The simulation is run 100 times for each set of parameters, as we have empirically determined that this constitutes a good balance between accuracy and computation time. The time required for

the procedure to conclude is hardware, protocol and parameter dependent, but we have observed that 10 wall-clock hours are typically enough. We stress that this approach is general, modular and freely available [79].

6.3.8 Finding absolute minimal hardware requirements

In order to find these requirements, which are the minimal parameter values enabling meeting the performance targets if the only other imperfection is photon loss in fiber, we perform a sweep of each parameter, starting at the state-of-the-art value and terminating when the targets are met. For each value of each parameter, we sweep also over the protocol parameters, i.e., the cut-off time, coincidence time window (for double-click entanglement generation with ion traps) and bright-state parameter (for single-click entanglement generation).

6.4 Data availability

The data presented in this work have been made available at <https://doi.org/10.4121/19746748>.

6.5 Code availability

The code that was used to perform the simulations and generate the plots in this chapter has been made available at <https://gitlab.com/softwarequitech/simulation-code-for-requirements-for-a-processing-node-quantum-repeater-on-a-real-world-fiber-grid>.

Author contributions

G.A. led the development of hardware models and the simulation of repeater protocols. F.F.S. led the development and execution of optimizations. G.A. and F.F.S. devised the target metric and proved the underlying theorems related to verifiable blind quantum computation. G.A., F.F.S., D.M., T.C., A.D., H.J. and J.R. contributed to the development of the code used in the simulations. A.T. contributed to the optimal execution of simulations on computing clusters. G.A., F.F.S. and S.W. wrote the manuscript. All authors revised the manuscript. S.W. conceived and supervised the project.

6.6 Setup

In this section, we elaborate on our modeling of the setup we study. We go over the topology of the fiber network we considered, the protocols employed by the repeater nodes, the modeling of the nodes themselves and of entanglement generation.

6.6.1 Fiber network and node placement

Deployment of quantum networks in the real world will likely make use of existent fiber infrastructure, as we have discussed in Chapter 2. In order to accurately account for this in our investigation of repeater hardware requirements, we used data of SURF's fiber network in our simulation. SURF is a network provider for education and research institutions in the Netherlands. The data we have access to consists of the physical location in which

nodes are placed, the length of the fibers connecting them, the measured attenuation of each fiber and their dispersion. We restricted the placement of quantum nodes and heralding stations to existing nodes in the network, and we assumed that they were connected by the shortest length of fiber possible. We note that in the case we studied this corresponds also to the least overall attenuation. Although dispersion was not considered in our models, an investigation of its effects would constitute an interesting extension to this work. There are four nodes in the shortest connection between Delft and Eindhoven in SURF's network, as depicted in Figure 6.6. This means we are restricted to placing a single




Figure 6.6: Satellite photo of the Netherlands overlaid with depiction of the shortest connection between the Dutch cities of Delft and Eindhoven in SURF's fiber network. The white circles represent locations where processing nodes and heralding stations can be placed, and are connected to one another through white fibers. The position of the circles in the figure roughly approximates their actual physical location. All distances are given in kilometers.

repeater between the end nodes, as a two-repeater setup would require five nodes in total, two for the repeaters and three for the heralding stations. A single-repeater setup, on the other hand, requires only three nodes, one for the repeater itself and two for heralding stations. One of the connection's nodes must therefore not be used, and there are two possible choices for how this can be done, as depicted in Figure 6.7. We applied our methodology to both of these paths and determined that the one on the left in Figure 6.7 requires a smaller improvement over current hardware. Therefore, all the results presented in Section 6.1 apply to it. For more details, see Section 6.14.2.

6.6.2 Repeater protocol

We now elaborate on the protocol executed by the nodes. We note that the repeaters we investigate are sequential, which means that they can only generate entanglement with one neighbor at a time.



6

Figure 6.7: Two possible choices for processing node (white circles, black circles within represent qubits) and heralding station placement in the SURF network's shortest connection between the Dutch cities of Delft and Eindhoven. In the path on the left, the Rotterdam node is unused, thereby directly connecting the Delft - Rotterdam and Rotterdam - Utrecht links. Similarly, the Den Bosch node is unused in the path on the right.

1. A request for end-to-end entanglement generation is placed at one of the end nodes.
2. This end node sends a classical message through the fiber to the other end node, in order to verify whether it is ready to initiate the entanglement generation protocol.
3. If that is the case, the second end node sends a confirmation message back, as well as an activation message for the repeater node.

The next step is the generation of elementary link states. We begin by generating entanglement on the Eindhoven - Nieuwegein link, which is longest, so as to minimize the time states remain in memory.

4. The neighboring Eindhoven and Nieuwegein nodes share classical messages sent through the fiber connecting them to ensure that both agree to generate entanglement.
5. Once they have established agreement, entanglement generation attempts begin and continue until success.
6. Steps 4 and 5 are repeated by the repeater and the Delft end node.
7. The repeater performs a Bell-state measurement on the two qubits it holds, thereby creating an entangled state held by the end nodes.
8. The outcome of this measurement is sent as a classical message to both end nodes.
9. The end nodes become aware that end-to-end entanglement has been established and perform the appropriate correction on the Bell state.

We also employ a cut-off protocol. If the generation of the second entangled state lasted longer than a predefined cut-off time, the first state, corresponding to the longer link, is discarded. Entanglement generation then restarts along the longer link.

Such a protocol involving sequential repeaters and a cut-off timer has been studied before, e.g., in [22]. The steps above are sufficient to generate one end-to-end entangled state. If the generation of multiple states had been requested, steps 4-9 would be repeated until enough pairs had been generated. We note that we do not simulate the application of the Bell-state correction, but instead record which correction should have been applied and handle it in post-processing.

Information on how we implemented such a protocol in a scalable and hardware-agnostic fashion can be found in Section 6.13.

6.6.3 Quantum-computing server

After the end-to-end entangled state has been generated, we assume the end node in Delft transfers its half of it to a powerful quantum-computing server. This is a similar setting as the one investigated in [68], where the authors consider an architecture in which a node contains two NV centers, one of them used for networking and the other for computing. We assume that the state transfer process is instantaneous and noiseless and that the server is always available to receive the state. Additionally, we assume that all quantum gates performed by the server are noiseless and instantaneous, and that qubits stored in the server are subject to depolarizing memory noise with a coherence time of $T = 100$ s.

6.6.4 Processing nodes

The quantum nodes we investigated are processing nodes, i.e. quantum nodes that are capable of storage and processing of quantum information. This processing is done through noisy quantum gates. The specific gate set available to the nodes depends on the particular hardware, but we model the gate noise of all of them with depolarizing channels. Measurements are also noisy, which is captured by a bit-flip channel, i.e. with some probability a $|0\rangle$ ($|1\rangle$) is read as 1 (0). Furthermore, as already mentioned, all the nodes we investigate are sequential, which means that they can only generate entanglement with one other node at a time.

We now elaborate on the details of our modeling for each of the three nodes we study.

6.6.5 Color centers

In Table 6.2, we present the baseline values of all color center hardware parameters relevant to our simulations, as well as references reporting their experimental demonstration.

Parameter	Noise	Duration/Time
Visibility	0.9 [49]	-
Probability of double excitation	0.06 [49]	-
$N_{1/e}$: Nuclear dephasing during electron initialization	5300 [49]	-
Dark count probability	1.5×10^{-7} [49]	-
σ_{phase} : Interferometric phase uncertainty (rad)	0.21 [49]	-
Photon detection probability excluding attenuation losses	5.1×10^{-4} [49]	-
Spin-photon emission	F = 1 [76]	3.8 μ s [76]
Electron readout	F=0.93(0), 0.995(1) [49]	3.7 μ s [47]
Carbon initialization	F=0.99 [51]	300 μ s [80]
Carbon Z-rotation	F=0.999 [81]	20 μ s [81]
Electron-carbon controlled X-rotation	F=0.97 [46]	500 μ s [46]
Electron initialization	F=0.995 [51]	2 μ s [82]
Electron single-qubit gate	F=0.995 [49]	5 ns [46]
Electron T1	-	1 hours [50]
Electron T2	-	0.5 s [49]
Carbon T1	-	10 hours [51]
Carbon T2	-	1 s [51]

Table 6.2: Baseline color center hardware parameters.

Color center nodes are modeled with a star topology, with the communication qubit in the middle. The memory qubits can all interact with the communication qubit, but not with

one another. The communication qubit owes its name to the fact that it is optically active, which means it can be used for light-matter entanglement generation. The spin states of the memory qubits are long-lived, so they are typically used for information storage. We model memory decoherence in color center qubits through amplitude damping and phase damping channels with T_1 and T_2 lifetimes. The effect of the amplitude (phase) damping channel after time t is given by equation (6.10) ((6.11)).

$$\begin{aligned} \rho \rightarrow & \left(|0\rangle\langle 0| + \sqrt{e^{-t/T_1}} |1\rangle\langle 1| \right) \rho \left(|0\rangle\langle 0| + \sqrt{e^{-t/T_1}} |1\rangle\langle 1| \right)^\dagger \\ & + \sqrt{1 - e^{-t/T_1}} |0\rangle\langle 1| \rho \left(\sqrt{1 - e^{-t/T_1}} |0\rangle\langle 1| \right)^\dagger \end{aligned} \quad (6.10)$$

$$\rho \rightarrow \left(1 - \frac{1}{2} \left(1 - e^{-t/T_2} e^{-t/(2T_1)} \right) \right) \rho + \frac{1}{2} \left(1 - e^{-t/T_2} e^{-t/(2T_1)} \right) Z \rho Z \quad (6.11)$$

The T_1 and T_2 lifetimes of the communication qubit are different from those of the memory qubits. An entangling gate is available in the form of a controlled X-rotation between the communication qubit and each memory qubit. Furthermore, arbitrary single-qubit rotations can be implemented on the communication qubit.

The constrained topology and gate set of the color center place some limitations on the quantum circuits to be executed. First of all, the typical Bell-state-measurement circuit must be adapted, as depicted in Figure 17 (d) of the Supplementary Information of [60]. Furthermore, since only the communication qubit can be used to generate light-matter entanglement, the repeater node must move its half of the first entangled state it generates from the communication qubit to a memory qubit in order to free it up to generate the second entangled state. The circuit for this move operation can be seen in Figure 17 (a) of the Supplementary Information of [60].

Finally, we note that it might be advantageous for the end node in Eindhoven, which generates entanglement with the repeater first and then has to wait, to map its half of the elementary link entangled state from the communication qubit to the memory qubit while it waits for the repeater node to generate entanglement with the Delft node. Note that this has nothing to do with the fact that the end node in Delft transfers its qubit to the powerful quantum-computing server after end-to-end entanglement is established. There is however a trade-off: while mapping the state means that it will be held in a qubit with a longer coherence time, it will also undergo extra decoherence due to the noise in the gates that constitute the circuit for the move operation. We have investigated this trade-off by applying our methodology to the two situations, and found that not mapping requires a smaller improvement over current hardware. Therefore, the color center results shown in the main text pertain to the situation in which the Eindhoven node does not map its half of the entangled state to a memory qubit. We must however note that this finding is specific to both the topology we study and the baseline hardware quality we consider. For more details on the comparison between mapping and not mapping, see Section 6.14.1.

The color center hardware model we employed builds on previous work [22, 78], and its NetSquid implementation has been validated against experiments [60]. This includes the model for the processor as well as for the entangled states generated through a single-click protocol. The main novelty introduced in this work regarding color center modelling is a model for the entangled states generated through the Barrett-Kok protocol [67]. This

is essentially the model introduced in Section 6.8, with the addition of induced dephasing noise. This addition accounts for the fact that every entanglement generation attempt induces dephasing noise on the color center's memory qubits [34]. We simulate this effect using a dephasing channel. The dephasing probability p , accumulated after possibly multiple entanglement generation attempts, is given by equation (6.12).

$$p = \frac{1 - (1 - 2p_{\text{single}})^k}{2}. \quad (6.12)$$

In this equation, p_{single} is the probability of dephasing after a entanglement generation attempt and k is the number of required entanglement generation attempts. In our simulations, we apply a dephasing channel of parameter p twice after entanglement has been successfully generated, to reflect the fact that each attempt requires the emission of two photons. p_{single} can be related to $N_{1/e}$, the number of electron spin pumping cycles after which the Bloch vector length of a nuclear spin in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ in the $X - Y$ plane of the Bloch sphere has shrunk to $1/e$ when the electron spin state has bright-state parameter $\alpha = 0.5$, through equation (6.13).

$$p_{\text{single}} = (1 - \alpha)(1 - e^{-1/N_{1/e}}). \quad (6.13)$$

$N_{1/e}$ can in turn be experimentally determined, and $N_{1/e} = 5300$ for state-of-the-art color center experiments [48].

6

The double-click model is the only component of our color center simulations that had not yet been compared to experimental data. With this in mind, we validated it against the experiment reported in [44]. There, the authors demonstrated heralded entanglement generation between two color centers separated by three meters using the Barrett-Kok protocol. After establishing entanglement, measurements of the two entangled qubits were performed to investigate whether the outcomes were correlated as expected. This was repeated for the X and Z bases, and for the states $|\Psi^\pm\rangle = 1/\sqrt{2}(|01\rangle \pm |10\rangle)$. We replicated this setup using our color center NetSquid model and ran the experiment 10000 times per measurement basis in order to gather relevant statistics. The results of this validation are shown in Figure 6.8. The results obtained with our simulation model broadly replicate the experimental results, although they do not lie within the statistical error bars. Overall, the simulation results are closer to the ideal case of perfect (anti-)correlation. This can be explained by the fact that our model for the double-click states is quite simple and hardware-agnostic, ignoring noise sources such as the probability of double photon emission. Further, the number of experimental data points is small, of the order of a total of 100 events for each of the plots in the figure. Nonetheless, considering the simplicity of the model, we believe that the level of agreement is satisfactory.

6.6.6 Trapped ions

In Table 6.3, we present the baseline values of all trapped-ion hardware parameters relevant to our simulations, as well as references to the articles reporting their experimental demonstration.

In this work, we present for the first time a NetSquid model for trapped-ion nodes in quantum networks. The trapped-ion nodes we model are based on the state of the art for

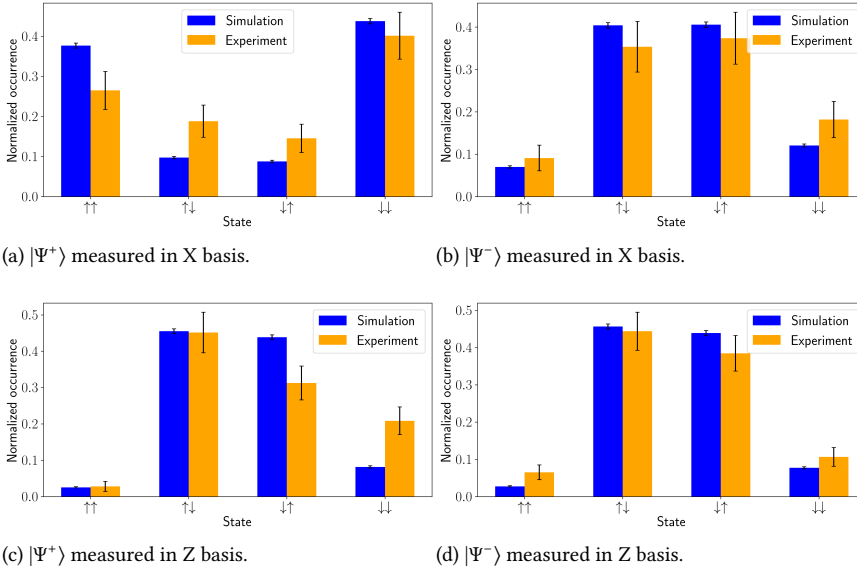


Figure 6.8: Comparison of measurement outcomes of the entangled state generated using the Barrett-Kok protocol in the experiment described in [44] and our simulation of the same scenario. The plots on the left (right) correspond to the case in which the state $|\Psi^+\rangle = 1/\sqrt{2}(|01\rangle + |10\rangle)$ ($|\Psi^-\rangle = 1/\sqrt{2}(|01\rangle - |10\rangle)$) is generated. The plots above (below) show the outcomes when measuring in the X (Z) basis. The error bars depict the standard error of the mean. Anti-correlation of the spin states is expected for every plot except for the one in the top left, for which we expect to see a correlation. The smaller dimension of the simulation error bars can be attributed to the number of executions of the protocol, which was of the order of 10000 per plot. This is two orders of magnitude more than what was performed experimentally.

trapped ions in a cavity, which consists of $^{40}\text{Ca}^+$ ions in a linear Paul trap [31, 54, 55, 64, 77, 83–89] (we note that promising results have also been achieved for trapped ions without cavities, these systems are however not considered in this work [90–94]). In our model they have all-to-all connectivity, their qubits all have the same coherence time and can all be used to generate light-matter entanglement. However, the node can only generate entanglement with one remote node at a time.

Decoherence in n trapped-ion qubits is modeled through a collective Gaussian dephasing channel that has the following effect on the n -qubit state ρ [31]:

$$\rho \rightarrow \int_{-\infty}^{\infty} K_r \rho K_r^\dagger p(r) dr, \quad (6.14)$$

where

$$K_r = \exp\left(-ir \frac{t}{\tau} \sum_{j=1}^n Z_j\right), \quad (6.15)$$

Z_j denotes a Pauli Z acting on qubit j , τ the coherence time and t the storage time, and

$$p(r) = \frac{1}{\sqrt{2\pi}} e^{-r^2/2}. \quad (6.16)$$

Parameter	Noise	Duration/Time
Visibility	0.89 [53]	-
Dark count probability	1.4×10^{-5} [54]	-
Photon detection probability excluding attenuation losses	0.111 [54–56]	-
Ion-photon emission	F = 0.99 [77]	50 μ s [52, 55]
Readout	F=0.999(0), 0.99985(1) [57]	1.5 ms [53]
Initialization	F=0.999 [58]	36 μ s [52]
Z-rotation	F=0.99 [59]	26.6 μ s [59]
Mølmer-Sørensen gate	F=0.95 [52]	107 μ s [52]
Coherence time	-	85 ms [52]

Table 6.3: Baseline trapped-ion hardware parameters. A detection time window of 17.5 μ s is assumed. For the visibility, a coincidence time window of 0.5 μ s is assumed (see Section 6.6.6 for further explanation). The photon detection probability excluding attenuation losses includes a 30% efficiency factor for quantum frequency conversion [56]. It is based on a detection efficiency of 0.43 for a 46(1) MHz drive laser and a detection time window of 17.5 μ s [55]. However, the number from [55] is based on a detector efficiency of 0.87(2) for photons at 854 nm. The detection efficiency at telecom frequency would instead be 0.75 using superconducting nanowire detectors [54], giving an additional conversion factor of 0.75/0.87. The dark count probability is based on a 0.8 Hz dark count rate for telecom superconducting nanowire detectors [54] multiplied by 17.5 μ s. The ion-photon emission fidelity has been corrected for the 1.5% infidelity due to dark counts in [77]. The initialization duration includes time for cooling sequences and repumping (3 ms of cooling for 230 photon generation attempts on one ion, with 40 μ s for repumping and optical pumping in 30 of the attempts and 20 μ s in 200 of the attempts, averaging at $\sim 36\mu$ s per attempt [59]).

6

This can be read as follows: the qubits dephase because they undergo Z-rotations at an unknown constant rate of $-2r$ per coherence time τ . This is modeled by sampling the Gaussian distribution for the dephasing rate, $p(r)$, for each ion trap each time its state is reset. The qubits are then time-evolved by applying unitary rotations in accordance with the sampled value for r . The baseline value $\tau = 85$ ms included in Table 6.3 is obtained from [52]. However, the value for the coherence time reported there is 62 ± 3 ms. The reason for this discrepancy is a difference in convention. To see this, we can evaluate equation (6.14) for $n = 1$, i.e., for a single qubit. In that case, we find

$$\rho \rightarrow \lambda\rho + (1 - \lambda)Z\rho Z, \quad (6.17)$$

where

$$\lambda = \frac{1}{2} \left(1 - e^{-2\left(\frac{t}{\tau}\right)^2} \right). \quad (6.18)$$

The single-qubit dephasing model used in [52] instead has

$$\rho \rightarrow \lambda'\rho + (1 - \lambda')Z\rho Z, \quad (6.19)$$

where

$$\lambda' = \frac{1}{2} \left(1 - e^{-\left(\frac{t}{\tau'}\right)^2} \right). \quad (6.20)$$

Here, τ' is the coherence time in their model. The models are exactly equivalent for $\tau = \sqrt{2}\tau'$. Therefore, the reported value $\tau' = 62 \pm 3$ ms corresponds to $\tau = 88 \pm 4$ ms. The value we use, $\tau = 85$ ms, represents a conservative interpretation of the result presented in [52]. Our model for the storage of quantum states in ionic qubits has been validated against experimental data from [52]. In this experiment, ion-photon entanglement is created with one ion in a two-ion device. Next, ion-photon entanglement is created with the other ion every $330\mu\text{s}$. Our simulation results are compared to the experimental results in Figure 6.9.




Figure 6.9: Validation of our trapped-ion decoherence model against an experiment in [52]. In the experiment, a trap with two ions first emits a photon entangled with the first ion, and then keeps emitting new photons entangled with the second ion every $330\mu\text{s}$. The figure shows the evolution of the fidelity to the perfect Bell state of the state shared by the first ion and the photon entangled to it as a function of time. Error bars of the simulation results represent the standard error of the mean and are sometimes hard to distinguish because of their size. The simulation has been conducted using the baseline coherence time $\tau = 85$ ms, which is the value obtained from [52]. The ion-photon emission fidelity has been set to $F = 0.97$ to tune the fidelity at time zero such that good agreement between the simulation and the experiment is obtained. All other parameters have been set to their perfect values.

The entangling gate available to the trapped-ion qubits as we model them is the Mølmer-Sørensen gate [65]. The gate set also includes arbitrary single-qubit Z-rotations and collective rotations around a tunable axis in the XY plane [64]. The Bell-state-measurement circuit is implemented as a Z-rotation of angle $\pi/4$ for one qubit and $-\pi/4$ for the other, a Mølmer-Sørensen gate and a measurement of both qubits in the computational basis. All gates are modeled as a perfect gate followed by depolarizing channels on all partaking qubits.


Just as for color centers, entanglement generation through the Barrett-Kok protocol is modeled using the model introduced in Section 6.8. A difference with color centers, however, is that the photons emitted by ions are typically temporally impure due to off-resonant scattering [86], resulting in low Hong-Ou-Mandel visibility and hence entangled-state fidelity. This can be counteracted by using a stringent detection time window and by imposing a coincidence time window. A click pattern is then only heralded as a success

in case both clicks fall within the detection time window and the time between the two clicks does not exceed the coincidence time window. The detection time window and coincidence time window can be tuned to increase the visibility, but at the cost of having a smaller success probability. In order to account for the effect of the detection time window, we employ a toy model for the temporal state of photons emitted from trapped-ion devices. This toy model does not accurately represent the true state of the emitted photons, but as we show in Figure 6.10, it can be used to capture the trade-off between success probability and visibility well. In this toy model, we model photons as mixtures of pure photons emitted at different times. The pure photons have one-sided exponential wavefunctions, and the emission time is also distributed according to a one-sided exponential. Under these assumptions, the detection probability, coincidence probability and visibility can all be exactly calculated as a function of the in total two parameters that describe these two exponentials. These calculations are performed in Section 6.9, and the results can be used in conjunction with the model in Section 6.8 to calculate the success probability and state. To show that this model can be used to capture the success probability and visibility with good accuracy, we have performed a joint least-square procedure for the detection-time probability density function, the coincidence probability and the visibility to match the two free parameters to the data presented in [86]. This data has been produced by emitting two photons from the same trapped-ion device, frequency converting these photons, and then making them interfere. In Figure 6.10 (a), we show the resulting theoretical results and compare them to the experimental results.

6

Instead of basing the parameters we use in our simulations on [86], we base them on data for the interference of photons emitted by two distinct ion traps [53], as this more accurately represents the scenario we investigate in this study. We determine again the two parameters that describe the two exponentials by fitting to the data using the exact same method as above. The half-life time of the fitted exponentials representing the wave function and the emission time were found to be $3.01 \mu\text{s}$ and $6.79 \mu\text{s}$ respectively, with the fits and the data shown in Figure 6.10 (b). This data has been taken using a detection time window of $17.5 \mu\text{s}$. Therefore, for consistency, we use a fixed detection time window of $17.5 \mu\text{s}$ throughout our simulations, and hence the parameters shown in Table 6.3 (such as, e.g., the photon detection probability excluding attenuation losses and the dark count probability) all assume a detection time window of $17.5 \mu\text{s}$. On the other hand, the coincidence time window is treated as a freely tunable parameter, allowing for a trade-off between rate and fidelity. The value for the visibility reported in Table 6.3 and Table II of the main text was obtained from the model in Section 6.9 using the fitted parameters reported above, a detection time window of $17.5 \mu\text{s}$ and a coincidence time window of $0.5 \mu\text{s}$.

We note that [86] includes a physically-motivated theoretical model for the trade-off between coincidence probability and visibility as a function of the detection and coincidence time windows. We have not used their model here as it requires numerical integration to evaluate, while our model can be rapidly evaluated using an analytical closed-form expression. Additionally, our goal here is not to predict the behaviour of a specific physical system but to accurately represent the trade-off between rate and fidelity without overfitting to experimental data. Finally, as our toy model does not attempt to closely capture the physics of any individual system, it can be considered to be system agnostic. It could



(a)

(b)

Figure 6.10: Comparison between data from two different experiments and the toy model introduced in Section 6.9. In both experiments, the detection probability density and coincidence probability were not conditioned on the successful detection of two photons. To account for this, we have multiplied both the detection-probability-density data and the coincidence-probability-density data by a different overall scaling factor. Both the scaling factors, the parameters of the two exponentials describing the photon state and an offset for the detection probability density have been determined using a least-squares procedure. The least-squares procedure has been performed jointly for the three data sets corresponding to the same experiment by summing the square errors of all three. Here, the largest weight has been given to the detection probability density (10^6), the second largest to the visibility (10^5), and the smallest to the coincidence probability (1). (a) Comparison to data from Meraner et al. [86]. In the experiment, the Rabi pulse was terminated after approximately $9 \mu\text{s}$, therefore we have only compared the first nine μs . Because of the terminated pulse the detection probability density falls to zero at approximately $12 \mu\text{s}$. Therefore, effectively the entire wave packet is detected. To reproduce this in our model, we have not implemented a detection time window (or equivalently, have set the detection time window to infinite). The fitted half-life times of the exponentials representing the wave function and emission time are $2.40 \mu\text{s}$ and $2.76 \mu\text{s}$ respectively. (b) Comparison to data from Krutyanskiy et al. [53] We base the modeling for the visibility and coincidence probability of ion traps in this paper on the fit shown here. A detection time window of $17.5 \mu\text{s}$ was used in the experiment. The detection-probability-density data used here corresponds to “node A” from [53]. The fitted half-life times of the exponentials representing the wave function and emission time are $3.01 \mu\text{s}$ and $6.79 \mu\text{s}$ respectively. Note that the x axes for the middle figures are the same as for the bottom figures.

thus be fitted to different types of photon sources, giving it a potentially broader scope of application.

6.6.7 Abstract nodes

The purpose of the abstract nodes is to provide a general model for processing nodes. Therefore, their modeling is kept simple and platform-agnostic: there is all-to-all connectivity between the qubits, all of them can be used to generate light-matter entanglement, they all have the same coherence time properties and all quantum gates are available. The Bell-state measurement circuit implemented by abstract nodes is the usual one: a controlled-NOT gate, followed by a Hadamard on the control qubit and a measurement of both qubits in the computational basis. We note that this model and its NetSquid implementation are not novel, having first been introduced in [79].

In order to quantify the level of accuracy that is sacrificed by considering a model with a higher degree of abstraction, we compare the performance of a single abstract-node repeater in the Delft-Eindhoven path to the equivalent color center and trapped ion setups. To do so, we require a method of converting hardware parameters from the more in-depth models to the abstract model. We therefore start by introducing this mapping.

Color center to abstract model mapping

The emission fidelity, visibility, dark count probability and probability of photon detection excluding attenuation losses are mapped without change from the color center model to the abstract model. An entanglement swap in an NV platform consists of one-qubit gates on both carbon and electron, two-qubit gates and measurement and initialization of the electron (see Figure 17 in Supplementary Note 5 of [60] for an image of the circuit). Imperfections in gates and initialization are modelled by depolarizing channels in the NV model, while the measurement error is modelled by probabilistic bit flips. In mapping to the abstract model we approximate the measurement error as a depolarizing channel. All the errors associated to the operations in the circuit are then multiplied together to obtain a single parameter s_q . $1 - s_q$ is used to parameterize a depolarizing channel applied after a perfect Bell-state measurement. The action of this depolarizing channel on a given state ρ as a function of s_q is given by equation (6.21), from which we can see that s_q is a measure of the quality of an entanglement swap.

$$\phi(\rho, s_q) = \left(\frac{1 + 3s_q}{4} \right) \rho + \frac{1 - s_q}{4} (X\rho X + Y\rho Y + Z\rho Z). \quad (6.21)$$

In our color center model, the coherence times of the carbon spins are different from those of the electron spin. This subtlety is lost in the abstract model, where we take the coherence time of all qubits to be the same as the carbon spins'. Other dephasing processes such as induced dephasing [34], which are present in our color center model, are ignored in the abstract model. In Table 6.4, we present the abstract model parameters obtained from the color center baseline hardware parameters as shown in Table 6.2.

Having introduced the process by which we map color center parameters to the abstract model, we now proceed with the results of validating the abstract model against the NV model. To do so, the following steps were taken: (i) define the values of the baseline hardware parameters for the more in-depth model and map them to the abstract model following the procedure described above, thus obtaining the corresponding abstract model baseline, (ii) run the simulation, (iii) improve both baselines using the improvement factor technique introduced in Section 6.3 and (iv) repeat steps (ii) and (iii) for improvement factors in the desired range.

Parameter	Noise	Duration/Time
Visibility	0.9	-
Dark count probability	1.5×10^{-7}	-
Photon detection probability excluding attenuation losses	5.1×10^{-4}	-
Spin-photon emission	$F = 1$	$3.8 \mu\text{s}$
Swap quality	0.83	$503.7 \mu\text{s}$
T1	-	10 hours
T2	-	1 s

Table 6.4: Baseline abstract model hardware parameters mapped from color center baseline shown in Table 6.2.

This analysis is done both for single and double-click entanglement generation, as we simulated color center repeaters running both protocols.

Color center validation

In Figure 6.11 we show the results of the validation for the abstract model against the NV model, for single-click (top) and double-click (bottom) entanglement generation.


The agreement is similar for both protocols. The rate of entanglement generation, shown on the plots on the right side, is identical for both models. The only source of difference timing-wise is in how long it takes to perform an entanglement swap, with color center taking slightly longer due to its more complex circuit. However, the low success probability of generating entanglement means that many attempts are required, rendering the time devoted to local operations negligible. Since the time taken per entanglement generation attempt is equal in both models, it is to be expected that the rate is identical.

For small improvement factors, there is a sizeable gap in the average teleportation fidelity achievable in each model, as shown on the plots on the left side. This fidelity is significantly larger for the abstract model. We conjecture that this is due to sources of noise that are present in the NV model but not in the abstract model. These include induced dephasing noise, probability of double photon excitation and deviations in interferometric phase, the last two being single-click specific. As parameters improve, the magnitude of these noise sources drops, and so does the gap between the fidelity achieved by the two setups.


Overall, the abstract model captures the behavior of the more in-depth NV model reasonably well. However, it does result in a more optimistic picture regarding the parameter quality required to achieve certain fidelity targets. For example, in the abstract model with double-click entanglement generation, an improvement factor of 5 suffices to reach an average teleportation fidelity of 0.7. This same target requires an improvement factor of 7 in the NV model. This supports the need for detailed hardware models, which take platform-specific limitations and sources of noise into account.

Trapped ion to abstract model mapping

The visibility, dark count probability, photon detection probability excluding attenuation losses and spin-photon emission parameters are mapped without change from the trapped ion model to the abstract model. The process by which the swap quality parameter is obtained is identical to the one described in Section 6.6.7. There is a notable difference



(a) Single-click.



(b) Double-click.

Figure 6.11: Performance of color center nodes and abstract nodes on the Delft - Eindhoven setup, with single-click (top) and double-click (bottom) entanglement generation. The leftmost point on both plots corresponds to the baseline hardware values. The points to the right were obtained by uniformly improving the hardware over this baseline. The error bars represent the standard error of the mean and are often smaller than the markers. “Average number s/success” is the average number of seconds per entangled pair that is successfully distributed.

between how memory decoherence is accounted for in the two models. In our trapped ion model, states stored in memory suffer from collective dephasing and no relaxation is considered. Our abstract model, on the other hand, considers a T_1 , T_2 memory noise model, as empirically it has been found to fit well to a large variety of physical systems. When mapping from trapped ion parameters to abstract model parameters, we take the abstract model's T_2 to be given by the collective dephasing coherence time of the trapped ion and we set T_1 to infinity, i.e. we consider no relaxation in the abstract model. The collective dephasing affecting trapped ion qubits follows a Gaussian shape, whereas the dephasing in the abstract model follows a simple exponential. In Table 6.5, we present the abstract model parameters obtained from the trapped ion baseline hardware parameters as shown in Table 6.3.

Parameter	Noise	Duration/Time
Visibility	0.89	-
Dark count probability	1.5×10^{-5}	-
Photon detection probability excluding attenuation losses	0.0288	-
Spin-photon emission	F = 0.99	50 μ s
Swap quality	0.94	1.91 ms
T1	-	-
T2	-	6 ms

Table 6.5: Baseline abstract model hardware parameters mapped from trapped ion baseline shown in Table 6.3.

Trapped ion validation

In this section, we investigate how well the simpler abstract model captures the behavior of the trapped ion model. We do this considering only double-click entanglement generation, as this was the only entanglement generation protocol we considered when performing trapped ion simulations.

In Figure 6.12 we show the results of the validation of the abstract model against the trapped ion model. The agreement in terms of the entanglement generation rate is perfect, with the rates overlapping for all values of the improvement factor. This is to be expected, since the end-to-end entanglement generation time is dominated by the time spent attempting to generate elementary links, and each attempt takes the same amount of time in both models. The average teleportation fidelity follows the same trend for both models, starting at very low values for current hardware parameters and quickly rising as hardware parameters are improved. We note that for low improvement factors, the abstract model achieves a higher fidelity. The opposite seems to be true for high improvement factors, although there the difference is small and does not exceed one error bar. This can be explained by the Gaussian nature of the trapped ion dephasing. In the trapped ion model, the probability of a state stored in memory dephasing over a given period of time t is $1 - e^{-t^2/T^2}$, with T being the coherence time. In the abstract model, this probability is $1 - e^{-t/T}$. This means that for $t/T < 1$, the probability of error for trapped ions is smaller, while the opposite is true for $t/T > 1$. At low values of the improvement factor, the success probability of entanglement generation is small, as are coherence times. Therefore,




Figure 6.12: Performance of trapped ion nodes and abstract nodes on the Delft - Eindhoven setup, double-click entanglement generation. The leftmost point on both plots corresponds to the baseline hardware values. The points to the right were obtained by uniformly improving the hardware over this baseline. The error bars represent the standard error of the mean.

the time a state is expected to stay in memory is likely larger than the coherence time, and we expect that the error rate is higher in the trapped ion model. As parameters improve, it becomes more likely that states remain in memory for periods of time smaller than the coherence time, which is the regime in which the error rate is higher in the abstract model. This is in line with what is observed in Figure 6.12. Overall, the agreement is better than what was observed in Section 6.6.7. There, owing to noise sources present in the color center model that were ignored in the abstract model, the latter performed better than the former. No noise sources were ignored when mapping from the trapped ion model to the abstract model, so this better agreement was to be expected. We conclude that the abstract model captures the behavior of the more detailed trapped ion model almost perfectly in the setup we considered.

6

6.6.8 Entanglement generation

For near-term parameters, the success probability of entanglement generation is very low. This means that many entanglement generation attempts are required, and that a simulation of this process would spend most of its time simulating failed attempts. This is computationally very inefficient, so we instead perform entangled state insertion, through a process we call *magic* [95]. This process was first introduced in [78].

Magic works as follows: once two nodes have decided to generate entanglement together, we sample from a geometric distribution in order to determine how many attempts would have been required to succeed. The success probability of this geometric distribution is limited by the product of the probabilities of emitting the photon in the correct mode, capturing it into the fiber, frequency-converting it, transmitting it through the fiber and detecting it at the detector. Furthermore, imperfections such as the imperfect indistinguishability of interfering photons and detector dark counts also impact the success probability. Their effect depends on whether a single-click or double-click protocol is used.

The elapsed time for the entanglement generation process is given by the product of

the sampled number of required attempts and the duration of one attempt, which is in turn given by the sum of the emission time and the photon travel time.

The state generated is given by an analytical model which is different for single and double-click entanglement generation. For more details, see Section 6.8.

6.7 Target metric

In this section, we explain the target metric used in this chapter. As discussed in the main text, there are two conditions on end-to-end entanglement distribution that define the target. The first is on the average fidelity with which qubits can be teleported using the generated entangled states, and the second is on the rate at which such states are generated. The target values for the teleportation fidelity and entangling rate are chosen such that the quantum link would be able to support Verifiable Blind Quantum Computation (VBQC) [35] when the server consists of a powerful quantum computer with a coherence time of 100 seconds. We show that if the targets are met, the client would be able to execute VBQC by preparing states at the powerful quantum computer using either quantum teleportation or remote state preparation (for remote state preparation, see Section 6.7.5).

The following results presented in this section are novel:

- the constraint equation that, when solved, guarantees VBQC is feasible (Theorems 1 and 4);
- the extension of the noise robustness theorem in [35] to guarantee that VBQC is feasible when the *average* error probability can be bounded instead of the *maximum* error probability, assuming that the error probabilities across different rounds are independent and identically distributed (Theorem 2 and Section 6.7.4);
- a modified version of the VBQC protocol [35] that is based on remote state preparation instead of qubit transmission (Protocol 1) and a proof that, in the absence of local noise, it is equivalent to the original protocol where some effective quantum channel is used for qubit transmission (thereby guaranteeing that the correctness of the original protocol is inherited; we note that we have not otherwise investigated the security of this protocol) (Theorem 3).

6.7.1 Teleportation fidelity

We consider the following quantum-teleportation protocol [35]. A one-qubit information state ρ is teleported using a two-qubit resource state σ shared by two parties. A Bell-state measurement is performed between the qubit holding the information state and one of the qubits in the resource state. If the outcome of the measurement corresponds to Bell state

$$|\Phi_{ij}\rangle \equiv X^i Z^j |\Phi^+\rangle \quad (6.22)$$

with $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ then the Pauli correction $X^i Z^j$ is performed on the one remaining qubit. Executing this protocol results in transmitting the information state through the teleportation channel Λ_σ .

Definition 1 *Teleportation channel.* The teleportation channel associated with the two-qubit state σ is given by the single-qubit quantum channel

$$\Lambda_\sigma(\rho) \equiv \sum_{i,j} (X^i Z^j \otimes \langle \Phi_{ij} |) (\sigma \otimes \rho) (X^i Z^j \otimes | \Phi_{ij} \rangle). \quad (6.23)$$

We note that if $\sigma = |\Phi_{00}\rangle\langle\Phi_{00}|$ then Λ_σ is the identity map. The average teleportation fidelity corresponding to the resource state σ is given by $F_{\text{tel}}(\sigma)$.

Definition 2 *Average teleportation fidelity.* The average teleportation fidelity associated with the two-qubit state σ is given by

$$F_{\text{tel}}(\sigma) \equiv \int_{\psi} \langle \psi | \Lambda_\sigma(|\psi\rangle\langle\psi|) | \psi \rangle d\psi, \quad (6.24)$$

where the integral is over the Haar measure.

We note that by the Haar measure, we here mean the uniform measure over single-qubit quantum states, i.e. the uniform measure on the unit sphere in \mathcal{C}^2 . It is the unique measure that is invariant under unitary transformations [96].

Finally, we note that if the sender and receiver agree on a unitary U , then teleportation can also be executed as follows. First, the sender applies U to the information state. Second, the sender teleports the resulting information state to the receiver. Last, the receiver applies the unitary U^\dagger to undo the original unitary and obtain the information state. The qubit is then transmitted through a rotated teleportation channel.

Definition 3 *Rotated teleportation channel.* The rotated teleportation channel associated with the two-qubit state σ and the unitary U is given by

$$\Lambda_{\sigma,U}(\rho) = U^\dagger \Lambda_\sigma(U\rho U^\dagger)U \quad (6.25)$$

We remark that the average teleportation fidelity is not affected by the introduction of the unitary U because of the invariance of the Haar measure, i.e.

$$F_{\text{tel}}(\sigma) = \int_{\psi} d\psi \langle \psi | \Lambda_{\sigma,U}(|\psi\rangle\langle\psi|) | \psi \rangle, \quad (6.26)$$

Using a unitary to turn a teleportation channel into a rotated teleportation channel can be advantageous when not every state on the Bloch sphere needs to be transmitted with equal fidelity, and σ is such that not all states can be transmitted with equal fidelity. By applying the unitary U , the Bloch sphere can potentially be rotated in such a way to make states for which high-fidelity transmission is desirable coincide with states that can be transmitted at high fidelity.

6.7.2 Requirements from VBQC

We consider the scenario where two nodes are connected using the one-repeater quantum connection studied in this work. These two nodes use the entanglement generated by this quantum connection to perform VBQC. Specifically, the first node (the client) utilizes

VBQC to execute a two-qubit computation on the quantum processor of the second node (the server) in a verified and blind fashion. It is assumed that the server is able to execute gates without noise and has a coherence time of 100 seconds. Our target metric is chosen such that it guarantees that the quantum connection is able to support this protocol.

A single round of the VBQC protocol involves the preparation of two qubits by the client at the server, and the execution of a series of quantum gates and measurements on those qubits by the server. The client can use the remote-state-preparation protocol [97] to use one entangled state to prepare one qubit at the server. Some rounds are computation rounds, the results of which are sent classically by the server to the client. All other rounds are test rounds. In a test round, some of the qubits transmitted to the server are traps; if the server tries to measure these qubits or performs another operation than the one specified by the client, this will become apparent from the returned computation results. However, tampering by the server is indistinguishable from noise. Only if noise is within certain bounds can the protocol be performed successfully.

This defines minimum requirements on the quantum connection used by the client to prepare the qubits at the server. First, the fidelity at which states can be prepared needs to be large enough. Second, the rate at which they can be prepared needs to be large enough as well. The reason for this is that after the first qubit is prepared at the server, it will undergo memory decoherence while waiting for the second qubit to be prepared.

Specifically, we consider the case of depolarizing memory.

Definition 4 *Depolarizing memory.* If a single-qubit quantum state ρ is stored in a depolarizing memory with coherence time T for a time t , it is subjected to a depolarizing channel

$$\mathcal{D}_p(\rho) = p\rho + (1-p)\frac{1}{2} \quad (6.27)$$

where the depolarizing parameter p is given by

$$p = e^{-\frac{t}{T}}. \quad (6.28)$$

The minimum requirements are then defined by the following theorem.

Theorem 1 *Requirements on entanglement generation for VBQC.* Assume a quantum link generates the two-qubit state σ between a client and a server with average rate R , and that the distribution times are independent and identically distributed. Furthermore, assume that qubits at the server are stored in a depolarizing memory with coherence time T . Lastly, assume that all local operations are noiseless and instantaneous. If the client prepares qubits at the server using the rotated teleportation channel $\Lambda_{\sigma,U}$ for some unitary U , then a unitary U exists such that the VBQC protocol proposed in [35], for a two-qubit deterministic quantum computation, can be executed in a way that is composable secure with exponentially small ϵ if

$$F_{tel}(\sigma) > \frac{1}{2} \left(1 + \frac{1}{\sqrt{2}} e^{\frac{1}{2RT}} \right). \quad (6.29)$$

Practically speaking, Theorem 1 means that VBQC with two qubits and no failure probability that is inherent to the computation is feasible in case equation (6.83) holds. A requirement is that the state σ is the same for each delivery of entanglement, and that the

distribution times are independent and identically distributed. We note that this is the case for the one-repeater setup studied in this chapter. After entanglement swapping at the repeater node takes place, the end-to-end entangled state is removed from the end nodes and the state of the network path is fully reset, making each entanglement delivery completely independent from the last, with identical distributions for both delivery times and errors. In general, the state that is delivered will depend on the amount of time entangled qubits are stored before entanglement swapping takes place at the repeater node, resulting in a state that is not the same each round. However, if the processing of the entangled state is not conditioned on the amount of storage time, the final state will effectively look like a constant mixture over all values that the storage time can take.

In this chapter, we consider two different sets of target teleportation fidelity and target rate, namely $(F_{\text{tel}}, R) = (0.8717, 0.1 \text{ Hz})$ and $(0.8571, 0.5 \text{ Hz})$. Both of these have been chosen to satisfy Eq. (6.83) for $T = 100$ seconds.

6.7.3 Proving Theorem 1

In [35], it is shown that the VBQC protocol is composable secure with exponentially small ϵ in case the noise is such that the failure probability of each individual test round can be upper bounded. Key to proving Theorem 1 is a relaxation of this condition: two-qubit VBQC is also feasible if instead the *average* failure probability of test rounds can be upper bounded, in case the failure probabilities are independent and identically distributed. This is stated in the following theorem.

Theorem 2 (*Local correctness of VBQC protocol on Noisy Devices*) *Let p denote the inherent error probability of the quantum computation, which is executed using a k -colorable graph state. Assume that, for every test round, the probability that at least one of the trap-measurement outcomes is incorrect is a random variable. Furthermore, assume that these are independent and identically distributed for all test rounds. Let q be the expected value of these random variables. The VBQC protocol presented in [35] is ϵ_{cor} -locally-correct with exponentially low ϵ_{cor} if $q < (1/k)(2p - 1)/(2p - 2)$.*

Theorem 2 is proven in Section 6.7.4 and allows us to derive the following lemma.

Lemma 1 *Two-qubit VBQC for deterministic computations is composable secure with exponential ϵ if the probabilities that the trap-measurement outcome is incorrect are independent and identically distributed for all test rounds and the average probability that the trap-measurement outcome in a single test round is incorrect, q , satisfies $q < 1/4$.*

Proof: First, we note that all two-qubit graph states are at least two-colorable, i.e., $k \leq 2$. Second, we note that for deterministic computations the inherent error probability of the computation is zero, i.e. $p = 0$. Then, from Theorem 2, it follows that if $q < 1/4$ is true, then the VBQC protocol is ϵ_{cor} -locally-correct with exponentially low ϵ_{cor} . Additionally, as shown in [35], the VBQC protocol is ϵ_{bl} -local-blind and ϵ_{ver} -local-verifiable with ϵ_{ind} -independent-verification, with ϵ_{bl} , ϵ_{ver} and ϵ_{ind} exponentially low. Therefore, as in [35], it follows that the protocol is composable secure with exponential ϵ . \square

During a test round, the client randomly designates one of the two qubits that it prepares at the server the “dummy” qubit and the other the “trap” qubit. The client that

remotely prepares the dummy qubit in $|d\rangle$, where d is chosen uniformly at random by the client from $\{0, 1\}$. It prepares the trap qubit in the state $|+\rangle_\theta$ defined by

$$|\pm_\theta\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm e^{i\theta}|1\rangle), \quad (6.30)$$

where the client chooses θ uniformly at random from $\Theta \equiv \{i\pi/4\}_{0 \leq i \leq 7}$. That is, the trap qubit will be in one of eight equidistant quantum states on the equator of the Bloch sphere. The server will perform a CZ gate between the two qubits, measure them in the basis $\{|+\delta\rangle, |-\delta\rangle\}$, and send the measurement outcomes to the client. Here, $\delta = \theta + r\pi$, where r is chosen uniformly at random by the client from $\{0, 1\}$. The test round is declared a success if the measurement on the trap qubit yields $d \oplus r$ and a failure otherwise. If the server is honest and there is no noise, test rounds are always successful. Otherwise, we show that the following Lemma holds:

Lemma 2 *If, during a test round of two-qubit VBQC, the trap qubit is prepared with fidelity F_{trap} and the dummy qubit is prepared with fidelity F_{dummy} , then the probability that the measurement outcome on the trap qubit is incorrect is given by*

$$p_{\text{fail}} = F_{\text{dummy}}(1 - F_{\text{trap}}) + F_{\text{trap}}(1 - F_{\text{dummy}}). \quad (6.31)$$

Proof: Consider the case $d = r = 0$. In that case, we can write

$$\rho_{\text{dummy, server}} = F_{\text{dummy}}|0\rangle\langle 0| + (1 - F_{\text{dummy}})|1\rangle\langle 1| + a|0\rangle\langle 1| + a^*|1\rangle\langle 0| \quad (6.32)$$

for some constant a and

$$\rho_{\text{trap, server}} = F_{\text{trap}}|+\rangle\langle +| + (1 - F_{\text{trap}})|-\rangle\langle -| + b|+\rangle\langle -| + b^*|-\rangle\langle +| \quad (6.33)$$

for some constant b . Here, we have made use of the fact that both $\{|0\rangle, |1\rangle\}$ and $\{|+\rangle, |-\rangle\}$ are complete bases for the single-qubit Hilbert space.

After receiving both states, the server will perform a CZ gate between the two qubits, and then measure the trap qubit in the $\{|+\rangle, |-\rangle\}$ basis. Whether the test round is successful or not depends on whether the expected outcome $d \oplus r = 0$, i.e. $|+\rangle$, is obtained from this measurement. In order to get the measurement statistics on the trap qubit, we can first trace out the dummy qubit. With that in mind, let's look at what happens with the term

$$\begin{aligned} & a\text{CZ}|0\rangle\langle 1|\rho_{\text{trap, server}}\text{CZ} + a^*\text{CZ}|1\rangle\langle 0|\rho_{\text{trap, server}}\text{CZ} \\ & = a|0\rangle\langle 1|\rho_{\text{trap, server}}Z + a^*|1\rangle\langle 0|Z\rho_{\text{trap, server}}. \end{aligned} \quad (6.34)$$

After the CZ has been performed, the off-diagonal terms of $\rho_{\text{dummy, server}}$ are still off diagonal. These will vanish when tracing out the dummy qubit and can therefore be safely ignored. Therefore, we make the substitution

$$\rho_{\text{dummy, server}} \rightarrow F_{\text{dummy}}|0\rangle\langle 0| + (1 - F_{\text{dummy}})|1\rangle\langle 1|. \quad (6.35)$$

Then, the effect of the CZ is easy to evaluate, giving

$$\rho_{\text{after CZ}} = F_{\text{dummy}}|0\rangle\langle 0|\rho_{\text{trap, server}} + (1 - F_{\text{dummy}})|1\rangle\langle 1|Z\rho_{\text{trap, server}}Z \quad (6.36)$$

which, after tracing out the dummy qubit, gives

$$\begin{aligned} \rho_{\text{trap, after CZ}} = & \left(F_{\text{dummy}} F_{\text{trap}} + (1 - F_{\text{dummy}})(1 - F_{\text{trap}}) \right) |+\theta\rangle\langle+\theta| \\ & + \left(F_{\text{dummy}}(1 - F_{\text{trap}}) + F_{\text{trap}}(1 - F_{\text{dummy}}) \right) |-\theta\rangle\langle-\theta| \\ & + c |+\theta\rangle\langle-\theta| + c^* |-\theta\rangle\langle+\theta|, \end{aligned} \quad (6.37)$$

where c is a function of b , F_{dummy} and F_{trap} . Applying a POVM with elements $|+\theta\rangle\langle+\theta|$ and $|-\theta\rangle\langle-\theta|$ then gives a failure probability of the test round of

$$p_{\text{fail}} = \text{Tr} \left(|-\theta\rangle\langle-\theta| \rho_{\text{trap, after CZ}} \right) = F_{\text{dummy}}(1 - F_{\text{trap}}) + F_{\text{trap}}(1 - F_{\text{dummy}}). \quad (6.38)$$

This calculation can be repeated for all three cases where $d = r = 0$ is false, each time giving the exact same outcome. \square

We now have a formula for the probability that a test round fails, given by equation (6.38). However, this formula depends on the fidelity with which specific states are transmitted over the teleportation channel. These states are randomly chosen during each test round ($|0\rangle$ or $|1\rangle$ for the dummy qubit, $|+\theta\rangle$ for the trap qubit). This means that, in general, the failure probability is not constant per round. Before we are able to use Lemma 1, we need to know something about the average failure probability per round. Additionally, we need to account for decoherence in the server's memory while waiting for the second qubit to be prepared at the server. Both are accounted for in the following lemma.

Lemma 3 *Assume a quantum link generates the two-qubit state σ between a client and a server with average rate R , and that the distribution times are independent and identically distributed. Additionally assume that a unitary U has been chosen such that dummy qubits can be transmitted through a rotated teleportation channel with average fidelity*

$$\bar{F}_{\text{dummy}} \equiv \frac{1}{2} \left(\langle 0 | \Lambda_{\sigma, U} (|0\rangle\langle 0|) |0\rangle + \langle 1 | \Lambda_{\sigma, U} (|1\rangle\langle 1|) |1\rangle \right) \quad (6.39)$$

and trap qubits with average fidelity

$$\bar{F}_{\text{trap}} \equiv \frac{1}{8} \sum_{\theta \in \Theta} \langle +\theta | \Lambda_{\sigma, U} (|+\theta\rangle\langle+\theta|) |+\theta\rangle. \quad (6.40)$$

Assume that the condition

$$\bar{F}_{\text{dummy}}(1 - \bar{F}_{\text{trap}}) + \bar{F}_{\text{trap}}(1 - \bar{F}_{\text{dummy}}) \leq \frac{1}{2} \quad (6.41)$$

holds. Furthermore, assume that qubits received by the server are stored in depolarizing quantum memory with coherence time T . Lastly, assume that all local operations are noiseless and instantaneous. In that case, for two-qubit VBQC, the average test-round failure probability is bounded by

$$q \leq e^{-\frac{1}{RT}} \left[\bar{F}_{\text{dummy}}(1 - \bar{F}_{\text{trap}}) + \bar{F}_{\text{trap}}(1 - \bar{F}_{\text{dummy}}) \right] + \frac{1}{2} (1 - e^{-\frac{1}{RT}}). \quad (6.42)$$

Proof: Let Δt be the time between the generation of the first and second entangled state. Then, the first qubit is stored for time Δt in depolarizing memory until the second qubit is prepared at the server. If the qubit was prepared at the server with fidelity F , the depolarizing noise will have the effect

$$F \rightarrow e^{-\frac{\Delta t}{T}} F + \frac{1}{2} (1 - e^{-\frac{\Delta t}{T}}). \quad (6.43)$$

We note that equation (6.31) is symmetric under interchange of F_{dummy} and F_{trap} . Therefore, we can assume that the dummy qubit is prepared first without loss of generality. Writing F_{dummy} and F_{trap} for the fidelities with which the qubits are teleported to the server (i.e. excluding the effect of memory decoherence), it follows that

$$p_{\text{fail}} = e^{-\frac{\Delta t}{T}} \left[F_{\text{dummy}}(1 - F_{\text{trap}}) + F_{\text{trap}}(1 - F_{\text{dummy}}) \right] + \frac{1}{2} (1 - e^{-\frac{\Delta t}{T}}). \quad (6.44)$$

Now, to calculate the average failure probability $q = \langle p_{\text{fail}} \rangle$, we note that F_{dummy} , F_{trap} and Δt are all independent random variables; the first depends on the choice of d (i.e. whether to prepare $|0\rangle$ or $|1\rangle$), the second depends on the choice of θ (i.e. which $|+\theta\rangle$ to prepare), and the last depends on the probability distribution for the entanglement delivery time. This allows us to write

$$q = \left\langle e^{-\frac{\Delta t}{T}} \right\rangle \left[\bar{F}_{\text{dummy}}(1 - \bar{F}_{\text{trap}}) + \bar{F}_{\text{trap}}(1 - \bar{F}_{\text{dummy}}) \right] + \frac{1}{2} (1 - \left\langle e^{-\frac{\Delta t}{T}} \right\rangle). \quad (6.45)$$

Because the exponential function is convex, Jensen's inequality [98] gives

$$\left\langle e^{-\frac{\Delta t}{T}} \right\rangle \geq e^{-\frac{\langle \Delta t \rangle}{T}}. \quad (6.46)$$

The times between the distribution of two entangled states are by assumption all independent and identically distributed, i.e., they are all copies of the same Δt . The (average) entangling rate is therefore simply equal to

$$R = \frac{1}{\langle \Delta t \rangle}, \quad (6.47)$$

and therefore we find

$$\left\langle e^{-\frac{\Delta t}{T}} \right\rangle \geq e^{-\frac{1}{RT}}. \quad (6.48)$$

In case equation (6.41) holds equation (6.48) can be combined with equation (6.45) to obtain equation (6.42). \square We

note that the use of Jensen's inequality above accounts for any kind of potential jitter in the delivery of entangled qubits to the server. Whatever the distribution on the waiting time Δt looks like and at however irregular intervals entanglement is delivered, Jensen's inequality will guarantee that Eq. (6.42) holds.

Now, we want to use the average teleportation fidelity F_{tel} instead of the quantities \bar{F}_{dummy} and \bar{F}_{trap} to bound q . The final building block towards obtaining such a bound and proving Theorem 1 is the following lemma.

Lemma 4 *There exists a unitary U such that*

$$\bar{F}_{dummy} = \bar{F}_{trap} = F_{tel}, \quad (6.49)$$

where \bar{F}_{dummy} is defined in equation (6.39), \bar{F}_{trap} in equation (6.40) and F_{tel} in equation (6.24) (with σ left implicit).

Proof: While \bar{F}_{dummy} and \bar{F}_{trap} are fidelity averages over specific subsets of the Bloch sphere, F_{tel} is an average over the entire Bloch sphere. This allows us to find the relationship

$$F_{tel} = \frac{1}{3}\bar{F}_{dummy} + \frac{2}{3}\bar{F}_{trap}. \quad (6.50)$$

To see how this relationship follows, we first note that the average fidelity over the entire Bloch sphere can be written as an average over any six states that form a regular octahedron on the Bloch sphere [99]. One example of such an octahedron is given by the six eigenstates of the Pauli operators, which gives

$$\begin{aligned} F_{tel} &= \frac{1}{6} \left(\langle 0 | \Lambda_{\sigma,U} (|0\rangle\langle 0|) |0\rangle \right) + \langle 1 | \Lambda_{\sigma,U} (|1\rangle\langle 1|) |1\rangle \\ &\quad + \langle +_0 | \Lambda_{\sigma,U} (|+_0\rangle\langle +_0|) |+_0\rangle + \langle -_0 | \Lambda_{\sigma,U} (|-_0\rangle\langle -_0|) |-_0\rangle \\ &\quad + \left\langle +\frac{\pi}{2} \middle| \Lambda_{\sigma,U} \left(\left| +\frac{\pi}{2} \right\rangle \left\langle +\frac{\pi}{2} \right| \right) \middle| +\frac{\pi}{2} \right\rangle + \left\langle -\frac{\pi}{2} \middle| \Lambda_{\sigma,U} \left(\left| -\frac{\pi}{2} \right\rangle \left\langle -\frac{\pi}{2} \right| \right) \middle| -\frac{\pi}{2} \right\rangle \right) \\ &= \frac{1}{6} \left(\langle 0 | \Lambda_{\sigma,U} (|0\rangle\langle 0|) |0\rangle \right) + \langle 1 | \Lambda_{\sigma,U} (|1\rangle\langle 1|) |1\rangle + \sum_{i=0}^4 \left\langle +\frac{i\pi}{2} \middle| \Lambda_{\sigma,U} \left(\left| +\frac{i\pi}{2} \right\rangle \left\langle +\frac{i\pi}{2} \right| \right) \middle| +\frac{i\pi}{2} \right\rangle \right). \end{aligned} \quad (6.51)$$

Another such octahedron is obtained by rotating these six eigenstates around the Z axis by an angle of $\pi/4$. This gives the relation

$$\begin{aligned} F_{tel} &= \frac{1}{6} \left(\langle 0 | \Lambda_{\sigma,U} (|0\rangle\langle 0|) |0\rangle \right) + \langle 1 | \Lambda_{\sigma,U} (|1\rangle\langle 1|) |1\rangle \\ &\quad + \sum_{i=0}^4 \left\langle +\frac{(2i+1)\pi}{4} \middle| \Lambda_{\sigma,U} \left(\left| +\frac{(2i+1)\pi}{4} \right\rangle \left\langle +\frac{(2i+1)\pi}{4} \right| \right) \middle| +\frac{(2i+1)\pi}{4} \right\rangle \right). \end{aligned} \quad (6.52)$$

Adding equations (6.51) and (6.52) together and dividing by two then gives

$$F_{tel} = \frac{1}{6} \left(\langle 0 | \Lambda_{\sigma,U} (|0\rangle\langle 0|) |0\rangle \right) + \langle 1 | \Lambda_{\sigma,U} (|1\rangle\langle 1|) |1\rangle + \frac{1}{12} \left(\sum_{\theta \in \Theta} \langle +\theta | \Lambda_{\sigma,U} (|+\theta\rangle\langle +\theta|) |+\theta\rangle \right), \quad (6.53)$$

which is equivalent to Eq. (6.50).

While the unitary U will leave the average over the entire Bloch sphere, F_{tel} , invariant, the same does not hold for \bar{F}_{dummy} . The unitary rotates the Bloch sphere and thus effectively turns \bar{F}_{dummy} into an average over any pair of antipodal points on the Bloch

sphere. Each pair of antipodal points can be described using only one of the two points. The average over pairs of antipodal points can therefore be described as a function f with as domain one half of the Bloch sphere. This function f maps each point on that half of the Bloch sphere to the average fidelity of that point and its antipodal point. Now, \bar{F}_{dummy} can be chosen to correspond to any of the values in f 's range. Additionally, the average of f over its domain is equal to the average fidelity over all points on the entire Bloch sphere, i.e. F_{tel} . By the mean value theorem, we can conclude that there is a value in the range of the function that equals the average of the function. That is, there exists a choice for the unitary U such that $\bar{F}_{\text{dummy}} = F_{\text{tel}}$. Then, equation (6.50) implies that if $\bar{F}_{\text{dummy}} = F_{\text{tel}}$, then $\bar{F}_{\text{trap}} = F_{\text{tel}}$. \square

Theorem 1 is then finally proven by combining Lemmas 1, 3, and 4.

6.7.4 Proving Theorem 2

In Section F of [35], the authors show that their VBQC protocol is robust to noise, assuming that the probability of error in each round can be upper-bounded by some maximum probability of error p_{max} . More specifically, they show that the protocol can be configured in such a way that it is ϵ_{cor} -locally-correct with exponentially small ϵ_{cor} .

Here we argue that if we assume that the error probabilities are independent and identically distributed across different rounds of the protocol, then the error probability in each round is effectively equal to the average probability of error. It then suffices that this average be bounded to obtain local correctness per the result of [35], as the error probability becomes constant and the maximum error probability is equal to the average error probability. We hereby prove Theorem 2.

We assume that for each round, there is a ‘‘true’’ probability of error. This true probability of error is a random variable, with a second-order probability distribution determining what values it takes and with what probabilities [100]. Let p_{error_i} be the probability of there being an error in round i , i.e., the value taken by the true probability of error in round i , drawn from the second-order probability distribution. By the law of total probability, this can be written as:

$$p_{\text{error}_i} = \int P(\text{error}|p = p_e)P(p = p_e) dp_e, \quad (6.54)$$

where $P(\text{error}|p = p_e)$ is the probability that there is an error given that the true probability of error takes the value p_e and $P(p = p_e)$ is the probability density that this happens. By definition, $P(\text{error}|p = p_e) = p_e$, therefore we can rewrite the equation as:

$$p_{\text{error}_i} = \int p_e P(p = p_e) dp_e = \bar{p}_e, \quad (6.55)$$

with \bar{p}_e being the expected value of the second-order probability distribution from which each round's probability of error is sampled. The second-order probability distribution can then be ignored, and the probability that an error occurs in a given round is simply given by a first-order probability. It follows that the probability of error in every round is \bar{p}_e , i.e., the average probability of error, so it suffices that the average probability of error be bounded.

6.7.5 Remote state preparation

Here, we introduce a modified version of the VBQC protocol [35] in which the client “sends” qubits to the server using remote state preparation (Protocol 1). Remote state preparation is experimentally simpler than teleportation. Therefore, it is likely that early VBQC demonstrations will be more feasible when using remote state preparation than when using teleportation. We show that, when local operations are noiseless, the modified protocol is equivalent to the protocol introduced in [35] but using some specific effective quantum channel to send qubits from the client to the server. This result is expressed in Theorem 3. Therefore, the correctness property carries over from the protocol in [35] to the modified protocol, showing that it is indeed possible to use remote state preparation to execute VBQC. Additionally, we show that the conclusions about the feasibility of VBQC found above (Theorem 1) also hold for the modified protocol. That is, when the rate and fidelity of entanglement generation are good enough to support VBQC through quantum teleportation with noiseless local operations, they are also good enough to support VBQC through remote state preparation with noiseless local operations. This result is expressed in Theorem 4.

As preliminaries to proving the above, we first introduce two definitions.

Definition 5 *U-NOT operation.* The U-NOT operation Y is defined as [101]

$$Y(\alpha|0\rangle + \beta|1\rangle) = \beta^*|0\rangle - \alpha^*|1\rangle. \quad (6.56)$$

That is, Y maps any qubit state to a state that is orthogonal to it.

We note that the U-NOT operation Y is anti-unitary and hence cannot be physically implemented [101]. It maps all states on the Bloch sphere to their antipodal points, which cannot be realized with rotations only. However, mapping a specific point on the Bloch sphere to its antipodal point can always be achieved by rotating the Bloch sphere by π around any axis that is orthogonal to the axis intersecting the point. Such a mapping is provided by the following definition.

Definition 6 *$|\psi\rangle$ -NOT operations.* The family of $|\psi\rangle$ -NOT operations $\mathcal{A}_{\phi,|\psi\rangle}$ is defined by

$$\mathcal{A}_{\phi,|\psi\rangle} \equiv e^{-i\phi} Y(|\psi\rangle) \langle\psi| + e^{i\phi} |\psi\rangle \langle Y(|\psi\rangle) | \quad (6.57)$$

The parameter ϕ in $\mathcal{A}_{\phi,|\psi\rangle}$ represents the freedom in choosing which axis to use for the π rotation that maps $|\psi\rangle$ to $Y(|\psi\rangle)$ and vice versa. We note that $\mathcal{A}_{\phi,|\psi\rangle}^\dagger = \mathcal{A}_{\phi,|\psi\rangle}$. Now, we define a modified version of the VBQC protocol that makes use of remote state preparation instead of quantum teleportation.

Protocol 1 *VBQC with remote state preparation.* This protocol is the same as the VBQC protocol presented in [35], except for the following.

- Before starting the protocol, the client and server agree on a one-qubit unitary operation U .

- Whenever the client would send a qubit v in the state $|\psi\rangle$ to the server, it instead measures its half of a two-qubit resource state shared with the server in the basis $\{U|\psi\rangle, Y(U|\psi)\}$. The outcome of this measurement is stored at the client as c_v , with $c_v = 0$ corresponding to outcome $U|\psi\rangle$ and $c_v = 1$ corresponding to outcome $Y(U|\psi)\rangle$. The server applies the operation U^\dagger to its local entangled qubit. This qubit held by the server is now considered the qubit as received from the client.
- In a computation round, the measurement outcome δ_v obtained from qubit v is bit flipped by the client in case $c_v = 1$. That is,

$$\delta_v \rightarrow \delta_v \oplus c_v \quad (\text{computation round}). \quad (6.58)$$

- In a test round, for each trap qubit v , the measurement outcome δ_v is bit flipped by the client in case $c_v = 1$, and once more for every neighboring dummy qubit w for which $c_w = 1$. That is,

$$\delta_v \rightarrow \delta_v \oplus c_v \oplus \bigoplus_{w \in N_G(v)} c_w \quad (\text{test round}). \quad (6.59)$$

Here, G is the computation graph used in the VBQC protocol and $N_G(v)$ is the neighbourhood of qubit v in graph G .

The outcomes c_v are never shared with the server.

6

Lemma 5 *Effective remote-state-preparation channel.* Let $|\psi\rangle$ be some pure single-qubit state and let σ be some two-qubit density matrix shared by Alice and Bob. Let $\phi_{|\psi\rangle}$ be some function mapping the single-qubit state $|\psi\rangle$ to a real number. Furthermore, let U be some single-qubit unitary operation. If the first of two qubits holding the state σ is measured in the basis $\{U|\psi\rangle, Y(U|\psi)\}$ with measurement outcome c ($c = 0$ corresponding to $U|\psi\rangle$, $c = 1$ corresponding to $Y(U|\psi)\rangle$) after which the operation $U^\dagger \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle}^c$ is applied to the second qubit and the first qubit is traced out, then this is equivalent to sending a qubit in the state $|\psi\rangle$ through the rotated effective remote-state-preparation $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ channel given by

$$\Lambda_{\phi_{|\psi\rangle}, \sigma, U}(|\psi\rangle) = U^\dagger \Lambda_{\phi_{|\psi\rangle}, \sigma}(U|\psi\rangle)U, \quad (6.60)$$

where $\Lambda_{\phi_{|\psi\rangle}, \sigma}$ is the effective remote-state-preparation channel given by

$$\begin{aligned} \Lambda_{\phi_{|\psi\rangle}, \sigma}(|\psi\rangle) &= \left(\langle \psi | \otimes 1 \right) \sigma \left(|\psi\rangle \otimes 1 \right) \\ &+ \left(\langle \psi | \otimes 1 \right) \left(\mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \otimes \mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \right) \sigma \left(\mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \otimes \mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \right) \left(|\psi\rangle \otimes 1 \right). \end{aligned} \quad (6.61)$$

Proof: In case the state $U|\psi\rangle$ is measured on the first qubit, i.e., $c = 0$, the unnormalized post-measurement state after tracing out the first qubit and applying $U^\dagger \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle}^0 = U^\dagger$ is

$$\rho'_{c=0} = U^\dagger \left(\langle \psi | U^\dagger \otimes 1 \right) \sigma \left(U|\psi\rangle \otimes 1 \right) U. \quad (6.62)$$

This measurement outcome is obtained with probability $p_{c=0} = \text{Tr}\{\rho_{c=0}\}$, and the corresponding normalized state is $\rho_{c=0} = \rho'_{c=0}/p_{c=0}$. In case the state $Y(U|\psi)\rangle$ (which is equal

op to global phase to $\mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} U|\psi\rangle$) is measured, i.e., $c = 1$, the unnormalized state after tracing out the first qubit and applying $U^\dagger \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle}^1 = U^\dagger \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle}$ is instead

$$\begin{aligned} \rho'_{c=1} &= U^\dagger \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} \left(\langle \psi | U^\dagger \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} \otimes 1 \right) \sigma \left(\mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} U|\psi\rangle \otimes 1 \right) \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} U \\ &= U^\dagger \left(\langle \psi | U^\dagger \otimes 1 \right) \left(\mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} \otimes \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} \right) \sigma \left(\mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} \otimes \mathcal{A}_{\phi_{|\psi\rangle}, U|\psi\rangle} \right) \left(U|\psi\rangle \otimes 1 \right) U \end{aligned} \quad (6.63)$$

with measurement probability $p_{c=1} = \text{Tr}\{\rho_{c=1}\}$ and normalized state $\rho_{c=1} = \rho'_{c=1}/p_{c=1}$. The resulting state can be described as a mixture between the states corresponding to the different measurement outcomes weighted by their respective probabilities, i.e.,

$$\rho = p_{c=0}\rho_{c=0} + p_{c=1}\rho_{c=1} = \rho'_{c=0} + \rho'_{c=1} = \Lambda_{\phi_{|\psi\rangle}, \sigma, U}(|\psi\rangle). \quad (6.64)$$

□

We note that the effective remote-state-preparation channel is not a true quantum channel, i.e., it is not a completely positive trace-preserving (CPTP) map between density matrices. In fact, it is only defined for pure states, and can not (straightforwardly) be rephrased as a linear operator on a density matrix. However, the output state is a valid density matrix with trace 1, as it should be as it is the result of a measurement on and unitary evolution of the resource state σ .

6

Theorem 3 *Equivalence of VBQC with remote state preparation. Assume all local operations at both the server and the client are noiseless. Then, there exists a function $\phi_{|\psi\rangle}$ that maps single-qubit states to real numbers such that Protocol 1 is equivalent to the unaltered VBQC protocol described in [35] using the rotated effective remote-state-preparation channel $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ to send qubits in pure states from the client to the server. Here, σ is the resource state used in Protocol 1.*

Proof: In Protocol 1, the client performs bit flips on the measurement outcomes received from the server. Every measurement the server performs is in a basis of the form $\{|+\theta\rangle, |-\theta\rangle\}$ (defined in equation (6.30)). These states are mapped to each other by the Pauli Z operator, which is a $|+\theta\rangle$ -NOT operation

$$Z = \mathcal{A}_{-\theta, |+\theta\rangle}. \quad (6.65)$$

Therefore, each measurement performed by the server in Protocol 1 of which the result is bit flipped in case some number c is equal to one (i.e., $\delta \rightarrow \delta \oplus c$ where δ is the measurement result) can effectively be replaced by a unitary operation Z^c followed by a measurement of which the result is not bit flipped. It is thus as if the server applies the operation Z^c , even though the server never actually learns the value of c . This equivalence is essential to the proof.

First, we show that a computation round in Protocol 1 is equivalent to a computation round in the unaltered VBQC protocol when sending the qubits using $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ in case a specific condition on $\phi_{|\psi\rangle}$ holds. In a computation round, for each of the qubits held by the server, it first performs the unitary operation U^\dagger . Then, it executes a number of CZ gates between the qubit and some other qubits. We remind the reader that CZ gates are symmetric in the two partaking qubits; we can thus always choose which qubit we consider the

control qubit and which we consider the target qubit as we find convenient. These gates are followed by a measurement in the basis $\{|+\theta\rangle, |-\theta\rangle\}$, where the angle θ is specified by the client. The outcome δ of the measurement is bit flipped by the client according to $\delta \rightarrow \delta \oplus c_v$. In shorthand, we will write the sequence as: U^\dagger , CZs, measurement, bit flip. We will show that this sequence is equivalent to a sequence that we can apply Lemma 5 to. As a first step, we use the equivalence stated in the first paragraph of this proof to replace the measurement followed by a bit flip by a measurement preceded by the operation Z^{c_v} . The sequence is thus equivalent to the sequence: U^\dagger , CZs, Z^{c_v} , measurement. As a second step, because Z commutes with CZ, we rewrite the sequence as: U^\dagger , Z^{c_v} , CZs, measurement.

Now, using equation (6.65), the sequence can be rewritten as follows: U^\dagger , $\mathcal{A}_{-\theta, |+\theta\rangle}^{c_v}$, CZs, measurement. To enable us to move the operator U^\dagger in this sequence, we represent the unitary U in general matrix form

$$U = \begin{bmatrix} a & b \\ -e^{i\varphi} b^* & e^{i\varphi} a^* \end{bmatrix}, \quad (6.66)$$

where $|a|^2 + |b|^2 = 1$ and $\varphi \in [0, 2\pi)$. This can be used to verify that

$$Y(U|\psi\rangle) = e^{-i\varphi} UY(|\psi\rangle). \quad (6.67)$$

Therefore, for every U , there exists a φ such that for every ϕ and every $|\psi\rangle$

$$\begin{aligned} U^\dagger \mathcal{A}_{\phi-\varphi, U|\psi\rangle} &= U^\dagger \left[e^{-i(\phi-\varphi)} e^{-i\varphi} UY(|\psi\rangle) \langle \psi | U^\dagger + e^{i(\phi-\varphi)} U|\psi\rangle \langle e^{-i\varphi} UY(|\psi\rangle) \rangle^\dagger \right] \\ &= \mathcal{A}_{\phi, |\psi\rangle} U^\dagger. \end{aligned} \quad (6.68)$$

From this, we conclude that there exists a φ (determined by U) such that the sequence on qubit v is equivalent to: $\mathcal{A}_{-(\theta+\varphi), U|+\theta_v\rangle}^{c_v}$, U^\dagger , CZs, measurement. At this point, we are able to invoke Lemma 5. From this lemma, it follows that the client performing its measurement followed by the server applying the above sequence is equivalent to the client sending the state $|+\theta_v\rangle$ through a channel $\Lambda_{\phi|\psi, \sigma, U}$ for which $\phi_{|+\theta\rangle} = -\theta - \varphi$, after which the server applies the sequence: CZs, measurement. This is exactly the sequence of operations in the unaltered VBQC protocol. Therefore it follows that a computation round in Protocol 1 is equivalent to a computation round in the unaltered VBQC protocol where the channel $\Lambda_{\phi|\psi, \sigma, U}$ is used to send qubits from the client to the server in case the condition $\phi_{|+\theta\rangle} = -\theta - \varphi$ is met.

It now remains to show the same equivalence between the two protocols for test rounds. For the trap qubit v , we can again replace the measurement followed by $c_v \oplus \bigoplus_{w \in N_G(v)} c_w \equiv \bar{c}$ bit flips by a measurement without bit flips preceded by the operator $Z^{\bar{c}}$. The sequence of operations on the trap then becomes: U^\dagger , CZ gates with dummy qubits, $Z^{\bar{c}}$, and then a measurement. Now, the identity

$$CZ(1 \otimes Z) = (X \otimes 1)CZ(X \otimes 1) \quad (6.69)$$

can be used to move every bit flip due to a measurement outcome in the preparation of a dummy qubit by the client to the corresponding qubit at the server. That is, each Z^{c_w} for

$w \in N_G(v)$ is moved to the qubit w . What remains at the trap qubit v itself is then exactly the same sequence of operations as in a computation round. From what we have shown above for computation rounds, it follows that we can treat trap qubits in test rounds of Protocol 1 as if they are trap qubits in test rounds of the unaltered VBQC protocol, where the qubits are sent from the client to the server using the channel $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ if $\phi_{|+\theta\rangle} = -\theta - \varphi$. It then remains only to show that the equivalence holds for the dummy qubits.

Now, we focus on one of the dummy qubits, which we denote w . Consider the scenario where the client attempts to send the qubit w in the state $|d\rangle$, where $d \in \{0, 1\}$, to the server as in Protocol 1. This qubit is the server's half of the resource state σ . The client measures its half of σ in the basis $\{U|d\rangle, Y(U|d)\rangle\}$, with measurement outcome c_w . At the server, the following sequence of operations is applied to the qubit w : U^\dagger , $\prod_{u \in N_G(w)} CZ_{w,u}$, measurement in the basis $\{|+\theta\rangle, |-\theta\rangle\}$ for some θ . Let us first consider the case where all $u \in N_G(w)$ are trap qubits. Then, by moving the effects of bit flips from trap qubits to dummy qubits as described above, every $CZ_{w,u}$ is effectively replaced by $(X^{c_w} \otimes 1)CZ_{w,u}(X^{c_w} \otimes 1)$. Because $X^2 = 1$, this has the effect of transforming the sequence into the following: U^\dagger , X^{c_w} , $\prod_{u \in N_G(w)} CZ_{w,u}$, X^{c_w} , measurement. The second occurrence of X^{c_w} changes the outcome of the measurement on the dummy qubit. However, the measurement outcome of the dummy qubits is of no consequence in the VBQC protocol (the outcome is sent by the server to the client and then discarded by the client). Therefore, we can effectively remove the second occurrence of X^{c_w} from the sequence. For the first occurrence, we note that X is both a $|1\rangle$ -NOT gate and a $|0\rangle$ -NOT gate,

$$X = \mathcal{A}_{|1\rangle} = -\mathcal{A}_{|0\rangle}. \quad (6.70)$$

Therefore, up to a global phase in case $d = 0$, the sequence becomes equivalent to: U^\dagger , $\mathcal{A}_{|d\rangle}$, $\prod_{u \in N_G(w)} CZ_{w,u}$, measurement. We note that the unitary U is here the same as for the trap qubit (it is the same for all qubits in Protocol 1). Therefore, we can invoke equation (6.68) again to rewrite the sequence as: $\mathcal{A}_{-\varphi, U|d\rangle}$, U^\dagger , $\prod_{u \in N_G(w)} CZ_{w,u}$, measurement. It then immediately follows from Lemma 5 that this is equivalent to the client sending the qubit w in the pure state $|d\rangle$ using a quantum channel $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ for which $\phi_{|d\rangle} = -\varphi$. After the server receives the qubit through this effective channel, the remaining sequence is: $\prod_{u \in N_G(w)} CZ_{w,u}$, measurement. This is the same as in the unaltered VBQC protocol, and therefore we can treat dummy qubits in test rounds of Protocol 1 as if they are dummy qubits in the unaltered VBQC protocol that are transmitted using $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ with the condition $\phi_{|0\rangle} = \phi_{|1\rangle} = -\varphi$, provided they are only adjacent to trap qubits in the computation graph G .

As final part of our proof, we show that the above derivation for dummy qubits still holds in case they are adjacent to other dummy qubits in the computation graph. Every CZ with a trap qubit results in two X^{c_w} s. When the dummy qubit is only adjacent to trap qubits, X^{c_w} s resulting from neighboring CZs then cancel out in the middle (because $X^2 = 1$), such that only operators at the beginning and ending of the entire sequence remain. However, a CZ with another dummy qubit does not give any X^{c_w} s. X^{c_w} s from CZs with trap qubits that enclose one or more CZs with dummy qubits can then no longer cancel against one another. A way out is offered by the following identity:

$$(1 \otimes X)CZ = CZ(Z \otimes X). \quad (6.71)$$

This means that X can be commuted through CZs at the cost of inducing a Z at the other qubit partaking in the CZ. Now, if a Z is induced on a dummy qubit, it can be commuted through all CZs the dummy partakes in and placed in front of the measurement. Here, it results in an effective bit flip on the measurement outcome. Since again the measurement outcomes at the dummy qubits are inconsequential, the operator can safely be ignored. This means that X^{c_w} s can safely commute through all the CZs with other dummy qubits, allowing them to cancel out as before and get again to a sequence where there is one X^{c_w} before all the CZs and one after. The sequence then is the same as when the dummy qubit would not be adjacent to other dummy qubits, and the same conclusion derived in the above paragraph holds.

Combining all the above, we conclude that Protocol 1 is equivalent to the VBQC protocol [35] using the channel $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ to send pure state from the client to the server. This holds for any function $\phi_{|\psi\rangle}$ that satisfies

$$\phi_{|+\theta\rangle} = -\theta - \varphi, \quad (6.72)$$

$$\phi_{|d\rangle} = -\varphi, \quad (6.73)$$

for any $d \in \{0, 1\}$, for any $\theta \in [0, 2\pi)$, and where φ depends on the choice of unitary U in Protocol 1 (it is the parameter appearing in equation (6.66)). There exists an infinite number of functions satisfying this condition (note that it is not required that the function is continuous; in fact it does not matter in the least how the function behaves away from $|d\rangle$ and $|+\theta\rangle$ as these are the only states that are ever sent through the channel), and therefore the theorem is proven. \square

Lemma 6 *Equivalence of remote state preparation and quantum teleportation. The average fidelity of the effective remote-state-preparation channel (equation (6.61)) corresponding to the two-qubit state σ ,*

$$F_{RSP}(\sigma) \equiv \int_{\psi} d\psi \langle \psi | \Lambda_{\phi_{|\psi\rangle}, \sigma}(|\psi\rangle) | \psi \rangle, \quad (6.74)$$

is independent of the function $\phi_{|\psi\rangle}$. Furthermore, it is equal to the average teleportation fidelity corresponding to the same state σ (equation (6.24)). That is,

$$F_{RSP}(\sigma) = F_{tel}(\sigma) \quad (6.75)$$

Proof: First we rewrite the average teleportation fidelity defined in equation (6.24) as

$$F_{tel}(\sigma) = \sum_{i,j} \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \Phi_{00} | \right) \left(X^i Z^j \otimes X^i Z^j \otimes 1 \right) \left(\sigma \otimes |\psi\rangle\langle\psi| \right) \left(X^i Z^j \otimes X^i Z^j \otimes 1 \right) \left(|\psi\rangle \otimes |\Phi_{00}\rangle \right). \quad (6.76)$$

Then we use the property

$$\langle \Phi_{00} | (1 \otimes |\psi\rangle) \rangle = \frac{1}{\sqrt{2}} \langle \psi | \quad (6.77)$$

to find

$$F_{tel}(\sigma) = \frac{1}{2} \sum_{i,j} \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \psi | \right) \left(X^i Z^j \otimes X^i Z^j \right) \sigma \left(X^i Z^j \otimes X^i Z^j \right) \left(|\psi\rangle \otimes |\psi\rangle \right). \quad (6.78)$$

Since the Haar measure is invariant under unitaries, the $X^i Z^j$ can be absorbed into the state $|\psi\rangle$, giving

$$F_{\text{tel}}(\sigma) = 2 \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \psi | \right) \sigma \left(| \psi \rangle \otimes | \psi \rangle \right). \quad (6.79)$$

Similarly we can rewrite $F_{\text{RSP}}(\sigma)$ as

$$\begin{aligned} F_{\text{RSP}}(\sigma) &= \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \psi | \right) \sigma \left(| \psi \rangle \otimes | \psi \rangle \right) \\ &\quad + \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \psi | \right) \left(\mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \otimes \mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \right) \sigma \left(\mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \otimes \mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle} \right) \left(| \psi \rangle \otimes | \psi \rangle \right). \end{aligned} \quad (6.80)$$

The second term can be rewritten as

$$\begin{aligned} &\int_{\psi} d\psi \left(e^{i\phi_{|\psi\rangle}} (Y(|\psi\rangle))^\dagger \otimes e^{i\phi_{|\psi\rangle}} (Y(|\psi\rangle))^\dagger \right) \sigma \left(e^{-i\phi_{|\psi\rangle}} Y(|\psi\rangle) \otimes e^{-i\phi_{|\psi\rangle}} Y(|\psi\rangle) \right) \\ &= \int_{\psi} d\psi \left((Y(|\psi\rangle))^\dagger \otimes (Y(|\psi\rangle))^\dagger \right) \sigma \left(Y(|\psi\rangle) \otimes Y(|\psi\rangle) \right) \\ &= \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \psi | \right) \sigma \left(| \psi \rangle \otimes | \psi \rangle \right). \end{aligned} \quad (6.81)$$

The last step here follows from the fact that an integral over all antipodal points on the Bloch sphere is itself just an integral over all points on the Bloch sphere. We thus find

$$F_{\text{RSP}}(\sigma) = 2 \int_{\psi} d\psi \left(\langle \psi | \otimes \langle \psi | \right) \sigma \left(| \psi \rangle \otimes | \psi \rangle \right). \quad (6.82)$$

□

Theorem 4 *Requirements on entanglement generation for VBQC through remote state preparation. Assume a quantum link generates the two-qubit state σ between a client and a server with average rate R . Furthermore, assume that qubits at the server are stored in a depolarizing memory with coherence time T . Lastly, assume that all local operations are noiseless and instantaneous. Then, a unitary U exists such that Protocol 1 can be executed to realize the VBQC protocol [35] for a two-qubit deterministic quantum computation in a way that is compositably secure with exponentially small ϵ if*

$$F_{\text{tel}}(\sigma) > \frac{1}{2} \left(1 + \frac{1}{\sqrt{2}} e^{\frac{1}{2RT}} \right). \quad (6.83)$$

Proof: By Theorem 3, there exists a function $\phi_{|\psi\rangle}$ such that Protocol 1 is equivalent to the VBQC protocol as presented in [35] where qubits are transmitted using the channel $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$. Therefore, we can simply repeat the proof of Theorem 1 but with the channel $\Lambda_{\sigma, U}$ replaced by $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$. This results then exactly in Eq. (6.83), but with $F_{\text{tel}}(\sigma)$ replaced by the average fidelity over $\Lambda_{\phi_{|\psi\rangle}, \sigma}$, i.e., $F_{\text{RSP}}(\sigma)$. Eq. (6.83) then follows directly from Lemma 6.

We note that in order to repeat the proof of Theorem 1 two properties of the effective remote-state-preparation channel need to hold. Specifically, they need to hold in order to reproduce Lemma 4. These are properties that hold for any linear CPTP map. $\Lambda_{\phi_{|\psi\rangle}, \sigma, U}$ however is not linear, but the properties can still be shown to hold. First, the average fidelity of the channel is invariant under unitary transformations. That is,

$$\int_{\psi} d\psi \langle \psi | \Lambda_{\phi_{|\psi\rangle}, \sigma, U(|\psi\rangle) | \psi \rangle = F_{\text{RSP}}(\sigma) \quad (6.84)$$

for any unitary U . This follows most evidently from Eq. (6.82), where the effect of including a unitary U would be just to replace $|\psi\rangle \rightarrow U|\psi\rangle$, which leaves the Haar measure invariant.

Second, it needs to be shown that $F_{\text{RSP}}(\sigma)$ can be evaluated by evaluating the fidelity of $\Lambda_{\text{RSP}, \sigma}$ only at six states on the Bloch sphere forming a regular octahedron. To this end, we use the fact that six states forming a regular octahedron are the union of three mutually unbiased bases and hence form a complex projective 2-design [102]. Therefore an integral over the Bloch sphere of which the integrand is a second-order polynomial in $|\psi\rangle\langle\psi|$ can be replaced by an average over those six states. We note that this cannot be applied to Eq. (6.80) directly, as the dependence of $\mathcal{A}_{\phi_{|\psi\rangle}, |\psi\rangle}$ on $|\psi\rangle$ means that the integrand is not necessarily a second-order polynomial. However, it can be applied directly to Eq. (6.82) to express $F_{\text{RSP}}(\sigma)$ as an average over the six states. Below, we show that the resulting expression is the same as taking the average over the six states directly in Eq. (6.80).

An octahedron is made up out of three pairs of antipodal points, so we denote the set of six states $\{|\psi_i\rangle, Y(|\psi_i\rangle)\}$ for $i = 0, 1, 2$. Then, we can write (6.82) as

$$F_{\text{RSP}}(\sigma) = \frac{1}{3} \left(\sum_i \left(\langle \psi_i | \otimes \langle \psi_i | \right) \sigma \left(| \psi_i \rangle \otimes | \psi_i \rangle \right) + \sum_i \left((Y(|\psi_i\rangle))^\dagger \otimes (Y(|\psi_i\rangle))^\dagger \right) \sigma \left((Y(|\psi_i\rangle)) \otimes (Y(|\psi_i\rangle)) \right) \right). \quad (6.85)$$

It now remains to show that this is the same expression as what one would get from directly averaging the channel fidelity over these six states. This direct average can be

written as

$$\begin{aligned}
& \frac{1}{6} \left(\sum_i \langle \psi_i | \Lambda_{\phi_{|\psi_i\rangle}, \sigma} (|\psi_i\rangle) | \psi_i \rangle + \sum_i (Y(|\psi_i\rangle))^\dagger \Lambda_{\phi_{|\psi_i\rangle}, \sigma} (Y(|\psi_i\rangle)) Y(|\psi_i\rangle) \right) \\
&= \frac{1}{6} \left(\sum_i \left(\langle \psi_i | \otimes \langle \psi_i | \right) \sigma \left(| \psi_i \rangle \otimes | \psi_i \rangle \right) + \sum_i \left((Y(|\psi_i\rangle))^\dagger \otimes (Y(|\psi_i\rangle))^\dagger \right) \sigma \left((Y(|\psi_i\rangle)) \otimes (Y(|\psi_i\rangle)) \right) \right) \\
&+ \frac{1}{6} \left(\sum_i \left(\langle \psi_i | \otimes \langle \psi_i | \right) \left(\mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \otimes \mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \right) \sigma \left(\mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \otimes \mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \right) \left(| \psi_i \rangle \otimes | \psi_i \rangle \right) \right) \\
&+ \sum_i \left((Y(|\psi_i\rangle))^\dagger \otimes (Y(|\psi_i\rangle))^\dagger \right) \left(\mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \otimes \mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \right) \sigma \left(\mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \otimes \mathcal{A}_{\phi_{|\psi_i\rangle}, |\psi_i\rangle} \right) \left((Y(|\psi_i\rangle)) \otimes (Y(|\psi_i\rangle)) \right) \\
&= \frac{1}{6} \left(\sum_i \left(\langle \psi_i | \otimes \langle \psi_i | \right) \sigma \left(| \psi_i \rangle \otimes | \psi_i \rangle \right) + \sum_i \left((Y(|\psi_i\rangle))^\dagger \otimes (Y(|\psi_i\rangle))^\dagger \right) \sigma \left((Y(|\psi_i\rangle)) \otimes (Y(|\psi_i\rangle)) \right) \right) \\
&+ \frac{1}{6} \left(\sum_i \left(e^{-i\phi_{|\psi_i\rangle}} (Y(|\psi_i\rangle))^\dagger \otimes e^{-i\phi_{|\psi_i\rangle}} (Y(|\psi_i\rangle))^\dagger \right) \sigma \left(e^{i\phi_{|\psi_i\rangle}} Y(|\psi_i\rangle) \otimes e^{i\phi_{|\psi_i\rangle}} Y(|\psi_i\rangle) \right) \right) \\
&+ \sum_i \left(e^{i\phi_{|\psi_i\rangle}} \langle \psi_i | \otimes e^{i\phi_{|\psi_i\rangle}} \langle \psi_i | \right) \sigma \left(e^{-i\phi_{|\psi_i\rangle}} | \psi_i \rangle \otimes e^{-i\phi_{|\psi_i\rangle}} | \psi_i \rangle \right) \\
&= \frac{1}{3} \left(\sum_i \left(\langle \psi_i | \otimes \langle \psi_i | \right) \sigma \left(| \psi_i \rangle \otimes | \psi_i \rangle \right) + \sum_i \left((Y(|\psi_i\rangle))^\dagger \otimes (Y(|\psi_i\rangle))^\dagger \right) \sigma \left((Y(|\psi_i\rangle)) \otimes (Y(|\psi_i\rangle)) \right) \right). \tag{6.86}
\end{aligned}$$

Therefore, we conclude that taking the average of the fidelity over a regular octahedron of $\Lambda_{\phi_{|\psi_i\rangle}, \sigma}$ is equivalent to taking the average over the entire Bloch sphere using the Haar measure. We note that the above argument also holds for $\Lambda_{\phi_{|\psi_i\rangle}, \sigma, U}$ for any unitary U . \square

6.8 Double-click model

In this section, we derive an analytical model for the entangled states created on elementary links when using the double-click protocol, also known as the Barrett-Kok protocol [103]. This model is used as one of the building blocks of our NetSquid simulations, as mentioned in Section 6.13. To the best of our knowledge, the analytical model is a novel result.

6.8.1 Model assumptions

The double-click protocol is a protocol for heralded entanglement generation on an elementary link. First, at each of the two nodes sharing the elementary link (designated A and B), a photon is emitted. This photon can be in one of two different photonic modes. For concreteness, we will here assume these two modes are horizontal and vertical polarization ($|H\rangle$ and $|V\rangle$, respectively), as is the case for the trapped-ion systems we consider in this work. However, depending on the hardware platform that is used, they could just as well be some other modes, e.g., different temporal modes (“early” and “late”), as is the case for the color-center systems we consider. Our model does not incorporate any effects

specific to the type of modes that are used, and therefore the assumption that the modes are polarization modes is made without loss of generality. The photon is emitted such that the mode that it is in is maximally entangled with the state of the emitter, i.e. such that the emitter - photon state after emission is (up to normalization) $|0H\rangle + |1V\rangle$. Then, the photons emitted at both nodes are sent to a midpoint station.

At the midpoint station, the photons from the two different nodes are interfered on a non-polarizing beam splitter. The two output modes are then passed through a polarizing beam splitter, of which each output mode is impinged on a single-photon detector. There are thus four single-photon detectors, two corresponding to horizontal polarization, and two corresponding to vertical polarization. This setup is illustrated in Figure 6.13. If a single photon is detected at one of the “horizontal” detectors and one at the “vertical” detectors, assuming photons in the same polarization emitted at different nodes are indistinguishable, the photons are projected on the state $|HV\rangle \pm |VH\rangle$. This results in the emitters being in the maximally entangled state $|\Psi^\pm\rangle = |01\rangle \pm |10\rangle$. The + state is obtained if the two detectors clicking are located behind the same polarizing beam splitter, while the - state is obtained if they are located behind different polarizing beam splitters. Note: if a different type of modes is used, this setup may look slightly different. For example, in case temporal modes are used, there is no need for polarizing beam splitters and using only two single-photon detectors is sufficient as the different modes can be distinguished based on the time at which they are detected.




Figure 6.13: Setup of midpoint station in double-click entanglement generation using polarization-encoded photons. Two photonic modes (a and b) are interfered on a non-polarizing 50-50 beam splitter (BS). The output modes (c and d) are then each led into a separate polarizing beam splitter (PBS). Each of the two output modes of each of the two polarizing beam splitters is caught at one of four detectors (D1, D2, D3 and D4).

In our simulations, we use an analytical model to describe the success probability and post-measurement state of the double-click scheme in the presence of several imperfections. The imperfections included in our model are

- Photon loss. Due to nonunit collection efficiency of emitters, attenuation losses in optical fiber and inefficiency of single-photon detectors, there is often only a small probability that an emitted photon is not lost before it partakes in the midpoint measurement. This is captured by the parameters p_A and p_B , where p_A (p_B) denotes the detection probability given that a photon is emitted at node A (B). These account both for attenuation losses and for the photon detection probability excluding attenuation losses.

- Imperfect indistinguishability. We assume photons emitted by the different nodes with the same polarization are not perfectly indistinguishable. This is captured using the Hong-Ou-Mandel visibility V [74, 75]. We assume the visibility is the same between two horizontally polarized photons as between two vertically polarized photons.
- Non-photon-number-resolving detectors. In our model, we distinguish between the case of photon-number-resolving detectors (case NR) and non-photon-number-resolving detectors (case NNR). If the used detectors are NR, when there are two or more photons at the same detector during a single midpoint measurement, all photons are registered individually. However, if detectors are NNR, they cannot distinguish between one or more photons. This model does not account for the case when photons can sometimes, but not always, be distinguished. Such behavior occurs in reality when e.g. two photons can only be resolved if the time between their detections is large enough.
- Detector dark counts. Sometimes, single-photon detectors report the presence of a photon when there is none. We model this using a fixed dark-count probability, p_{dc} . During a midpoint measurement, each single-photon detector gives a single dark count with probability p_{dc} , and gives none with probability $1 - p_{dc}$. Note that, in reality, for NR detectors, there is also a nonzero probability for multiple dark counts to occur during a single midpoint measurement in the same detector. Therefore, for NR detectors, treating dark counts this way will only lead to an approximation. The approximation can be expected to be accurate if the probability of multiple dark counts is negligible. For NNR detectors, this way of treating dark counts does not lead to an approximation but is perfectly accurate; multiple dark counts cannot be distinguished from one dark count, and therefore the probability of two or more dark counts and the probability of one dark count can be safely absorbed into one number, which is p_{dc} .
- Imperfect emission. It is possible that, directly after emission, the emitter and photon are not in the maximally entangled state $|\phi\rangle = \frac{1}{\sqrt{2}}(|0H\rangle + |1V\rangle)$. To capture this, the state is modelled as a Werner state of the form $\rho_{\text{emit}} = q|\phi\rangle\langle\phi| + (1-q)\frac{1}{4}$. For each node, the parameter q is chosen such that $F_{\text{em A}} (F_{\text{em B}})$ is the emission fidelity $q + (1-q)/4 = \frac{1}{4}(1+3q)$ at node A (B).

6.8.2 POVMs

To derive an analytical model, we notice that the midpoint station effectively implements a single-click midpoint measurement on each of the two different photonic modes (horizontal and vertical) separately. To make use of this, we write the photonic states as Fock states on the two different modes, such that $|H\rangle = |1\rangle_H |0\rangle_V$ and $|V\rangle = |0\rangle_H |1\rangle_V$. Distinguishing also between photons arriving from side A and side B, this allows us to write the pre-measurement state as a state in the Hilbert space that is obtained from taking the tensor product between the Hilbert spaces of the emitters and the horizontally and vertically polarized photons. That is, $\mathcal{H}_{\text{pre-measurement}} = \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_{H_A} \otimes \mathcal{H}_{H_B} \otimes \mathcal{H}_{V_A} \otimes \mathcal{H}_{V_B}$.

Since we are not interested in the post-measurement state of the photons, we can model the measurement as a POVM³. The POVM elements of the double-click midpoint station can then be derived from the single-click measurement operators as

$$M_{\text{double click, }ijkl} = 1_A \otimes 1_B \otimes [M_{\text{single click, }ij}]_{H_A H_B} \otimes [M_{\text{single click, }kl}]_{V_A V_B}. \quad (6.87)$$

Here, $M_{\text{single click, }ij}$ is the POVM element corresponding to i clicks in the first detector and j clicks in the second detector of a single-click setup. Thus, keeping in line with the naming of Figure 6.13, $M_{\text{double click, }ijkl}$ is the POVM element corresponding to i clicks in detector 1, j clicks in detector 3, k clicks in detector 2, and l clicks in detector 4, such that detectors 1 and 3, and 2 and 4 correspond to the same polarization, and detector 1 and 2, and 3 and 4, correspond to the same polarizing beam splitter. The single-click POVM elements can be obtained from Section D.5.2 of the supplemental material of [78]. In doing so, we identify the square absolute value of the overlap of the two photon wave functions, $|\mu|^2$ in [78], with the Hong-Ou-Mandel visibility V . The reason for this is that these two are the same when both photons are in a pure state [75] (we note that the photons in our model are only mixed in the polarization degree of freedom, the wave packets themselves are pure and therefore we can safely make the substitution). We modify the single-click POVM elements from [78] to account for dark counts as follows (dropping the “single click” subscript):

$$\begin{aligned} M'_{10} &= M_{10}(1 - p_{\text{dc}})^2 + M_{00}p_{\text{dc}}(1 - p_{\text{dc}}), \\ M'_{20} &= M_{20}(1 - p_{\text{dc}}) + M_{10}p_{\text{dc}}(1 - p_{\text{dc}}), \end{aligned} \quad (6.88)$$

and similarly for M'_{01} and M'_{02} . Note that we have absorbed the POVM element $M'_{30} = M_{20}p_{\text{dc}}(1 - p_{\text{dc}})$ into the POVM element M'_{20} , since for neither NR and NNR detectors will the occurrence of two and the occurrence of three detections be discriminated; for NNR detectors, the different detection events cannot be resolved, while for NR detectors, both the presence of two and of three detections will lead to heralded failure. Other POVM elements (M'_{00} , M'_{11} , M'_{21} , ...) are not needed for our analysis, since having no detection in one of the polarizations, or having two detections at different detectors for one of the modes, is always heralded as a failure.

The double-click protocol heralds two different measurement outcomes as success, namely outcome “detectors behind same polarizing beam splitter” and outcome “detectors behind different polarizing beam splitters”. These two outcomes are henceforth abbreviated “same PBS” and “different PBS”. To determine the probability of each occurring and the corresponding post-measurement states, we need to write down the POVM elements corresponding to these two outcomes. Here, we note that in the case NR, the presence of multiple detections in a single detector is always heralded as a failure, while in the case NNR, multiple detections cannot be distinguished from a single detection. This gives the

³Note that we are interested in the post-measurement state of the emitters. However, as long as the state of the photons is traced out immediately after the measurement, a POVM is sufficient to accurately determine the post-measurement state.

POVM elements (only writing the part acting on $\mathcal{H}_{H_A} \otimes \mathcal{H}_{H_B} \otimes \mathcal{H}_{V_A} \otimes \mathcal{H}_{V_B}$)

$$\begin{aligned}
M_{\text{same PBS, NR}} &= M'_{01} \otimes M'_{01} + M'_{10} \otimes M'_{10}, \\
M_{\text{different PBS, NR}} &= M'_{01} \otimes M'_{10} + M'_{10} \otimes M'_{01}, \\
M_{\text{same PBS, NNR}} &= \sum_{n,m=1,2} \left(M'_{0n} \otimes M'_{0m} + M'_{n0} \otimes M'_{m0} \right), \\
M_{\text{different PBS, NNR}} &= \sum_{n,m=1,2} \left(M'_{0n} \otimes M'_{m0} + M'_{0n} \otimes M'_{m0} \right).
\end{aligned} \tag{6.89}$$

6.8.3 Results without coincidence window

To derive formulas for the success probability and post-measurement state, we explicitly calculate the probabilities and post-measurement states of the above POVM elements on the six-qubit space using the symbolic-mathematics Python package SymPy [104]. The corresponding code can be found in the repository holding our simulation code [105]. The results are obtained by first initializing Werner states for each node and applying amplitude-damping channels with loss parameter $1 - p_A$ on the \mathcal{H}_{H_A} and \mathcal{H}_{V_A} subspaces and $1 - p_B$ on the \mathcal{H}_{H_B} and \mathcal{H}_{V_B} subspaces. Then the probability and post-measurement state for both the “same PBS” and “different PBS” measurement outcomes are calculated from this pre-measurement state in both the cases NR and NNR. The result can be written as

$$\begin{aligned}
p_{\text{double click}} &= p_T + p_{F1} + p_{F2} + p_{F3} + p_{F4}, \\
\rho_{\text{double click}} &= q_{\text{em}} \left(p_T |\Psi^\pm\rangle\langle\Psi^\pm| + p_{F1} \frac{|01\rangle\langle 01| + |10\rangle\langle 10|}{2} + p_{F2} \frac{|00\rangle\langle 00| + |11\rangle\langle 11|}{2} \right) \\
&\quad + \left((1 - q_{\text{em}})(p_T + p_{F1} + p_{F2}) + p_{F3} + p_{F4} \right) \frac{1}{4},
\end{aligned} \tag{6.90}$$

where $p_{\text{double click}}$ is the success probability and $\rho_{\text{double click}}$ is the unnormalized post-measurement state. The different constants are defined as

$$\begin{aligned}
q_{\text{em}} &= \frac{1}{9}(4F_{\text{em } A} - 1)(4F_{\text{em } B} - 1), \\
p_T &= \begin{cases} \frac{1}{2} p_A p_B V (1 - p_{\text{dc}})^4 & \text{if NR,} \\ \frac{1}{2} p_A p_B V (1 - p_{\text{dc}})^2 & \text{if NNR,} \end{cases} \\
p_{F1} &= \begin{cases} \frac{1}{2} p_A p_B (1 - V) (1 - p_{\text{dc}})^4 & \text{if NR,} \\ \frac{1}{2} p_A p_B (1 - V) (1 - p_{\text{dc}})^2 & \text{if NNR,} \end{cases} \\
p_{F2} &= \begin{cases} 0 & \text{if NR,} \\ \frac{1}{2} p_A p_B (1 + V) p_{\text{dc}} (1 - p_{\text{dc}})^2 & \text{if NNR,} \end{cases} \\
p_{F3} &= \begin{cases} 2[p_A(1 - p_B) + (1 - p_A)p_B] p_{\text{dc}} (1 - p_{\text{dc}})^3 & \text{if NR,} \\ 2[p_A(1 - p_B) + (1 - p_A)p_B] p_{\text{dc}} (1 - p_{\text{dc}})^2 & \text{if NNR,} \end{cases} \\
p_{F4} &= 4(1 - p_A)(1 - p_B) p_{\text{dc}}^2 (1 - p_{\text{dc}})^2.
\end{aligned} \tag{6.91}$$

Furthermore, the Bell states are defined by

$$|\Psi^\pm\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle). \quad (6.92)$$

The different terms in the equations can be interpreted as corresponding to different possible detection cases, with p_i being the probability of case i occurring, and the density matrix that it multiplies with the state that is created in that case. The different cases are as follows.

- **case T.** This is a “true” heralded success. That is, two photons were detected at the midpoint station (probability $p_A p_B$) in different polarizations (probability $\frac{1}{2}$), and they behaved as indistinguishable photons (i.e. they interfered) (probability V). Finally, there cannot have been dark counts in any of the detectors, except for the detectors at which the photons were detected in the case of non-number-resolving detectors (as this doesn’t change the outcome). The resulting density matrix is one corresponding to the Bell state $|\Psi^\pm\rangle$ (+ for both detections at the same polarizing beam splitter, - for both detections at different polarizing beam splitters).
- **case F1.** This is the first “false” heralded success (i.e. a false positive; a “success” detection pattern is observed without there being a maximally entangled state). Again, two photons arrived at the midpoint station and were detected in different polarizations (probability $\frac{1}{2} p_A p_B$). However, they did not behave as indistinguishable photons (i.e. they did not interfere) (probability $(1 - V)$). Since the photons are in different polarizations, the post-measurement state will be classically anticorrelated $\frac{1}{2}(|01\rangle\langle 01| + |10\rangle\langle 10|)$.
- **case F2.** This is the second “false” heralded success. Two photons arrived at the midpoint station (probability $p_A p_B$) and are detected at the exact same detector. Additionally, a dark count occurs, causing a click pattern that is heralded as a success. For this, both photons need to be detected in the same polarization (probability $\frac{1}{2}$) and end up at the same detector. If they behave as indistinguishable photons (probability V), they will bunch together due to Hong-Ou-Mandel interference and will be guaranteed to go to the same detector. If they do not behave as indistinguishable photons (probability $1 - V$), there is a $\frac{1}{2}$ probability that they happen to go to the same detector. Combining, this gives a factor $V + \frac{1}{2}(1 - V) = \frac{1}{2}(1 + V)$. Since the photons are detected with the same polarization, the post-measurement state will be classically correlated $\frac{1}{2}(|00\rangle\langle 00| + |11\rangle\langle 11|)$. Note that this case cannot occur when detectors are NR, as detecting both photons at the same detector is then heralded as a failure.
- **case F3.** This is the third “false” heralded success. Only one photon arrives at the midpoint station (probability $p_A(1 - p_B) + (1 - p_A)p_B$), and a dark count makes the detector click pattern look like a success. For this, either of the two detectors corresponding to the polarization the photon is not detected in must undergo a dark

count, while the remaining detectors do not undergo dark counts, which occurs with probability $2p_{\text{dc}}(1 - p_{\text{dc}})^n$, where $n = 3$ in case NR and $n = 2$ in case NNR (because then it doesn't matter whether there is a dark count in the detector that detects the photon). There is no information about correlation between the photons, therefore the post-measurement state is maximally mixed.

- **case F4.** This is the fourth and final “false” heralded success. No photons arrive at the midpoint station (probability $(1 - p_A)(1 - p_B)$), and the detector click pattern is created solely by dark counts. Since there are four distinct click patterns resulting in a heralded success, these dark counts occur with probability $4p_{\text{dc}}^2(1 - p_{\text{dc}})^2$. There is no information about correlation between the photons, therefore the post-measurement state is maximally mixed.

Finally, to understand the role of the parameter q_{em} , note that when either of the initial emitter - photon states is maximally mixed instead of entangled, there is no correlation between the emitter and the photon. Therefore, whatever detection event takes place at the midpoint station, no information about correlation between the emitters is revealed. The post-measurement state is thus a maximally mixed state in this case. The probability that both nodes send an entangled photon instead of a maximally mixed state is exactly q_{em} , and thus the probability that the post-measurement state is maximally mixed regardless which of the above cases takes place is $1 - q_{\text{em}}$.

6

6.8.4 Results with coincidence window

When performing the double-click protocol, click patterns can be accepted as a success or instead rejected based on the time at which the two detector clicks are registered. A first reason for this is that each round of the double-click protocol only lasts a finite amount of time. That is, there is a detection time window corresponding to each round, and only clicks occurring during the detection time window can result in a heralded success for that specific round. If there is a nonzero probability that photons are detected outside of the detection time window, e.g. because their wave functions are stretched very long, this can be captured in the model by adjusting the detection probabilities appropriately (p_A and p_B).

However, there can also be a second reason. Sometimes, it is desirable to implement a coincidence time window. In this case, when two clicks occur within the correct detectors and within the detection time window, a success is only heralded if the time between the two clicks is smaller than the coincidence time window. While this lowers the success probability of the double-click protocol, it can increase the Hong-Ou-Mandel visibility V (thereby increasing the fidelity of entangled states created using the protocol).

To account for protocols that implement a coincidence time window, we here introduce three new parameters into our model.

- $p_{\text{ph-ph}}$, the probability that two photon detections that occur within the detection time window occur less than one coincidence time window away from each other.
- $p_{\text{ph-dc}}$, the probability that a photon detection and a dark count that occur within the detection time window occur less than one coincidence time window away from each other.

- $p_{\text{dc-dc}}$, the probability that two dark counts that occur within the detection time window occur less than one coincidence time window away from each other.

These parameters will be functions of photon-detection-time probability-density functions and the coincidence time window (we calculate them for a simplified model of the photon state in Section 6.9) Then, we make the following adjustments to the above results to account for the coincidence time window:

$$\begin{aligned}
 p_T &\rightarrow p_{\text{ph-ph}}p_T, \\
 p_{F1} &\rightarrow p_{\text{ph-ph}}p_{F1}, \\
 p_{F2} &\rightarrow p_{\text{ph-dc}}p_{F2}, \\
 p_{F3} &\rightarrow p_{\text{ph-dc}}p_{F3}, \\
 p_{F4} &\rightarrow p_{\text{dc-dc}}p_{F4}.
 \end{aligned} \tag{6.93}$$

The reason for this is as follows. p_T corresponds to an event where two photons are detected, leading to a heralded success. When using a coincidence time window, the two photons are only close enough in time to lead to a heralded success with probability $p_{\text{ph-ph}}$. The same logic holds for p_{F1} . Probability p_{F3} corresponds to a photon detection and a dark count leading to a heralded success; that now only happens if the photon detection and dark count are within one coincidence time window, which is exactly $p_{\text{ph-dc}}$. And probability p_{F4} corresponds to a heralded success due to two dark counts. These dark counts also should not be separated by too much time, giving a factor $p_{\text{dc-dc}}$. Less straightforward to adjust is p_{F2} in the NNR case. It corresponds to an event where two photons are detected within the same detector, but they are not independently resolved. The probability that the time stamp assigned to this detection is within a coincidence window from a dark count occurring in another detector, may not be exactly $p_{\text{ph-dc}}$. However, we do expect it to be a reasonable approximation, and therefore we use $p_{\text{ph-dc}}$ to avoid introducing a fourth new parameter to the model.

6.9 Effect of detection and coincidence time windows

In the double-click protocol, success is declared only if there are clicks in two detectors that measure different polarization modes. These clicks typically occur at random times, and a prerequisite for success is that certain conditions on the detection times are met. First, in any practical experiment, detection time windows have to be of finite duration. If a click only occurs after the detection time window closes, it is effectively not detected. Thus, success is only declared if two clicks occur within the detection time window. Second, it is sometimes beneficial to also condition success on the time difference between the two clicks. In that case, a success is only declared if the time between the clicks does not exceed the coincidence time window. This can help boost the Hong-Ou-Mandel visibility of the photon interference and thereby increase the fidelity of entangled states.

In Section 6.8, we present a model that allows for the calculation of the success probability of the double-click protocol and the two-qubit state that it creates. The coincidence probabilities between two photons, two dark counts and a photon and a dark count are free parameters in this model, just as the visibility and the photon detection probability. To accurately account for the detection time window and coincidence time window in this

model, these parameters need to be given appropriate values. In this section, we introduce a simplified model for the photon state that allows us to calculate the required values. We use this simplified model to simulate double-click entanglement generation with trapped-ion devices, as described in Section 6.6.6. To the best of our knowledge, this is a novel result.

Definition 7 *Detection time window.* If a detection time window of duration $T > 0$ is used in the double-click protocol, success is only heralded if both detector clicks occur within the time interval $[0, T]$.

Definition 8 *Coincidence time window.* If a coincidence time window of duration $\tau > 0$ is used in the double-click protocol, success is only heralded if the time between both detector clicks does not exceed τ .

Definition 9 *Photon state described by $(p_{em}(t), \psi_{t_0}(t))$.* Let $p_{em}(t)$ be a function such that

$$\int_0^{\infty} dt p_{em}(t) = 1 \quad (6.94)$$

and let $\psi_{t_0}(t)$ be a function such that

$$\int_{t_0}^{\infty} dt |\psi_{t_0}(t)|^2 = 1. \quad (6.95)$$

Then, $p_{em}(t)$ can be interpreted as a probability density function for the photon emission time, and $\psi(t)$ can be interpreted as the temporal wave function of a photon emitted at $t = 0$. The tuple $(p_{em}(t), \psi(t))$ then describes a mixed photon state

$$\rho = \int_0^{\infty} dt_0 p_{em}(t_0) |\psi_{t_0}\rangle\langle\psi_{t_0}| \quad (6.96)$$

where

$$|\psi_{t_0}\rangle = \int_{t_0}^{\infty} dt \psi_{t_0}(t) a_t^\dagger |0\rangle \quad (6.97)$$

with a_t^\dagger the photon's creation operator at time t .

The temporal impurity of a state described by $(p_{em}(t), \psi_{t_0}(t))$ (if $p_{em}(t)$ is not a delta function) can reduce the Hong-Ou-Mandel visibility of photons. The reason for this is that photons that are emitted at very different times have small overlap. If two photons are detected close together, they were probably not emitted at very different times (depending on their distributions). Using a coincidence time window is then effectively applying a temporal purification to the photons, allowing for an increase in visibility.

Definition 10 *Double-exponential photon state (a, b) .* The double-exponential photon state described by (a, b) , where both a and b are constants with dimension time^{-1} , is the photon state described by $(p_{em}(t), \psi_{t_0}(t))$ where

$$p_{em}(t) = a e^{-at} \Theta(t) \quad (6.98)$$

and

$$\psi_{t_0}(t) = \sqrt{2b}e^{-b(t-t_0)}\Theta(t-t_0). \quad (6.99)$$

Here, $\Theta(t)$ is the Heaviside step function. That is, both the emission-time probability density function and the pure photon wavefunctions are one-sided exponentials.

In this section, we model all photons emitted by processing nodes as having a double-exponential state. The pure wave functions of photons emitted by spontaneous decay of an excited state to a ground state in a two-level system are described well as one-sided exponentials [106]. An example of a system where photons are emitted this way is NV centers [48]. Similarly, pure wave functions of photons emitted using cavity-enhanced Raman transitions (using a constant Rabi pulse), as is the case for the trapped-ion systems we study in this chapter, also look approximately exponential [107]. We note that such trapped-ion systems are exactly the use case in this chapter for the simplified model presented here. For solid-state sources such as color centers, temporal impurity of photons is not a limiting factor [106]. However, for cavity-enhanced Raman transitions, off-resonant scattering causes the photon to only be emitted at a random time after a trajectory through the ion-state manifold [86, 107]. We model the resulting temporal impurity using the function $p_{\text{em}}(t)$. We note that we do not expect this function to be exponential for cavity-enhanced Raman transitions. For instance, the function should include a $\delta(0)$ delta-function contribution to account for the probability that not a single off-resonant scattering takes place. However, in the toy model presented here, we will assume $p_{\text{em}}(t)$ is a one-sided exponential so that we have a model with a small number of parameters in which exact closed-form expressions can be obtained for the relevant quantities. As shown in Section 6.6.6, this model can be fitted well to experimental data for interference between photons emitted by ion-cavity systems.

Lemma 7 *Detection-time probability density function.* Consider the case where a photon with double-exponential state (a, b) is emitted directly on a photon detector. Assume this photon detector is perfect except that it has a possibly nonunit detection efficiency η (with a flat response). The probability density function for the photon being detected at time t is given by

$$p(t) = \frac{2ab\eta}{a-2b} (e^{-2bt} - e^{-at})\Theta(t). \quad (6.100)$$

This probability density function may be subnormalized, as it is also possible that no photon is detected.

Proof: A perfect detector implements a POVM with operators $E_t = a_t^\dagger |0\rangle\langle 0| a_t$. Instead, a detector with efficiency factor η implements a POVM with operators $E'_t = \eta E_t$ and $F = 1 - \eta$, where F corresponds to no photon detection taking place. The probability density that the photon is detected at time t is then the probability density corresponding to the POVM operator E'_t , given by

$$p(t) = \text{Tr}(E'_t \rho). \quad (6.101)$$

For a photon state described by $(p_{\text{em}}(t), \psi_{t_0}(t))$, the density matrix is

$$\rho = \int_0^\infty dt_0 \int_{t_0}^\infty dt_1 \int_{t_0}^\infty dt_2 p_{\text{em}}(t_0) \psi_{t_0}(t_1) \psi_{t_0}^*(t_1) a_{t_1}^\dagger |0\rangle\langle 0| a_{t_2} \quad (6.102)$$

This can be evaluated using the cyclic property of the trace to give

$$\begin{aligned}
 p(t) &= \eta \int_0^\infty dt_0 \int_{t_0}^\infty dt_1 \int_{t_0}^\infty dt_2 p_{\text{em}}(t_0) \psi_{t_0}(t_1) \psi_{t_0}^*(t_1) \langle 0 | a_t a_{t_1}^\dagger | 0 \rangle \langle 0 | a_{t_2} a_t^\dagger | 0 \rangle \\
 &= \eta \int_0^\infty dt_0 \int_{t_0}^\infty dt_1 \int_{t_0}^\infty dt_2 p_{\text{em}}(t_0) \psi_{t_0}(t_1) \psi_{t_0}^*(t_1) \delta(t - t_1) \delta(t - t_2) \\
 &= \eta \int_0^\infty dt_0 p_{\text{em}}(t_0) |\psi_{t_0}(t)|^2.
 \end{aligned} \tag{6.103}$$

For a double-exponential photon state (a, b) , this becomes

$$\begin{aligned}
 p(t) &= 2ab\eta e^{-2bt} \int_0^\infty dt_0 e^{-(a-2b)t_0} \Theta(t - t_0) \\
 &= 2ab\eta e^{-2bt} \Theta(t) \int_0^t dt_0 e^{-(a-2b)t_0} \\
 &= 2ab\eta e^{-2b} \Theta(t) \frac{1}{a-2b} (1 - e^{-(a-2b)t}) \\
 &= \frac{2ab\eta}{a-2b} (e^{-2bt} - e^{-at}) \Theta(t).
 \end{aligned} \tag{6.104}$$

□

Definition 11 *Coincidence probability.* When using a detection time window T and coincidence time window of τ in the double-click protocol, the coincidence probability is the probability that given that there are two clicks within the detection time window, the clicks are also within one coincidence time window.

Our goal now is to find the coincidence probability for two double-exponential photons. This requires us to calculate the probability that two photons arrive within a time τ of one another, conditioned on each of the photons being successfully detected within the time interval $[0, T]$. To this end, we calculate the probability density function for the detection time of a double-exponential photon conditioned on the photon being successfully detected. This requires us to calculate the detection probability of the photon, i.e., the probability that it is successfully detected within the detection time window. The detection probability is also an important result in itself, as it is required by the model presented in Section 6.8 (it takes the role of p_A and p_B in this model).

Theorem 5 *Detection probability.* If a detection time window of duration T is used, then the probability that a photon with double-exponential state (a, b) is detected within the time window is given by

$$p_{\text{det}}(T) = \eta \left[1 - \frac{a}{a-2b} e^{-2bT} + \frac{2b}{a-2b} e^{-aT} \right]. \tag{6.105}$$

Proof: $p_{\text{det}}(T)$ is given by the probability that the photon is detected in the time interval $[0, T]$. This probability can be calculated from the probability density function

$$\begin{aligned}
 p_{\text{det}}(T) &= \int_0^T dt p(t) \\
 &= \frac{2ab\eta}{a-2b} \int_0^T dt (e^{-2bt} - e^{-at}) \\
 &= \frac{2ab\eta}{a-2b} \left[\frac{1}{2b} (1 - e^{-2bT}) - \frac{1}{a} (1 - e^{-aT}) \right] \\
 &= \frac{\eta}{a-2b} \left[a(1 - e^{-2bT}) - 2b(1 - e^{-aT}) \right] \\
 &= \eta \left[1 - \frac{a}{a-2b} e^{-2bT} + \frac{2b}{a-2b} e^{-aT} \right].
 \end{aligned} \tag{6.106}$$

□

Corollary 1 *When no detection time window is used, i.e., when the duration of the detection time window $T \rightarrow \infty$, then the photon detection probability is equal to the detector's detection efficiency η .*

Proof: When we take $T \rightarrow \infty$ in equation (6.105), we find $p_{\text{det}}(T) \rightarrow \eta$. □ This corresponds to the situation when the entire photon is within the detection time window and the only reason why the photon would not be detected is detector inefficiency. One can think of $p_{\text{det}}(T)/\eta$ as the “additional efficiency factor” due to not capturing the entire photon in the detection time window.

Lemma 8 *Conditional detection-time probability density function. Consider the case where a photon with double-exponential state (a, b) is emitted directly on a photon detector. Assume this photon detector is perfect except that it has a possibly nonunit detection efficiency η (with a flat response). The probability density function for the photon being detected at time t , if the photon is in the double-exponential state (a, b) , is given by*

$$p_T(t) = \Theta(t)\Theta(T-t) \frac{p(t)}{p_{\text{det}}(T)}. \tag{6.107}$$

Unlike $p(t)$, this probability density function is always normalized.

Proof: Let X be the continuous random variable corresponding to the detection time of the photon. Let it take the value -1 if no photon is detected, such that the corresponding probability density function $f_X(x)$ is normalized and X is well-defined as a random variable. It follows from Corollary 1 that the probability that this happens is $1 - \eta$. Therefore, the probability density function can then be written as

$$f_X(x) = (1 - \eta)\delta(x + 1) + \Theta(x)p(x). \tag{6.108}$$

Additionally, we define a discrete random variable Y_T which takes the value 1 if the photon is detected within the detection time window T and 0 if not. By Theorem 5, the probability distribution of Y_T is given by

$$f_{Y_T}(y) = \begin{cases} p_{\text{det}}(T) & \text{if } y = 1 \\ 1 - p_{\text{det}}(T) & \text{if } y = 0. \end{cases} \quad (6.109)$$

Note that X and Y_T are not independent random variables. Now, the conditional probability density function that we are interested in is

$$p_T(x) \equiv f_{X|Y_T}(x, 1) = \frac{f_{X,Y_T}(x, 1)}{f_{Y_T}(1)} = \frac{f_{X,Y_T}(x, 1)}{p_{\text{det}}(T)}. \quad (6.110)$$

Here, $f_{X,Y_T}(x, y)$ is the mixed joint density of the continuous random variable X and the discrete random variable Y_T . When X takes a value between 0 and T , then Y_T takes the value 1 with unit probability. Otherwise, Y_T takes the value 0 with unit probability. Therefore,

$$f_{X,Y_T}(x, 1) = \Theta(x)\Theta(T-x)f_X(x) = \Theta(x)\Theta(T-x)p(x). \quad (6.111)$$

Substituting this into equation (6.110) allows us then finally to write

$$p_T(t) = \Theta(t)\Theta(T-t)\frac{p(t)}{p_{\text{det}}(T)}. \quad (6.112)$$

□

Theorem 6 *Coincidence probability of two photons. The coincidence probability for two photon detections, if both photons are in the double-exponential state (a, b) , a detection time window of duration T is used and the coincidence time window of duration τ is used, is given by*

$$\begin{aligned} p_{\text{ph-ph}}(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 &= \frac{a^2}{a^2 - 4b^2} (1 - e^{-2b\tau}) - \frac{4b^2}{a^2 - 4b^2} (1 - e^{-a\tau}) \\ &+ \frac{a^2}{(a-2b)^2} (1 - e^{2b\tau}) e^{-4bT} + \frac{4b^2}{(a-2b)^2} (1 - e^{a\tau}) e^{-2aT} \\ &- \frac{4ab}{(a-2b)^2} \left(1 - \frac{ae^{2b\tau} + 2be^{a\tau}}{a+2b} \right) e^{-(a+2b)T}. \end{aligned} \quad (6.113)$$

Proof: By definition, the coincidence probability is given by

$$p_{\text{ph-ph}}(T, \tau) = \iint_{|t_1 - t_2| \leq \tau} dt_1 dt_2 p_T(t_1) p_T(t_2). \quad (6.114)$$

By Lemma 8, this implies

$$p_{\text{ph-ph}}(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 = \frac{1}{\eta^2} \iint_{|t_1 - t_2| \leq \tau} dt_1 dt_2 p(t_1) p(t_2). \quad (6.115)$$

The region of integration is $|t_1 - t_2| \leq \tau$, i.e., $-\tau \leq t_1 - t_2 \leq \tau$. The integrand is symmetric under the interchange of t_1 and t_2 . Therefore, the region $0 \leq t_1 - t_2 \leq \tau$ will give exactly the same contribution as $-\tau \leq t_1 - t_2 \leq 0$. This has the following physical interpretation: the probability of photon 2 arriving a time Δt after photon 1 is the same as the probability of photon 1 arriving Δt after photon 2. This can be used to simplify the integral somewhat, giving

$$p_{\text{ph-ph}}(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 = \frac{2}{\eta^2} \iint_{0 \leq t_1 - t_2 \leq \tau} dt_1 dt_2 p(t_1) p(t_2). \quad (6.116)$$

It follows from Lemma 7 that each $p(t)$ carries an overall factor $\Theta(t)\Theta(T)$. This can be absorbed into the integration limits to give

$$\begin{aligned} p_{\text{ph-ph}}(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 &= \frac{2}{\eta^2} \int_0^T dt_1 \int_{t_1}^{\min(t_1+\tau, T)} dt_2 p(t_1) p(t_2) \\ &= 2 \left(\frac{1}{\eta^2} \int_0^{T-\tau} dt_1 \int_{t_1}^{t_1+\tau} dt_2 p(t_1) p(t_2) + \frac{1}{\eta^2} \int_{T-\tau}^T dt_1 \int_{t_1}^T dt_2 p(t_1) p(t_2) \right). \end{aligned} \quad (6.117)$$

We calculate these two integrals one by one, using Lemma 7. The first is

$$\begin{aligned} &\left(\frac{a-2b}{2ab} \right)^2 \left(\frac{1}{\eta^2} \int_0^{T-\tau} dt_1 \int_{t_1}^{t_1+\tau} dt_2 p(t_1) p(t_2) \right) \\ &= \int_0^{T-\tau} dt_1 (e^{-2bt_1} - e^{-at_1}) \int_{t_1}^{t_1+\tau} (e^{-2bt_2} - e^{-at_2}) \\ &= \int_0^{T-\tau} dt_1 (e^{-2bt_1} - e^{-at_1}) \left(\frac{1}{2b} e^{-2bt_1} (1 - e^{-2b\tau}) - \frac{1}{a} e^{-at_1} (1 - e^{-a\tau}) \right) \\ &= \frac{1 - e^{-2b\tau}}{2b} \int_0^{T-\tau} dt_1 (e^{-4bt_1} - e^{-(a+2b)t_2}) - \frac{1 - e^{-a\tau}}{a} \int_0^{T-\tau} dt_1 (e^{-(a+2b)t_1} - e^{-2at_1}) \\ &= \frac{1 - e^{-2b\tau}}{2b} \left(\frac{1 - e^{4b\tau} e^{-4bT}}{4b} - \frac{1 - e^{(a+2b)\tau} e^{-(a+2b)T}}{a+2b} \right) \\ &\quad - \frac{1 - e^{-a\tau}}{a} \left(\frac{1 - e^{(a+2b)\tau} e^{-(a+2b)T}}{a+2b} - \frac{1 - e^{2a\tau} e^{-2aT}}{2a} \right) \\ &= \frac{a-2b}{a+2b} \left(\frac{1 - e^{-2b\tau}}{8b^2} - \frac{1 - e^{-a\tau}}{2a^2} \right) + \frac{-e^{2a\tau} + e^{a\tau}}{2a^2} e^{-2aT} \\ &\quad + \frac{e^{2b\tau} - e^{4b\tau}}{8b^2} e^{-4bT} + \frac{1}{a+2b} \left(\frac{e^{(a+2b)\tau} - e^{a\tau}}{2b} + \frac{e^{(a+2b)\tau} - e^{2b\tau}}{a} \right) e^{-(a+2b)T}. \end{aligned} \quad (6.118)$$

In the last step, terms with the same exponents of T were collected. Resolving the prefactor gives

$$\begin{aligned}
& \frac{1}{\eta^2} \int_0^{T-\tau} dt_1 \int_{t_1}^{t_1+\tau} dt_2 p(t_1) p(t_2) \\
&= \frac{a^2}{2(a^2 - 4b^2)} (1 - e^{-2b\tau}) - \frac{2b^2}{2(a^2 - 4b^2)} (1 - e^{-a\tau}) + \frac{2b^2}{(a-2b)^2} (-e^{2a\tau} + e^{a\tau}) e^{-2aT} \\
&+ \frac{a^2}{2(a-2b)^2} (e^{2b\tau} - e^{4b\tau}) e^{-4bT} \\
&+ \frac{2ab}{(a-2b)^2(a+2b)} (a(e^{(a+2b)\tau} - e^{a\tau}) + 2b(e^{(a+2b)\tau} - e^{2b\tau})) e^{-(a+2b)T}.
\end{aligned} \tag{6.119}$$

The second term is

$$\begin{aligned}
& \left(\frac{a-2b}{2ab}\right)^2 \left(\frac{1}{\eta^2} \int_{T-\tau}^T dt_1 \int_{t_1}^T dt_2 p(t_1) p(t_2)\right) \\
&= \int_{T-\tau}^T dt_1 (e^{-2bt_1} - e^{-at_1}) \int_{t_1}^T (e^{-2bt_2} - e^{-at_2}) \\
&= \int_{T-\tau}^T dt_1 (e^{-2bt_1} - e^{-at_1}) \left(\frac{1}{2b}(e^{-2bt_1} - e^{-2bT}) - \frac{1}{a}(e^{-at_1} - e^{-aT})\right) \\
&= \frac{1}{2ab} (2be^{-aT} - ae^{-2bT}) \int_{T-\tau}^T dt_1 (e^{-2bt_1} - e^{-at_1}) + \frac{1}{2b} \int_{T-\tau}^T dt_1 e^{-4bt_1} \\
&+ \frac{1}{a} \int_{T-\tau}^T dt_1 e^{-2at_1} - \frac{a+2b}{2ab} \int_{T-\tau}^T dt_1 e^{-(a+2b)t_1} \\
&= \frac{1}{2ab} (2be^{-aT} - ae^{-2bT}) \left[\frac{e^{-2bT}}{2b} (e^{2b\tau} - 1) - \frac{e^{-aT}}{a} (e^{a\tau} - 1) \right] + \frac{e^{-4bT}}{8b^2} (e^{4b\tau} - 1) \\
&+ \frac{e^{-2aT}}{2a^2} (e^{2a\tau} - 1) - \frac{e^{-(a+2b)T}}{2ab} (e^{(a+2b)\tau} - 1) \\
&= \frac{1}{2a^2} (e^{2a\tau} - 2e^{a\tau} + 1) e^{-2aT} + \frac{1}{8b^2} (e^{4b\tau} - 2e^{2b\tau} + 1) e^{-4bT} \\
&+ \frac{1}{2ab} (e^{a\tau} + e^{2b\tau} - e^{(a+2b)\tau} - 1) e^{-(a+2b)T}.
\end{aligned} \tag{6.120}$$

Again resolving the prefactor, we find

$$\begin{aligned}
& \frac{1}{\eta^2} \int_{T-\tau}^T dt_1 \int_{t_1}^T dt_2 p(t_1) p(t_2) \\
&= \frac{2b^2}{(a-2b)^2} (e^{2a\tau} - 2e^{a\tau} + 1) e^{-2aT} + \frac{a^2}{2(a-2b)^2} (e^{4b\tau} - 2e^{2b\tau} + 1) e^{-4bT} \\
&+ \frac{2ab}{(a-2b)^2} (e^{a\tau} + e^{2b\tau} - e^{(a+2b)\tau} - 1) e^{-(a+2b)T}.
\end{aligned} \tag{6.121}$$

Finally, substituting equations (6.119) and (6.121) into equation (6.117) yields equation (6.113). \square

Theorem 7 *Coincidence probability of two dark counts. The coincidence probability for two detector dark counts if the detection time window is T and the coincidence time window is τ is given by*

$$p_{dc-dc}(T, \tau) = 1 - \left(\frac{T - \tau}{T} \right)^2, \quad (6.122)$$

assuming that dark counts occur uniformly throughout the detection time window.

Proof: Given that there is a dark count within the detection time window T , the probability density function for the time at which it occurs is given by

$$d_T(t) = \frac{1}{T} \Theta(t) \Theta(T - t) \quad (6.123)$$

because we assume the dark counts are uniformly distributed. The coincidence probability is then given by

$$\begin{aligned} p_{dc-dc}(T, \tau) &= \iint_{|t_1 - t_2| \leq \tau} dt_1 dt_2 d_T(t_1) d_T(t_2) \\ &= \frac{2}{T^2} \int_0^T dt_1 \int_{t_1}^{\min(t_1 + \tau, T)} dt_2 \\ &= \frac{2}{T^2} \left(\int_0^{T-\tau} dt_1 \int_{t_1}^{t_1 + \tau} dt_2 + \int_{T-\tau}^T dt_1 \int_{t_1}^T dt_2 \right) \\ &= \frac{2}{T^2} \left(\int_0^{T-\tau} dt_1 \tau + \int_{T-\tau}^T dt_1 (T - t_1) \right) \\ &= \frac{2}{T^2} \left((T - \tau)\tau + T\tau - \frac{1}{2}T^2 + \frac{1}{2}(T - \tau)^2 \right) \\ &= \frac{1}{T^2} (2T\tau - \tau^2) \\ &= 1 - \left(\frac{T - \tau}{T} \right)^2. \end{aligned} \quad (6.124)$$

\square

Theorem 8 *Coincidence probability of a photon and a dark count. The coincidence probability for one photon detection and one dark count, if the photon is in the double-exponential state (a, b) , the detection time window is T and the coincidence time window is τ is given by*

$$\begin{aligned} p_{ph-dc}(T, \tau) \frac{p_{det}(T)}{\eta} &= \frac{a}{2b(a - 2b)T} \left[1 + 2b\tau - e^{-2b\tau} + e^{-2bT} (1 - 2b\tau - e^{2b\tau}) \right] \\ &\quad - \frac{2b}{a(a - 2b)T} \left[1 + a\tau - e^{-a\tau} + e^{-aT} (1 - a\tau - e^{a\tau}) \right]. \end{aligned} \quad (6.125)$$

assuming that dark counts occur uniformly throughout the detection time window.

Proof: For the photon, we again have the probability density function $p_T(t)$ as given by Lemma 8, while for the dark count we have the probability density function $d_T(t)$ as given by equation (6.123). The coincidence probability is then given by

$$p_{\text{ph-dc}}(T, \tau) = \iint_{|t_1 - t_2| \leq \tau} dt_1 dt_2 p_T(t_1) d_T(t_2). \quad (6.126)$$

When calculating the other coincidence probabilities, we were able to use a symmetry argument to simplify the integral. However, because $p_T(t) \neq d_T(t)$, we are unable to do so here. Assuming for the moment that $\tau \leq \frac{1}{2}T$ and the fact that both probability density functions are proportional to $\Theta(t)\Theta(T-t)$, we can write

$$p_{\text{ph-dc}}(T, \tau) = \left(\int_0^\tau dt_1 \int_0^{t_1+\tau} dt_2 + \int_\tau^{T-\tau} dt_1 \int_{t_1-\tau}^{t_1+\tau} dt_2 + \int_{T-\tau}^T dt_1 \int_{t_1-\tau}^T dt_2 \right) p_T(t_1) d_T(t_2). \quad (6.127)$$

This becomes

$$\begin{aligned} & p_{\text{ph-dc}}(T, \tau) \frac{p_{\text{det}}(T)T}{\eta} \frac{a-2b}{2ab} \\ &= \left(\int_0^\tau dt_1 \int_0^{t_1+\tau} dt_2 + \int_\tau^{T-\tau} dt_1 \int_{t_1-\tau}^{t_1+\tau} dt_2 + \int_{T-\tau}^T dt_1 \int_{t_1-\tau}^T dt_2 \right) (e^{-2bt_1} - e^{-at_1}) \\ &= \left(\int_0^\tau dt_1 (t_1 + \tau) + \int_\tau^{T-\tau} dt_1 2\tau + \int_{T-\tau}^T dt_1 (T + \tau - t_1) \right) (e^{-2bt_1} - e^{-at_1}). \end{aligned} \quad (6.128)$$

We calculate these three terms individually. Before doing this, we note that we can use integration by parts to calculate

$$\int_x^y dt e^{-zt} t = -\frac{1}{z} [te^{-zt}]_{t=x}^{t=y} + \frac{1}{z} \int_x^y e^{-zt} dt = \left[\frac{e^{-zt}}{z} \left(t + \frac{1}{z} \right) \right]_{t=y}^{t=x}. \quad (6.129)$$

Then, the first term becomes

$$\begin{aligned} \int_0^\tau dt_1 (t_1 + \tau) (e^{-2bt_1} - e^{-at_1}) &= \left[\left(\tau + \frac{1}{2b} + t_1 \right) \frac{e^{-2bt_1}}{2b} \right]_0^\tau - \left[\left(\tau + \frac{1}{a} + t_1 \right) \frac{e^{-at_1}}{a} \right]_0^\tau \\ &= \frac{1}{4b^2} + \frac{\tau}{2b} - \frac{1}{a^2} - \frac{\tau}{a} - \left(\frac{1}{2b} + 2\tau \right) \frac{e^{-2b\tau}}{2b} + \left(\frac{1}{a} + 2\tau \right) \frac{e^{-a\tau}}{a}. \end{aligned} \quad (6.130)$$

The second yields

$$2\tau \int_\tau^{T-\tau} dt_1 (e^{-2bt_1} - e^{-at_1}) = 2\tau \frac{e^{-2b\tau}}{2b} - 2\tau \frac{e^{2b(\tau-T)}}{2b} - 2\tau \frac{e^{-a\tau}}{a} + 2\tau \frac{e^{a(\tau-T)}}{a}. \quad (6.131)$$

The final one yields

$$\begin{aligned}
 \int_{T-\tau}^T dt_1 (T + \tau - t_1) (e^{-2bt_1} - e^{-at_1}) &= \left[\left(T + \tau - \frac{1}{2b} - t_1 \right) \frac{e^{-2bt_1}}{2b} \right]_{T-\tau}^{T-\tau} \\
 &\quad - \left[\left(T + \tau - \frac{1}{a} - t_1 \right) \frac{e^{-at_1}}{a} \right]_{T-\tau}^{T-\tau} \\
 &= \left[\frac{1}{2b} - \tau + e^{2b\tau} \left(2\tau - \frac{1}{2b} \right) \right] \frac{e^{-2bT}}{2b} \\
 &\quad - \left[\frac{1}{a} - \tau + e^{a\tau} \left(2\tau - \frac{1}{a} \right) \right] \frac{e^{-aT}}{a}.
 \end{aligned} \tag{6.132}$$

We note that the second term cancels fully against the first and the third. When adding all together, we find

$$\begin{aligned}
 p_{\text{ph-dc}}(T, \tau) \frac{p_{\text{det}}(T) T}{\eta} \frac{a-2b}{2ab} &= \frac{1}{2b} \left[\frac{1}{2b} + \tau - \frac{e^{-2b\tau}}{2b} + e^{-2bT} \left(\frac{1}{2b} - \tau - \frac{e^{2b\tau}}{2b} \right) \right] \\
 &\quad - \frac{1}{a} \left[\frac{1}{a} + \tau - \frac{e^{-a\tau}}{a} + e^{-aT} \left(\frac{1}{a} - \tau - \frac{e^{a\tau}}{a} \right) \right]
 \end{aligned} \tag{6.133}$$

and thus

$$\begin{aligned}
 p_{\text{ph-dc}}(T, \tau) \frac{p_{\text{det}}(T)}{\eta} &= \frac{a}{2b(a-2b)T} \left[1 + 2b\tau - e^{-2b\tau} + e^{-2bT} (1 - 2b\tau - e^{2b\tau}) \right] \\
 &\quad - \frac{2b}{a(a-2b)T} \left[1 + a\tau - e^{-a\tau} + e^{-aT} (1 - a\tau - e^{a\tau}) \right].
 \end{aligned} \tag{6.134}$$

This procedure can be repeated when making the assumption $\tau \geq \frac{T}{2}$. In that case, the exact same formula is found. Therefore, equation (6.125) is valid for any $0 \leq \tau \leq T$. \square

Definition 12 *Visibility.* When using a detection time window T and coincidence time window of τ in the double-click protocol, the Hong-Ou-Mandel visibility is defined as

$$V(T, \tau) = 1 - \frac{P(\text{two photons detected at different detectors} \mid \text{same mode})}{P(\text{two photons detected at different detectors} \mid \text{different modes})}. \tag{6.135}$$

Here, $P(\text{two photons detected at different detectors} \mid \text{same mode})$ is the probability that if both incoming photons are in the same mode, they will be both be detected, and these detection events occur at different detectors. On the other hand,

$P(\text{two photons detected at different detectors} \mid \text{different modes})$ is the probability that if both incoming photons are in different modes (e.g., different polarizations), they will both be detected, and these detection events occur at different detectors. Dark counts are not considered photon detections for the definitions of these probabilities.

Note that when the two photons are in the same mode, they are able to interfere. Then, if the two photons are pure and have the same temporal profile, they are perfectly indistinguishable and will never be detected at different detectors because of the Hong-Ou-Mandel effect [74]. On the other hand, if the two photons are in different modes (e.g., different polarizations), they are not able to interfere. We note that the definition here given is in line with the definition given for the Hong-Ou-Mandel visibility in the main text.

Theorem 9 *Visibility.* The Hong-Ou-Mandel visibility for a double-click protocol with two photons that are both in the double-exponential state (a, b) is given by

$$\begin{aligned}
 V(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 p_{\text{ph-ph}}(T, \tau) &= \frac{a}{a+2b} (1 - e^{-2b\tau}) + \frac{2ab^2}{(a-2b)^2(a-b)} (1 - e^{2(a-b)\tau}) e^{-2aT} \\
 &+ \frac{a^2}{(a-2b)^2} (1 - e^{2b\tau}) e^{-4bT} \\
 &- \frac{16ab^2}{(a-2b)^2(a+2b)} (1 - e^{a\tau}) e^{-(a+2b)T}.
 \end{aligned} \tag{6.136}$$

6

Proof: First, we evaluate $P(\text{two photons detected at different detectors} \mid \text{different modes})$. Because the photons do not interfere, this probability is just the probability that both photons are detected within the detection time window and within one coincidence time window, multiplied by a factor of $\frac{1}{2}$ as the probability for both photons going to different detectors is the same as the probability for both photons going to the same detector. The probability for a single photon falling within the detection time window is the detection probability (Theorem 5), and the probability of both photons being detected within a single coincidence time window is the photon-photon coincidence probability (Theorem 6). Thus,

$$P(\text{two photons detected at different detectors} \mid \text{different modes}) = \frac{1}{2} p_{\text{det}}(T)^2 p_{\text{ph-ph}}(T, \tau). \tag{6.137}$$

The second probability can be evaluated as [86]

$$\begin{aligned}
 &P(\text{two photons detected at different detectors} \mid \text{same mode}) \\
 &= \frac{\eta^2}{4} \int_0^\infty dt_1 \int_0^\infty dt_2 \iint_{|t'_1 - t'_2| \leq \tau} dt'_1 dt'_2 p_{\text{em}}(t_1) p_{\text{em}}(t_2) \left| \psi_{t_1}(t'_1) \psi_{t_2}(t'_2) - \psi_{t_1}(t'_2) \psi_{t_2}(t'_1) \right|^2.
 \end{aligned} \tag{6.138}$$

From combining equations (6.103), (6.107) and (6.114), we see that

$$\eta^2 \int_0^\infty dt_1 \int_0^\infty dt_2 \iint_{|t'_1 - t'_2| \leq \tau} p_{\text{em}}(t_1) p_{\text{em}}(t_2) \left| \psi_{t_1}(t'_1) \right|^2 \left| \psi_{t_2}(t'_2) \right|^2 = p_{\text{det}}(T)^2 p_{\text{ph-ph}}(T, \tau). \tag{6.139}$$

We use this, together with the fact that for double-exponential photons it holds that $\psi_0^*(t) = \psi_0(t)$, to find

$$P(\text{two photons detected at different detectors} \mid \text{same mode}) = \frac{1}{2} p_{\text{det}}(T)^2 p_{\text{ph-ph}}(T, \tau) - \frac{\eta^2}{2} \int_0^\infty dt_1 \int_0^\infty dt_2 \iint_{|t'_1 - t'_2| \leq \tau} dt'_1 dt'_2 p_{\text{em}}(t_1) p_{\text{em}}(t_2) \psi_{t_1}(t'_1) \psi_{t_1}(t'_2) \psi_{t_2}(t'_2) \psi_{t_2}(t'_1). \quad (6.140)$$

We can then work out equation (6.135) to find

$$V(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 p_{\text{ph-ph}}(T, \tau) = \int_0^\infty dt_1 \int_0^\infty dt_2 \int_0^T dt'_1 \int_0^T dt'_2 \Theta(|t'_1 - t'_2| - \tau) p_{\text{em}}(t_1) p_{\text{em}}(t_2) \psi_{t_1}(t'_1) \psi_{t_1}(t'_2) \psi_{t_2}(t'_2) \psi_{t_2}(t'_1). \quad (6.141)$$

The integrand is symmetric under interchange of t'_1 and t'_2 . This allows us to consider only the region $0 \leq t'_2 - t'_1 \leq \tau$, giving

$$V(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 p_{\text{ph-ph}}(T, \tau) = 2 \int_0^\infty dt_1 \int_0^\infty dt_2 \int_0^T dt'_1 \int_{t'_1}^{\min(t'_1 + \tau, T)} dt'_2 p_{\text{em}}(t_1) p_{\text{em}}(t_2) \psi_{t_1}(t'_1) \psi_{t_1}(t'_2) \psi_{t_2}(t'_2) \psi_{t_2}(t'_1). \quad (6.142)$$

Now, we notice that each $\psi_{t_0}(t)$ is proportional to $\Theta(t - t_0)$. This can be absorbed into the limit of integration for t_1 and t_2 , yielding

$$\begin{aligned} V(T, \tau) \left(\frac{p_{\text{det}}(T)}{\eta} \right)^2 p_{\text{ph-ph}}(T, \tau) &= 2 \int_0^T dt'_1 \int_0^{t'_1} dt_1 \int_0^{t'_1} dt_2 \int_{t'_1}^{\min(t'_1 + \tau, T)} dt'_2 p_{\text{em}}(t_1) p_{\text{em}}(t_2) \psi_{t_1}(t'_1) \psi_{t_1}(t'_2) \psi_{t_2}(t'_2) \psi_{t_2}(t'_1) \\ &= 8a^2 b^2 \int_0^T dt'_1 e^{-2bt'_1} \int_0^{t'_1} dt_1 e^{-(a-2b)t_1} \int_0^{t'_1} dt_2 e^{-(a-2b)t_2} \int_{t'_1}^{\min(t'_1 + \tau, T)} dt'_2 e^{-2bt'_2} \\ &= \frac{4a^2 b}{(a-2b)^2} \int_0^T dt'_1 e^{-2bt'_1} (e^{-2bt'_1} - e^{-2b\min(t'_1 + \tau, T)}) (1 - e^{-(a-2b)t'_1})^2 \\ &= \frac{4a^2 b}{(a-2b)^2} \left[\int_0^T e^{-2bt} - e^{-2b\tau} \int_0^{T-\tau} e^{-2bt} - e^{-2bT} \int_{T-\tau}^T \right] (e^{-2bt} - 2e^{-at} + e^{-2(a-b)t}) dt. \end{aligned} \quad (6.143)$$

In the last step, we split up the integration region into a part where $t'_1 + \tau$ is smaller and a part where T is smaller. Furthermore, for brevity, we renamed t'_1 to t . We first calculate

the first integral to find

$$\begin{aligned} \int_0^T dt \left(e^{-4bt} - 2e^{-(a+2b)t} + e^{-2at} \right) &= \frac{1}{4b}(1 - e^{-4bT}) - \frac{2}{a+2b}(1 - e^{-(a+2b)T}) + \frac{1}{2a}(1 - e^{-2aT}) \\ &= \frac{1}{4b} - \frac{2}{a+2b} + \frac{1}{2a} - \frac{1}{2a}e^{-2aT} - \frac{1}{4b}e^{-4bT} + \frac{2}{a+2b}e^{-(a+2b)T}. \end{aligned} \quad (6.144)$$

The second yields

$$\begin{aligned} e^{-2b\tau} \int_0^{T-\tau} dt \left(-e^{-4bt} + 2e^{-(a+2b)t} - e^{-2at} \right) \\ = e^{-2b\tau} \left(-\frac{1}{4b}(1 - e^{-4b(T-\tau)}) + \frac{2}{a+2b}(1 - e^{-(a+2b)(T-\tau)}) - \frac{1}{2a}(1 - e^{-2a(T-\tau)}) \right) \\ = \left(-\frac{1}{4b} + \frac{2}{a+2b} - \frac{1}{2a} \right) e^{-2b\tau} + \frac{e^{2(a-b)\tau}}{2a} e^{-2aT} + \frac{e^{2b\tau}}{4b} e^{-4bT} - \frac{2e^{a\tau}}{a+2b} e^{-(a+2b)T}. \end{aligned} \quad (6.145)$$

The final one yields

$$\begin{aligned} e^{-2bT} \int_{T-\tau}^T dt \left(-e^{-2bt} + 2e^{-at} - e^{-2(a-b)t} \right) \\ = -\frac{e^{2b\tau} - 1}{2b} e^{-4bT} + \frac{2(e^{a\tau} - 1)}{a} e^{-(a+2b)T} - \frac{e^{2(a-b)\tau} - 1}{2(a-b)} e^{-2aT}. \end{aligned} \quad (6.146)$$

Now, it is just a matter of adding these three terms together and taking the prefactor into account. We collect terms separately for each different exponent with a T . The part of the expression that is independent of T yields

$$\begin{aligned} (1 - e^{-2b\tau}) \frac{4a^2b}{(a-2b)^2} \left(\frac{1}{4b} - \frac{2}{a+2b} + \frac{1}{2a} \right) \\ = \frac{a}{(a-2b)^2} \left(a - \frac{8ab}{a+2b} + 2b \right) \\ = (1 - e^{-2b\tau}) \frac{a}{(a-2b)^2(a+2b)} (a(a+2b) - 8ab + 2b(a+2b)) \\ = (1 - e^{-2b\tau}) \frac{a}{a+2b}. \end{aligned} \quad (6.147)$$

For the terms proportional to e^{-2aT} we find

$$\begin{aligned} \frac{4a^2b}{(a-2b)^2} (1 - e^{2(a-b)\tau}) \left(\frac{1}{2a} - \frac{1}{2(a-b)} \right) e^{-2aT} &= \frac{2ab}{(a-2b)^2} (1 - e^{2(a-b)\tau}) \left(\frac{a}{a-b} - 1 \right) e^{-2aT} \\ &= \frac{2ab}{(a-2b)^2} (1 - e^{2(a-b)\tau}) \frac{b}{a-b} e^{-2aT} \\ &= \frac{2ab^2}{(a-2b)^2(a-b)} (1 - e^{2(a-b)\tau}) \frac{b}{a-b} e^{-2aT}. \end{aligned} \quad (6.148)$$

For the terms proportional to e^{-4bT} we find

$$\frac{4a^2b}{(a-2b)^2}(1-e^{2b\tau})\left(\frac{1}{4b}-\frac{1}{2b}\right)e^{-4bT}=\frac{a^2}{2(a-2b)^2}(1-e^{2b\tau})e^{-4bT}. \quad (6.149)$$

Finally, for the terms proportional to $e^{-(a+2b)T}$ we find

$$\begin{aligned} \frac{4a^2b}{(a-2b)^2}(1-e^{a\tau})\left(\frac{2}{a+2b}-\frac{2}{a}\right)e^{-(a+2b)T} &= \frac{8ab}{(a-2b)^2(a+2b)}(1-e^{a\tau})(a-(a+2b))e^{-(a+2b)T} \\ &= \frac{-16ab^2}{(a-2b)^2(a+2b)}(1-e^{a\tau})e^{-(a+2b)T} \end{aligned} \quad (6.150)$$

Adding these four different contributions together then yields equation (6.136). \square

We note that Lemma 8 and Theorems 6 and 9 are compared to experimental results obtained with a trapped-ion device in Figure 6.10.

6.10 Single-click model

In this section, we present an analytical model for the entangled states created on elementary links when using a single-click entanglement generation protocol [66]. This model is used as one of the building blocks of our NetSquid simulations, as mentioned in Section 6.13, and based on the models previously introduced in [46–48]. The novelty of the model presented here lies in combining features of the three previous models [46–48] and in additionally considering the possibility that non-number-resolving detectors may be used (the three cited papers assume the use of number-resolving detectors).

6.10.1 Model assumptions

We model a single-click protocol for entanglement generation on an elementary link between two nodes, which we designate A and B. The protocol starts with the preparation of an optically-active matter qubit at each of the nodes in the following state:

$$|\psi_m\rangle = \sqrt{\alpha}|\uparrow\rangle + \sqrt{1-\alpha}|\downarrow\rangle, \quad (6.151)$$

where the subscript m stands for matter, $|\uparrow\rangle$ is a bright state, i.e., a state that rapidly decays via photon emission after being excited, and α is the bright-state parameter, i.e. the fraction of the matter qubit's state that is in $|\uparrow\rangle$. Excitation and subsequent radiative decay of $|\uparrow\rangle$ entangles the state of the matter qubit with the presence $|1\rangle$ or absence $|0\rangle$ of a photon (subscript p):

$$|\psi_m, \psi_p\rangle = \sqrt{\alpha}|\uparrow\rangle|1\rangle + \sqrt{1-\alpha}|\downarrow\rangle|0\rangle \quad (6.152)$$

The photons are then sent to a heralding station where they are interfered. Detection of a single photon heralds the generation of a matter-matter entangled state.

In our analytical model, we account for the following imperfections when computing the success probability and entangled state generated with the protocol:

- Double excitation of the matter qubit. Resonant laser light incident on the optically-active matter qubit triggers its excitation. It is possible that this excitation happens

two times as the laser shines on the matter qubit, leading to the emission of two photons. This happens with probability p_{dexc} . We note that the excitation could in theory also happen multiple times, but, as detailed in Section 6.10.2, the effect this would have on the state would be the same as if two photons were emitted, so we can absorb the probability of more than one excitation into one quantity.

- Photon phase uncertainty. The photons interfering at the midpoint acquire a phase during transmission over the fiber, and the difference between the phases of the two interfering photons influences the entangled state that is generated [46]. The dephasing probability of the state p_{phase} can be computed from the standard deviation of the difference between the acquired phases σ_{phase} [47]:

$$p_{\text{phase}} = \frac{1}{2} \left(1 - e^{-\sigma_{\text{phase}}^2/2} \right). \quad (6.153)$$

Furthermore, we also account for photon loss, imperfect indistinguishability, non-photon-number-resolving detectors and detector dark counts, as described in 6.8. Finally, we account for the possibility of asymmetry in the placement of the heralding station, the attenuation of the fibers connecting the nodes to the heralding station and the bright-state parameters of the nodes.

6

6.10.2 Results

Here we present the derivation of the entangled matter-matter state generated in our model of the single-click protocol. We split this derivation into four situations, as done in [46–48]. Each of them corresponds to one of the different configurations of the states of the matter qubits that can result in a heralded success.

1. Both matter qubits are in the bright state. In this case, both matter qubits emit a photon, so this situation is heralded as a success if:
 - (a) Only one of the emitted photons survives. For NR detectors, there cannot be dark counts in either of the detectors (as otherwise two photons would be detected in the detector which also saw the actual emitted photon, and the event would be heralded as a failure). For NNR detectors, the requirement is just that there is no dark count in the detector that did not detect the emitted photon. The probability of this case happening is:

$$p_{1a} = \begin{cases} \alpha_A \alpha_B (1 - p_{\text{dc}})^2 (p_A (1 - p_B) + p_B (1 - p_A)) & \text{if NR,} \\ \alpha_A \alpha_B (1 - p_{\text{dc}}) (p_A (1 - p_B) + p_B (1 - p_A)) & \text{if NNR.} \end{cases} \quad (6.154)$$

- (b) No emitted photon is detected, and there is a dark count in one of the detectors. This case is the same irrespective of whether or not the detectors are NR. There is a factor of two because this can happen in either detector. The probability of this case happening is:

$$p_{1b} = 2\alpha_A \alpha_B (1 - p_A)(1 - p_B)(1 - p_{\text{dc}})p_{\text{dc}}. \quad (6.155)$$

- (c) Both emitted photons make it to the midpoint and are detected, but they bunch and go into the same detector. Furthermore, there is no dark count in the other detector. There is a factor of two because this can happen in either detector. This is heralded as a failure if the detectors are NR. The probability of this case happening is:

$$p_{1c} = \begin{cases} 0 & \text{if NR,} \\ \alpha_A \alpha_B p_A p_B p_{\text{same dets}} (1 - p_{\text{dc}}) & \text{if NNR,} \end{cases} \quad (6.156)$$

where $p_{\text{same dets}}$ is the probability that the two photons go to the same detector, as derived in case F2 of Section 6.8:

$$p_{\text{same dets}} = 1 - \frac{1 - V}{2}. \quad (6.157)$$

2. Matter qubit A is in the bright state, matter qubit B is not. In this case, only matter qubit A emits a photon, so this situation is heralded as a success if:

- (a) The emitted photon survives. For NR detectors, there cannot be dark counts in either of the detectors (as otherwise two photons would be detected in the detector which also saw the actual emitted photon, and the event would be heralded as a failure). For NNR detectors, the requirement is just that there is no dark count in the detector that did not detect the emitted photon.

$$p_{2a} = \begin{cases} \alpha_A (1 - \alpha_B) (1 - p_{\text{dc}})^2 p_A & \text{if NR,} \\ \alpha_A (1 - \alpha_B) (1 - p_{\text{dc}}) p_A & \text{if NNR.} \end{cases} \quad (6.158)$$

- (b) The emitted photon does not survive, and there is a dark count in one of the detectors. This case is the same irrespective of whether the detectors are NR. There is a factor of two because this can happen in either detector.

$$p_{2b} = 2\alpha_A (1 - \alpha_B) (1 - p_A) (1 - p_{\text{dc}}) p_{\text{dc}}. \quad (6.159)$$

3. Matter qubit B is in the bright state, matter qubit A is not. In this case, only matter qubit B emits a photon. This identical to case 2, interchanging A and B.
4. Neither of the matter qubits are in the bright state. No photon is emitted. In this case, we only get a success if there is a dark count in one of the detectors, but not the other. This case is the same irrespective of whether the detectors are NR. There is a factor of two because this can happen in either detector.

$$p_4 = 2(1 - \alpha_A)(1 - \alpha_B)(1 - p_{\text{dc}})p_{\text{dc}}. \quad (6.160)$$

The overall success probability of the protocol p_{suc} is then given by adding up the probability that each of the cases above happens, $p_{\text{suc}} = p_1 + p_2 + p_3 + p_4$, with $p_1 = p_{1a} + p_{1b} + p_{1c}$, $p_2 = p_{2a} + p_{2b}$ and $p_3 = p_{3a} + p_{3b}$.

The unnormalized density matrix of the generated state ρ can then be obtained by taking the model introduced in [47] and replacing the probabilities appropriately. The result is the following:

$$\rho = \begin{pmatrix} p_1 & 0 & 0 & 0 \\ 0 & p_2 & \pm \sqrt{V} p_2 p_3 & 0 \\ 0 & \pm \sqrt{V} p_2 p_3 & p_3 & 0 \\ 0 & 0 & 0 & p_4 \end{pmatrix}, \quad (6.161)$$

where the sign depends on which detector clicked.

Two more dephasing channels are then applied in succession to the state in order to account for the effect of double photon excitation and photonic phase drift.

The first channel, corresponding to double excitation, is applied to both matter qubits, with probability $p_{\text{dexc}}/2$. The light pulse used to excite the bright state to a short-lived excited state is not instantaneous, so there is a chance that the matter qubit decays back down to the original state and be re-excited before the pulse is complete. The first emitted photon will be lost to the environment because it is impossible to distinguish it from the laser light used to excite the matter qubit [47]. It must then be traced out, resulting in a loss of coherence between the two matter qubit states. However, detection of the second emitted photon will falsely herald entanglement, so we apply a dephasing channel with probability $p_{\text{dexc}}/2$ to account for the possibility that more than one photon is emitted.

The second one, corresponding to the photonic phase drift, is applied to only one of them, with probability p_{phase} . The difference in the phases acquired by the two interfering photons results in a phase difference between the two components of the resulting Bell state [47]. Applying a dephasing channel to only one of the matter qubits, with the correct probability given by p_{phase} , has the same effect.

6.11 Optimization method

In this section we provide more details regarding our optimization methodology. As mentioned in the main text, this methodology is based on genetic algorithms, which come in several different flavors. Our particular implementation is heavily based on the one introduced in [79], to which we refer the interested reader. The only novelty introduced here is the use of a different termination criterion, which is explained in detail in the following section. We note also that the code for our implementation, together with the tools required for integration with NetSquid simulations, is publicly accessible at [108].

6.11.1 Termination criteria for genetic algorithms

The matter of choosing termination criteria for genetic algorithms (and, more generally, evolutionary algorithms) has been the object of some study (see, e.g., [109] for a review). If the algorithm is terminated too soon, good solutions might remain undiscovered. On the other hand, running the algorithm for too long in case good solutions have already been found leads to wasting computational resources. Typically-used termination criteria can be split into two groups [109]:

1. Direct termination criteria: these can be obtained directly from the optimization, without any extra data analysis. Examples include setting a maximum number of

generations for the optimization or imposing a threshold value on the value of the best solution's cost;

2. Derived termination criteria: these are *a posteriori* criteria, requiring that some data analysis be performed on the outcome of the optimization. Examples include setting a threshold on the standard deviation of the costs of the individuals in the population or on the gap between the best and worst individuals in a given generation.

The authors of [109] applied an evolutionary algorithm to a particular cost function with different termination criteria. They found that the only reliable termination criteria fitting into the groups above were one in which the algorithm terminated after a fixed, predetermined number of generations, which we name GEN, and one in which the best solution had not varied by more than a predetermined value after a predetermined number of generations, which we name VAR. For all the other criteria tested, the algorithm did not terminate even though the optimal solution had already been found. GEN and VAR both have the drawback of depending on hyperparameters for which a good choice can only be made with knowledge of the problem at hand. By this we mean that the number of generations or accepted variation in the best solution per generation that guarantee termination are problem-dependent.

With this in mind, we opted to employ VAR as the termination criterion for our optimization runs. We made this choice as using VAR results in a more systematic, performance-dependent process for the decision of terminating the optimization. By this we mean that even though GEN guarantees termination (by definition) it does so in an arbitrary way by deciding to stop the optimization without any regard for its evolution. As suggested in [109] we terminated the algorithm if the best solution's cost averaged over the past fifteen generations had not varied by more than a given value. In contrast with the work of [109], we measured the variation in percentual terms. For each setup, we ran the optimization process ten times, each for two hundred generations. Then, to determine what the tolerance for the variation should be, we swept across its values, starting at 1% and with a step of 1%. The chosen tolerance was the one that guaranteed termination for all ten of the optimization runs, and the best solution (i.e., the ones showed in this work), was then the best cost found across the ten different runs, up until termination. We note that the tolerance can be different for different setups. We further note that this offline implementation of VAR is not good for saving computational resources, as the optimization must anyway be run for a large number of generations, with some of them being discarded. It was however simpler to integrate into our workflow, which weighed heavily since we were more constrained in working hours than in computing hours.

6.11.2 Cost function

Our goal with this work was to find the minimal requirements for a quantum repeater enabling Verifiable Blind Quantum Computation between two nodes separated by fiber of length 226.5 km. This implies solving a multi-objective optimization problem, as we want to minimize hardware-parameter improvement while simultaneously ensuring that performance targets are met. There are various ways of approaching such problems, one of them being scalarization. This consists of adding the cost functions corresponding to different objectives together, so that effectively only one scalar quantity has to be opti-

mized. Through this process, we arrive at the cost function C introduced in the Methods, which we reproduce in equation (6.162).

$$C = w_1 \left(1 + (F_{min} - F)^2\right) \Theta(F_{min} - F) + w_2 \left(1 + (R_{min} - R)^2\right) \Theta(R_{min} - R) + w_3 H_C(x_{1_c}, \dots, x_{N_c}) \quad (6.162)$$

We recall that H_C is the hardware cost, w_i are the weights of the objectives, Θ is the Heaviside function and F and R are the average teleportation fidelity and entanglement generation rate achieved by the parameter set, respectively. F_{min} and R_{min} are the minimal performance requirements. Scalarization conveniently transforms multi-objective optimization into their much simpler single-objective counterparts, but it does so by stowing away the problem in defining the weights w_1 , w_2 and w_3 assigned to each of the objectives. Different choices in the weights can lead to different outcomes from the optimization procedure. In this work, just as in [79], we wanted the performance targets to be hard requirements, i.e. a set of hardware parameters that did not fulfill them should not be assigned a low cost. To ensure this, we picked $w_1, w_2 \gg w_3$, such that $w_1 \left(1 + (F_{min} - F)^2\right) \Theta(F_{min} - F)$, $w_2 \left(1 + (R_{min} - R)^2\right) \Theta(R_{min} - R) \gg w_3 H_C(x_{1_c}, \dots, x_{N_c})$. We set $w_1 = w_2 = 1 \times 10^{20}$ and $w_3 = 1$. No particular heuristic was used to select these numbers. They were picked because they ensure that the cost assigned to not meeting the performance targets is much higher than the hardware cost, effectively making the performance targets hard requirements.

6

As mentioned in the main text, we picked the hardware cost function because it reflects the concept of progressive hardness, i.e. that parameters become harder to improve as they approach their perfect value. Furthermore, it satisfies a composability property regarding the probability of no-imperfection. To see this, consider a parameter's probability of no-imperfection p that can be expressed as the product of two other parameters' probabilities of no-imperfection, $p = p_a p_b$. p could for example be the probability that a photon emitted in the correct mode is collected into a fiber, while p_a and p_b are the probabilities that the photon is emitted with the right wavelength and collected into the fiber, respectively. Improving p by a factor of k takes it to $\sqrt[k]{p} = \sqrt[k]{p_a p_b} = \sqrt[k]{p_a} \sqrt[k]{p_b}$, which is equivalent to improving p_a and p_b separately by the same factor k . Therefore, hardware improvement as measured by this function is invariant to the granularity at which parameters are considered.

The last aspect we would like to highlight regarding the cost function is its squared difference terms, i.e. $1 + (F_{min} - F)^2$ and $1 + (R_{min} - R)^2$. These were introduced in [110] and are used to steer the algorithm towards sets of hardware parameters that are more likely to meet the performance targets. They do this by ensuring that parameter sets which fail to meet the targets by a large margin are assigned a higher cost, being therefore less likely to progress into further generations.

6.11.3 Probabilities of no-imperfection

For some parameters, such as the probability that a photon is not lost when coupling to a fiber, the conversion to probability of no-imperfection is obvious. For others, such as coherence times, this is not so. Therefore, we show in Table 7.5 the probability of no-imperfection for all parameters considered in our hardware models. We proceed with the derivation of the probability of no-imperfection for some of the less obvious cases.

Parameter	Probability of no-imperfection
Photon detection probability excluding attenuation losses p_{det}	p_{det}
Probability of double excitation p_{dexc}	$1 - p_{\text{dexc}}$
Gate depolarizing probability p_{dep}	$1 - p_{\text{dep}}$
Number of entanglement generation attempts before dephasing $N_{1/e}$	$(1 + e^{-1/N_{1/e}})/2$
T_1	e^{-t/T_1}
T_2	e^{-t/T_2}
Ion coherence time T_c	e^{-t/T_c^2}
Emission fidelity F_{em}	$1/3(4F_{\text{em}} - 1)$
Swap quality s_q	s_q
Visibility V	V
Dark count probability p_{dc}	$1 - p_{\text{dc}}$

Table 6.6: Probabilities of no-imperfection for parameters we optimized over in this work. Some parameters were merged for brevity, e.g. the probability of no-imperfection presented for T_2 holds for the abstract model T_2 , the carbon spin T_2 and the electron spin T_2 . In the probability of no-imperfection for each of the coherence times, t is the time spent in memory.

As mentioned in the main text, and explained in detail in Supplementary Note 4c of [60], the initialization of an color center's electron spin state induces dephasing of its carbon spin states through their hyperfine coupling. This is typically modelled as the carbon spin states dephasing with some probability each time entanglement generation is attempted [34]. This probability can be related to $N_{1/e}$ as $p = 1/2(1 - e^{-1/N_{1/e}})$. The corresponding probability of no-imperfection is then $p_{ne} = 1 - p = (1 + e^{-1/N_{1/e}})/2$.

T_1 represents the timescale over which qubit relaxation occurs, with the probability of amplitude damping occurring over a period of time t being given by $p_{ad} = 1 - e^{-t/T_1}$. The associated probability of no-imperfection is e^{-t/T_1} . Improving T_1 by a factor of k then corresponds to improving the probability of no-imperfection to $\sqrt[k]{e^{-t/T_1}}$. Some algebra reveals that this is equivalent to multiplying T_1 by a factor of k , and that this holds irrespective of the chosen timescale.

T_2 represents the timescale over which qubit dephasing occurs, with the probability of a Z error occurring over a period of time t being given by $p_z = (1 - e^{-t/T_2})/2$. The associated probability of no-imperfection is $p_{ne} = \frac{1 + e^{-t/T_2}}{2}$. To first order, this can be written as $p_{ne} = e^{-t/2T_2}$, and some algebra again reveals that with this approximation improving T_2 by a factor of k is equivalent to multiplying it by the same factor.

The ion coherence time T_c also represents a timescale for dephasing, but in this the case the probability of a Z error occurring is given by $1 - \frac{1}{2}(1 + e^{-2t^2/T_c^2})$. To first order, the probability of no-imperfection can thus be written as $p_{ne} = e^{-t^2/T_c^2}$. In this case, improving T_c by a factor of k is equivalent to multiplying it by \sqrt{k} .

We model noise in photon emission as a depolarizing channel of fidelity F_{em} . The

action of the depolarizing channel on a perfect Bell state $|\Phi^+\rangle$ can be written as follows:

$$|\Phi^+\rangle\langle\Phi^+| \rightarrow (1 - p_{\text{dep}})|\Phi^+\rangle\langle\Phi^+| + p_{\text{dep}}\frac{\mathbb{I}}{4},$$

where p_{dep} is the associated depolarizing probability and \mathbb{I} is the identity matrix. This can be rewritten as:

$$|\Phi^+\rangle\langle\Phi^+| \rightarrow (1 - p_{\text{dep}})|\Phi^+\rangle\langle\Phi^+| + p_{\text{dep}}\frac{1}{4}(|\Phi^+\rangle\langle\Phi^+| + |\Phi^-\rangle\langle\Phi^-| + |\Psi^+\rangle\langle\Psi^+| + |\Psi^-\rangle\langle\Psi^-|).$$

Since the Bell states are orthogonal to each other, it follows that $p = \frac{4}{3}(1 - F_{\text{em}})$ and that the corresponding probability of no-imperfection is $\frac{1}{3}(4F_{\text{em}} - 1)$.

The derivation of the probability of no-imperfection for the remaining parameters should be self-evident and is therefore omitted here.

6.11.4 Optimizing over tunable parameters

As discussed in the Methods, the entanglement generation and distribution protocols employed in our simulations include parameters that can be freely varied. We name these *tunable* parameters. They affect the behavior and performance of the setups we investigated, and as a consequence also the minimal hardware requirements. The tunable parameters should thus be chosen such that the best possible performance is extracted from a given set of hardware parameters, minimizing the cost function. The values of the tunable parameters that allow for this are the optimal values. This is however not trivial, as different sets of hardware parameters perform best with different tunable parameters. To illustrate this, we again go over the tunable parameters considered in our simulations.

We start with the *cut-off time*. This is the maximum duration for which a state can be held in memory before being discarded. For details on the implementation of a cut-off timer in our simulations, see Section 6.13. If the cut-off time is very short, states will not be held in memory for long, and therefore the end-to-end fidelity will be high. On the other hand, states will also be frequently discarded and regenerated, which means that establishment of end-to-end entanglement will take longer. In contrast, a very long cut-off is equivalent to no cut-off, in the sense that states are never discarded. This maximizes the entanglement generation rate at the expense of lower state fidelity.

The second tunable parameter is the *bright-state parameter*, which is relevant in the single-click entanglement generation protocol. This is the fraction of the superposition that is in the optically-active state, and therefore corresponds to the probability that a photon is emitted. A larger bright-state parameter corresponds to a higher probability of entanglement generation, but at the expense of a lower fidelity, as it also introduces a component orthogonal to the Bell basis in the generated entangled state. For more details on single-click entanglement generation see Section 6.10.

The final tunable parameter is the *coincidence time window*, which is part of our trapped ion double-click entanglement generation model. Two detection events arising from the correct detectors are only heralded as a success if the time elapsed between the events is smaller than the coincidence time window. It acts as a temporal filter, lowering the protocol's success probability but increasing the visibility and hence the fidelity of the

generated entangled states. For more details on our modeling of a coincidence time window, see Section 6.9.

These three parameters can be used to trade-off rate against fidelity, and their optimal values are different for different sets of hardware parameters. For example, if the coherence time is short and the detection probability is high, it will likely be beneficial to have a short cut-off time. The opposite is true if the coherence time is long and the detection probability is low.

In order to find good values for the tunable parameters, we included them as parameters to be optimized by the genetic-algorithm-based optimization machinery. We imposed that the values the cut-off time can take are in the interval between $0.1 T_C$ and T_C , where T_C is the coherence time (collective dephasing coherence time for trapped ions, T_2 for abstract nodes and carbon T_2 for NV centers). The expected entanglement generation time grows exponentially as the cut-off time is reduced, so the lower bound was imposed to prevent the simulation taking unreasonably long to run. Furthermore, we anyway expect that a too low cut-off time would not allow the rate target to be met, so we can be reasonably sure that no cheap hardware requirements are missed by imposing this constraint. The upper bound is imposed as we observed that not imposing it made it hard for the algorithm to converge, due to the reduced sensitivity of the target metrics to high values of the cut-off time. As discussed above, employing a very long cut-off time is effectively equivalent to not employing one at all. Therefore, in that regime the choice of cut-off time becomes irrelevant, and the set of parameters minimizing the cost function is chosen independently of it. We have empirically observed that the cut-off time tends to converge to around 65% of the relevant coherence time, which is fairly distant from both bounds we imposed. A back-of-the-envelope calculation can also be performed to argue that it is unlikely that allowing for cut-off times which are larger than the memory's coherence time would be useful. We do this by computing the end-to-end fidelity in a single sequential-repeater setup under the following assumptions:

- The cut-off time is equal to the memory dephasing time;
- There are no other noise sources.

The worst case scenario in this setup in terms of fidelity occurs when the second entangled state takes exactly cut-off time seconds to be generated, resulting in both qubits of the first entangled pair to dephase for a time equal to their dephasing time. The dephasing probability is in this case given by $p_Z = \frac{1-e^{-2}}{2}$. Assuming that the state that had been generated was $|\Phi^+\rangle$, the post-dephasing state is a mixture of $|\Phi^+\rangle$ and $|\Phi^-\rangle$:

$$\rho = (1 - p_Z)|\Phi^+\rangle\langle\Phi^+| + p_Z|\Phi^-\rangle\langle\Phi^-|. \quad (6.163)$$

This has a fidelity of 0.57 with the target Bell state $|\Phi^+\rangle$, corresponding to a teleportation fidelity of 0.71. This value is much lower than our lowest teleportation fidelity target, 0.8571, even with no noise sources besides dephasing noise on the memory. It is then unlikely that picking even higher cut-off times would lead to finding better solutions to our optimization problem.

In the single-repeater setup we investigated, there are four bright-state parameters to be chosen, corresponding to the four different fiber segments between processing nodes

and heralding stations. We imposed that αp_{det} had to be equal for all of them, with α is the bright-state parameter and p_{det} the probability that a photon is not lost in the fiber connecting the node to the midpoint station. This condition guarantees balanced entanglement-generation success probabilities across all segments, which is a good heuristic for segments connecting to the same heralding station, as it maximizes the fidelity of the generated states [48]. Imposing it also for segments connecting to different heralding stations was done in order to reduce the size of the search space.

There are also two coincidence time window parameters to be chosen, corresponding to the two elementary links. We imposed that they must have the same value in order to make the search space smaller.

6.12 Simulation performance

Each execution of our quantum-network simulations simulates the delivery of n end-to-end entangled states. When the protocols running on the end nodes learn through classical communication between nodes that n states were successfully distributed, they abort and the simulation terminates. In Figure 6.14, we show how the runtime of our simulation of the Delft - Eindhoven setup scales with the number n . As expected, the scaling is linear.




Figure 6.14: Performance of our simulation of the Delft - Eindhoven setup with abstract model nodes and a cut-off timer using a machine running 40 Intel Xeon Gold cores at 2.1 GHz and 192 GB of RAM. The runtime scales linearly with the number of entangled pairs being distributed. Distributing 100 times, which we have empirically determined is enough to evaluate the performance of a given parameter set with reasonable accuracy, takes roughly one second. The data point corresponding to n pairs was obtained by running the corresponding simulation 500 times. The error bars represent the standard error of the mean.

A simulation with $n = 100$, which we have empirically determined is enough to evaluate the performance of a given parameter set with reasonable accuracy, takes roughly 1 s. To

be more concrete, when running a color-center double-click simulation using the minimal hardware parameters presented in Section 6.1, we find that after distributing 100 pairs a teleportation fidelity F_{tel} of 0.8774 ± 0.0035 and a rate of $0.106 \text{ Hz} \pm 0.003$ are obtained.

We note that Figure 6.14 was obtained by running the simulation without a cut-off. Although the runtime still grows linearly with the number of distributed entangled pairs with a cut-off, its inclusion does mean that the simulation runtime grows exponentially as the the cut-off time becomes shorter. This is because the expected number of necessary entanglement generation attempts also grows exponentially, as seen in Figure 6.15.




Figure 6.15: Performance of our simulation of the Delft - Eindhoven setup with abstract model nodes and a cut-off timer using a laptop running a quad-core Intel i7-8665U processor at 1.9 GHz and 8 GB of RAM. The runtime scales exponentially as the cut-off time is reduced. The data point corresponding to n pairs was obtained by running the corresponding simulation 20 times. The error bars represent the standard error of the mean.

As discussed in Section 6.3 and in Section 6.11, the optimization methodology we employ requires running our simulation for many different sets of parameters. We now estimate a lower bound on the time required to perform optimization in one setup. We run the optimization algorithm for 200 iterations. In each of these, there are 200 different parameter sets, and the distribution of 100 entangled pairs is simulated for each. The computing nodes in the high-performance-computing cluster we use have 128 cores, which means that the simulation for 128 of the 200 parameter sets can be executed in parallel. Assuming that there is no cut-off, or that it is large enough not to significantly impact the simulation runtime, this means that we can expect 1 generation to be run in roughly 2.5 seconds. The data processing and file input and output required to generate new sets of parameters take a comparable amount of time, making $T = 200 \times 5 \text{ s}$, roughly seventeen minutes, a good estimate for the time required to perform optimization for one setup. We must however stress that this is a very optimistic lower bound, because as Figure 6.15

makes clear, the use of a cut-off has a huge impact on the runtime of the simulation. We have observed that optimization of most of the setups we studied required 10 to 20 hours to terminate.

6.13 Framework for simulating quantum repeaters

In this section, we discuss the framework that we use to evaluate the performance of quantum repeaters. This framework is presented in this work for the first time.

The code that we have used to simulate all the quantum networks in this chapter is publicly available [105]. The repository contains code that has a much broader applicability than simulating the networks of up to three nodes presented here. In fact, the simulations can be used to assess the performance of quantum-repeater chains with any number of nodes, and any spacing between nodes. The currently supported types of nodes are those containing NV centers, ion traps or abstract quantum processors, and the currently supported types of entanglement generation between neighboring nodes are the single-click and double-click protocols. The simulation code depends on a number of other public repositories [95, 111–115], all of which were developed in tandem with the code for this chapter and will be explained in more detail below.

6.13.1 Services

The primary functional unit of our quantum-network simulations is the “service”, which is defined by an input, an output, and its intended function. An example of a service that can be defined on a node is the measurement service. It takes as input a request to measure a qubit, and the intended function is that the qubit is measured. As output, the service returns the measurement outcome.

A service is distinct from its implementation, which is a protocol. Protocols make sure the intended function is fulfilled and generate the appropriate output. Different protocols can fulfill the same function. For example, in the case of the measurement service, a protocol that simulates a direct measurement of the required qubit (e.g., a fluorescence measurement) could be activated. Another possible implementation would be a protocol that first swaps the quantum state of the required qubit to some different physical qubit (that perhaps allows for higher-fidelity measurements), and then simulates a measurement of that qubit. The distinction between service and its implementation is illustrated in Figure 6.16, which emphasizes that the same high-level functionality can be implemented using different physical systems.

Treating services and their implementation separately has two distinct advantages for our simulations. First, it allows us to easily run the same protocols on different types of simulated hardware. Take as example performing an entanglement swap in the broader context of a repeater chain. To do so, the repeater protocol will place a request with the local entanglement-swap service. The repeater protocol does not need to know how the swap is implemented. On an abstract quantum processor, it can be implemented using a CNOT gate, while on an ion trap, it can be implemented using a Mølmer–Sørensen gate. Second, it allows for a modular stack of protocols, where protocols implementing a specific service can easily be interchanged. In the example of the repeater protocol, requests are made of an entanglement-generation service before the swap can be performed. If the




Figure 6.16: The black box represents a service, defined by a set of inputs, a set of outputs and some promised functionality. The protocols interacting with the service need not know how this functionality is implemented. Therefore, different implementations can be swapped in and out. In the figure, color centers and trapped ions are depicted to emphasize that the same high-level functionality can be executed by different physical systems.

protocol runs on an NV node, entanglement could either be generated using a single-click or double-click protocol. Switching between these two modes is easily realized by changing the protocol that implements the entanglement-generation service. Again, the repeater-node protocol does not need to be adapted.

The main interface of the repeater chain itself is also defined by a service. The service implemented by the repeater chain is a link-layer service [78, 116], which provides robust entanglement generation between the end nodes of the chain. These requests should be put on the end nodes of the chain, which activates a protocol that uses a messaging service to put requests on the SWAP-ASAP repeater services defined on the repeater nodes of the network. When the end-node protocols confirm they share entanglement (using a protocol that tracks entanglement in the network based on the classical communication shared by nodes), an appropriate output message is returned by the service. This is the cue that we use in our simulations to collect the density matrix of the created state and the time it took to create it.

6.13.2 SWAP-ASAP protocol

A SWAP-ASAP repeater chain is one in which repeater nodes perform an entanglement swap as soon as they hold two entangled qubits that were generated with different neighbors. This is in contrast to e.g. nested repeater schemes [117, 118]. We have implemented two different SWAP-ASAP repeater protocols. The first is suitable for repeater chains of any length and node spacing, and for repeater nodes that can generate entanglement with either one or both neighbors at the same time. The second, on the other hand, has been tailored more specifically to the one-repeater scenario studied in this chapter. It assumes that entanglement generation is limited to a single neighbor at a time. First a request is is-

sued to generate entanglement over a single connection. Once that has finished, a request is issued for the second link, and a swap is executed as soon as entanglement is confirmed. In case the connections are not of equal length, entanglement generation takes place on the longer link first. The reason for this is that the longer connection is expected to be the connection on which entanglement generation takes longer. By finishing the longer link first, the total time that entanglement needs to be stored in quantum memory is minimized. The second protocol is the one used to generate the results reported in this chapter.

To generate entanglement over elementary links, the repeater protocols issue requests with the entanglement service. In the protocol that we use to implement this service, these requests are queued. The number of requests that are processed simultaneously is hardware-dependent, and is a free parameter in our simulations. When handling a request for entanglement the protocol will, before doing anything else, issue a request to an agreement service. This service is in charge of synchronizing neighboring nodes that want to generate entanglement together. This is needed as typically both nodes need to be actively involved in generating entanglement for a state to be created between the two. In our simulations, we use an implementation of the agreement service where even-numbered nodes in the chain always initiate entanglement generation. These nodes will send a classical message to their neighbors when a request for agreement is made, and then wait for those nodes to send a classical reply indicating readiness, after which entanglement generation can start. On the other hand, when a request is made on an odd node, it will check whether a classical message has been received by the neighboring even node in the past. If so, it will reply indicating readiness. Otherwise, the request for agreement will be rejected. In that case, the entanglement service can try to process the next request in the queue, and see if agreement can be reached with this node again at some later time.

In case agreement is reached between two nodes, the entanglement protocols of the nodes will start entanglement generation. In our simulations, this is done using analytical models that decide after how much time an entangled state should be created between the nodes, and what this state should look like. This process is known as *magic* [95] and is further discussed in Section 6.6.8.

Finally, there is a cut-off protocol active on repeater nodes. It discards qubits that have been stored in quantum memories for too long. The exact amount of time after which states are discarded is called the cut-off time, and is a tunable parameter that allows for a trade-off between end-to-end entangling rate and fidelity. Every node runs an entanglement-tracking protocol that keeps track of both any local entangled qubits, and what entangled states currently exist in the network at large. Whenever the entanglement service registers a new qubit at the entanglement tracker, the cut-off protocol starts a timer. When the timer goes off, the cut-off protocol checks whether the entangled qubit still exists in local memory. If so, the entanglement tracker is told to discard the qubit. The entanglement tracker will also communicate classically with the entanglement trackers of other nodes in the network to inform them that the qubit has been discarded. If an entanglement tracker learns that a qubit has been discarded that was entangled with one of its local qubits, it responds by discarding that qubit as well. Links corresponding to discarded qubits must be regenerated. We note that the cut-off protocol does not run on the end nodes of the repeater chain. This is to prevent the possibility of one end node believing end-to-end entanglement has been achieved, while the other end node has in

actuality discarded its qubit (but the classical message has not yet reached the first end node).

6.13.3 Configuring quantum networks

In our simulations, quantum networks are made up of nodes. Each node represents a single physical location, and contains an object that we refer to as “driver”. This object provides a mapping between services and protocols that implement those services. The driver allows access to services without knowledge of their implementations. Each node has its own driver. Apart from drivers, nodes hold components that represent quantum hardware, which allow for the storage and/or manipulation of quantum states. The protocols running on the node can use this quantum hardware to implement specific services. The nodes in our simulations are ready-made packages with both driver and hardware included. In order to use them in a quantum network, they just need to be initialized (thereby specifying their parameters) and connected to other nodes.

The simulations performed for this chapter contain three different types of nodes. The first is the NV node. It holds an NV quantum processor, which is imported from the Python package NetSquid-NV [112]. The second is the ion-trap node. This node holds an ion-trap quantum processor, imported from the Python package NetSquid-TrappedIons [114]. Finally, there is the abstract node, which contains an abstract quantum processor imported from the Python package NetSquid-AbstractModel [119]. On initialization, each of these takes hardware parameters specific to the type of hardware being simulated, and a number of parameters used to configure the protocols used at the node. For example, the cut-off time needs to be specified, and in case of single-click heralded entanglement generation, the bright-state parameter as well.

Nodes are connected by two types of connections. These connections are themselves also ready-made packages, and can be found in the Python package NetSquid-PhysLayer [113]. The first type is the classical connection, which represents optical fiber that can be used to send classical messages. The second type is the heralded connection. It represents a midpoint station connected to two nodes by optical fiber, where optical Bell-state measurements can be performed on incoming photons. Such a connection can be used to perform heralded entanglement generation. As discussed above, we do not simulate the process of heralded entanglement generation itself, but instead use analytical models to magically create entangled states. However, the heralded connections still perform an important role as placeholders. Parameters passed to the heralded connection when configuring the network are later retrieved by the analytical models to decide how long it should take before a state is created, and what that state should be exactly. One key parameter specified in the heralded connection is whether single-click or double-click heralded entanglement distribution is used. In the simulations presented in this chapter neighboring nodes are always connected by both a classical connection and a heralded connection.

To put together nodes and connections for the creation of quantum networks, and to configure their parameters, we make use of the Python package NetSquid-NetConf [111]. The tools provided in this package allow for the writing of human-readable configuration files. These configuration files contain entries for all the different nodes and connections in the network. Their type is specified (such as “NV node” or “heralded connection”), as well as their parameters and how they are connected. These configuration files can also

be used to vary some of the parameters, allowing to e.g. perform a parameter scan over one of them and observe its effect on the network performance.

6.14 Extra optimization results

In this section, we present results of optimizations we performed that were not presented in the main text, but might still be of interest.

6.14.1 To move or not to move

As mentioned in the main text, the communication qubit of color centers has typically shorter coherence times than the memory qubits. For the baseline hardware parameters we investigated, the communication qubit had $T_1 = 1$ hours [50] and $T_2 = 0.5$ s [49], whereas the memory qubits had $T_1 = 10$ hours and $T_2 = 1$ s [51]. It might then be worthwhile for the end node that generates entanglement with the repeater first, i.e., the Eindhoven node, to move its half of the entangled state to memory while waiting for end-to-end entanglement to be established, even though that comes at the cost of more noise being introduced in this operation. A diagram of the circuit used for this operation can be found in Supplementary Note 5 B of [60]. To investigate this, we applied our methodology to two single-repeater color-center setups performing double-click entanglement generation. In one of the setups, which we name "move scenario", once the first elementary link is established, the end node performs the move operation while the waiting for the second link to be established. In the other setup, which we name "no-move scenario", the state is kept in the electron spin until end-to-end entanglement is established. The hardware requirements for these two scenarios are shown in Figure 6.17. The move scenario requires that the two-qubit gate be significantly improved, which is to be expected as the move operation requires the application of two of these gates [60]. On the other hand, the move scenario does not require an improvement on the electron spin's coherence time, in contrast with the no-move scenario. This is also not surprising, as in the move scenario entanglement is not stored in the electron spin for a significant amount of time.

The overall cost associated to the no-move scenario is slightly lower than the cost of the move scenario, so all the NV center results presented in the main text were obtained in the no-move scenario. We stress that this finding, although relevant for our particular case study, is not general. It might be that different baselines, different goals or different setups would lead to laxer hardware requirements for the move scenario.

6.14.2 Architecture comparison

As discussed in detail in Section 6.6, the fiber network we study contains four nodes in the shortest path connecting the Dutch cities of Delft and Eindhoven. This means that there is some freedom in how to place the two heralding stations and repeater node required for a single-repeater setup, as shown in Figure 6.7. In order to decide how to make this placement, we determined the minimal hardware requirements for achieving an entanglement generation rate $R = 0.1$ Hz and a teleportation fidelity $F_T = 0.8717$, enabling VBQC between Delft and Eindhoven, for both possibilities. These requirements are shown in Figure 6.18. The requirements are qualitatively similar for both architectures, with the photon detection probability excluding attenuation losses and induced noise on memory




Figure 6.17: Directions along which color-center hardware must be improved to achieve entanglement generation rate $R = 0.1$ Hz and teleportation fidelity $F_T = 0.8717$, enabling VBQC between Delft and Eindhoven, assuming that a double-click entanglement generation protocol is used. The blue (orange) line corresponds to the direction of hardware improvement in case the Eindhoven end node (does not) move their half of the entangled state to the memory qubit. Note the use of a logarithmic scale.




Figure 6.18: Directions along which color-center hardware must be improved to achieve entanglement generation rate $R = 0.1$ Hz and teleportation fidelity $F_T = 0.8717$, enabling VBQC between Delft and Eindhoven, assuming that a double-click entanglement generation protocol is used. The blue (orange) line corresponds to the direction of hardware improvement for the architecture shown on the left (right) in Figure 6.7. Note the use of a logarithmic scale.

qubits (see Section 6.6 for details on our modeling of color-center based repeaters) being the parameters requiring the most improvement. The architecture on the left in Figure 6.7 required more modest improvements overall, so this was the architecture considered in our work.

6.14.3 Connecting Delft and Eindhoven without a repeater

The main contribution of this work was the investigation of the hardware requirements for enabling 2-qubit VBQC between two cities separated by 226.5 km of optical fiber using a single repeater node. We investigated two sets of performance targets compatible with this goal, namely (i) $R = 0.1$ Hz, $F_T = 0.8717$ and (ii) $R = 0.5$ Hz, $F_T = 0.8571$. While (ii) is impossible to achieve via direct transmission, i.e., without a repeater, due to fiber loss, this is not the case for (i) if a single-click entanglement generation protocol is employed. In Figure 6.19 we show directions along which color-center hardware would have to be improved to meet (i) without using a repeater. For comparison, we also reproduce the improvement directions for color-center hardware to meet the same targets with a

repeater employing double-click entanglement generation, because this was the repeater setup requiring smallest improvements as measured by our cost function.




Figure 6.19: Directions along which hardware must be improved to achieve entanglement generation rate $R = 0.1$ Hz and teleportation fidelity $F_T = 0.8717$, enabling VBQC between Delft and Eindhoven. The blue (orange) line corresponds to the direction of hardware improvement for the case in which a repeater is (is not) used. The repeater scenario employs a double-click entanglement generation protocol, whereas in the direct transmission case single-click entanglement generation is employed. Note the use of a logarithmic scale.

The direct transmission setup requires less improvement in all parameters. In fact, the only parameter that requires significant improvement is photon detection probability excluding attenuation losses, although still less than what is required for the repeater setup. The reason for this is that the elementary link state generated with the single-click protocol and state-of-the-art parameters already has high enough fidelity, so the only constraint is that these states are generated fast enough. The required value for the photon detection probability excluding attenuation losses is 0.39, less than the 0.73 required for the repeater with double-click entanglement generation case, but still above the limit imposed by the zero-phonon line.

These results indicate that performing VBQC over this particular setup might best be done without a repeater, but nevertheless do not detract from the main goal of the chapter, which was to investigate hardware requirements if a repeater were to be used.

6.14.4 Hardware requirements for repeaters with single and double-click entanglement generation

We investigated also how the hardware requirements for color centers running single and double-click entanglement generation protocols differ. We considered a rate target of $R = 0.1$ Hz and an average teleportation fidelity target of $F_T = 0.8717$. These two sets of hardware requirements are presented in Figure 6.20.




Figure 6.20: Hardware requirements for executing 2-qubit VBQC using a color-center repeater performing double-click (orange) and single-click entanglement generation (blue). These are the requirements for achieving an entanglement generation rate of $R = 0.1$ Hz and an average teleportation fidelity of $F_T = 0.8717$.

The hardware requirements are more stringent for a color-center repeater performing single-click entanglement generation. This is due to the fairly demanding fidelity target, which does not leave much room for noise in a protocol that inherently generates imperfect entangled states. We must however stress that this conclusion is specific to this particular setup and these performance targets, and does not imply that double-click should in general be chosen over single-click. In fact, one need only look at the second set of performance targets we considered in the main text to understand this point. These targets are impossible to achieve using a color-center repeater performing double-click entanglement generation, but are feasible if single-click is employed.

6.14.5 Hardware improvement costs

In Table 6.7 we present the cost of hardware improvement associated with the minimal hardware requirements found for every setup we investigated.

Platform	Target	Setup	Protocol	Cost
Color center	R = 0.1 Hz	Standard	Double-click, no-move	26.2
		Fiber network	Single-click, no-move	82.6
			Single-click, move	165.5
			Double-click, no-move	59.8
			Double-click, move	100.1
	Fiber network (repeaterless)	Single-click	20.5	
	Alternative fiber network	Double-click, no-move	116.1	
	R = 0.5 Hz	Fiber network	Single-click, no-move	153.3
Single-click, move			227.3	
Trapped ions	R = 0.1 Hz	Fiber network	Double-click	171.1
Abstract	R = 0.1 Hz	Fiber network, color center baseline	Double-click	40.7
		Fiber network, trapped ion baseline		50.1
		Fiber network, converted from color center		37.2
		Fiber network, converted from trapped ion		121.0

Table 6.7: Improvement cost, as defined in Section 6.11, of minimal hardware requirements for all setups we investigated.

References

- [1] J. Amirloo, M. Razavi, and A. H. Majedi, *Quantum key distribution over probabilistic quantum repeaters*, *Physical Review A* **82**, 032304 (2010).
- [2] F. K. Asadi, N. Lauk, S. Wein, N. Sinclair, C. O'Brien, and C. Simon, *Quantum repeaters with individual rare-earth ions at telecommunication wavelengths*, *Quantum* **2**, 93 (2018).
- [3] N. K. Bernardes, L. Praxmeyer, and P. van Loock, *Rate analysis for a hybrid quantum repeater*, *Physical Review A* **83**, 012323 (2011).
- [4] J. Borregaard, P. Komar, E. Kessler, A. S. Sørensen, and M. D. Lukin, *Heralded quantum gates with integrated error detection in optical cavities*, *Physical Review Letters* **114**, 110502 (2015).

- [5] D. E. Bruschi, T. M. Barlow, M. Razavi, and A. Beige, *Repeat-until-success quantum repeaters*, Physical Review A **90**, 032306 (2014).
- [6] Z.-B. Chen, B. Zhao, Y.-A. Chen, J. Schmiedmayer, and J.-W. Pan, *Fault-tolerant quantum repeater with atomic ensembles and linear optics*, Physical Review A **76**, 022329 (2007).
- [7] O. Collins, S. Jenkins, A. Kuzmich, and T. Kennedy, *Multiplexed memory-insensitive quantum repeaters*, Physical review letters **98**, 060502 (2007).
- [8] S. Guha, H. Krovi, C. A. Fuchs, Z. Dutton, J. A. Slater, C. Simon, and W. Tittel, *Rate-loss analysis of an efficient quantum repeater architecture*, Physical Review A **92**, 022357 (2015).
- [9] L. Hartmann, B. Kraus, H.-J. Briegel, and W. Dür, *Role of memory errors in quantum repeaters*, Physical Review A **75**, 032310 (2007).
- [10] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, *Quantum repeater with encoding*, Physical Review A **79**, 032325 (2009).
- [11] K. Nemoto, M. Trupke, S. J. Devitt, B. Scharfenberger, K. Buczak, J. Schmiedmayer, and W. J. Munro, *Photonic quantum networks formed from nv -centers*, Scientific reports **6**, 1 (2016).
- [12] M. Razavi, M. Piani, and N. Lütkenhaus, *Quantum repeaters with imperfect memories: Cost and scalability*, Physical Review A **80**, 032301 (2009).
- [13] M. Razavi and J. H. Shapiro, *Long-distance quantum communication with neutral atoms*, Physical Review A **73**, 042303 (2006).
- [14] C. Simon, H. De Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, *Quantum repeaters with photon pair sources and multimode memories*, Physical review letters **98**, 190503 (2007).
- [15] S. E. Vinay and P. Kok, *Practical repeaters for ultralong-distance quantum communication*, Physical Review A **95**, 052336 (2017).
- [16] Y. Wu, J. Liu, and C. Simon, *Near-term performance of quantum repeaters with imperfect ensemble-based quantum memories*, Physical Review A **101**, 042301 (2020).
- [17] N. Sangouard, C. Simon, J. Minář, H. Zbinden, H. De Riedmatten, and N. Gisin, *Long-distance entanglement distribution with single-photon sources*, Physical Review A **76**, 050301 (2007).
- [18] N. Sangouard, C. Simon, B. Zhao, Y.-A. Chen, H. De Riedmatten, J.-W. Pan, and N. Gisin, *Robust and efficient quantum repeaters with atomic ensembles and linear optics*, Physical Review A **77**, 062301 (2008).
- [19] J. Borregaard, H. Pichler, T. Schröder, M. D. Lukin, P. Lodahl, and A. S. Sørensen, *One-way quantum repeater based on near-deterministic photon-emitter interfaces*, Physical Review X **10**, 021071 (2020).

- [20] D. Luong, L. Jiang, J. Kim, and N. Lütkenhaus, *Overcoming lossy channel bounds using a single quantum repeater node*, Applied Physics B **122**, 1 (2016).
- [21] F. Rozpędek, K. Goodenough, J. Ribeiro, N. Kalb, V. C. Vivoli, A. Reiserer, R. Hanson, S. Wehner, and D. Elkouss, *Parameter regimes for a single sequential quantum repeater*, Quantum Science and Technology **3**, 034002 (2018).
- [22] F. Rozpędek, R. Yehia, K. Goodenough, M. Ruf, P. C. Humphreys, R. Hanson, S. Wehner, and D. Elkouss, *Near-term quantum-repeater experiments with nitrogen-vacancy centers: Overcoming the limitations of direct transmission*, Physical Review A **99**, 052330 (2019).
- [23] P. van Loock, W. Alt, C. Becher, O. Benson, H. Boche, C. Deppe, J. Eschner, S. Höfling, D. Meschede, P. Michler, F. Schmidt, and H. Weinfurter, *Extending Quantum Links: Modules for Fiber- and Memory-Based Quantum Repeaters*, Advanced Quantum Technologies **3**, 1900141 (2020).
- [24] L. Kamin, E. Shchukin, F. Schmidt, and P. van Loock, *Exact rate analysis for quantum repeaters with imperfect memories and entanglement swapping as soon as possible*, (2022), arXiv:2203.10318 .
- [25] S. Abruzzo, S. Bratzik, N. K. Bernardes, H. Kampermann, P. Van Loock, and D. Bruß, *Quantum repeaters and quantum key distribution: Analysis of secret-key rates*, Physical Review A **87**, 052315 (2013).
- [26] J. B. Brask and A. S. Sørensen, *Memory imperfections in atomic-ensemble-based quantum repeaters*, Physical Review A **78**, 012350 (2008).
- [27] S. Muralidharan, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, *Ultrafast and fault-tolerant quantum communication across long distances*, Physical review letters **112**, 250501 (2014).
- [28] M. Pant, H. Krovi, D. Englund, and S. Guha, *Rate-distance tradeoff and resource costs for all-optical quantum repeaters*, Physical Review A **95**, 012304 (2017).
- [29] T. D. Ladd, P. van Loock, K. Nemoto, W. J. Munro, and Y. Yamamoto, *Hybrid quantum repeater based on dispersive cqed interactions between matter qubits and bright coherent light*, New Journal of Physics **8**, 184 (2006).
- [30] P. Van Loock, T. Ladd, K. Sanaka, F. Yamaguchi, K. Nemoto, W. Munro, and Y. Yamamoto, *Hybrid quantum repeater using bright coherent light*, Physical review letters **96**, 240501 (2006).
- [31] M. Zwerger, B. Lanyon, T. Northup, C. Muschik, W. Dür, and N. Sangouard, *Quantum repeaters based on trapped ions with decoherence-free subspace encoding*, Quantum Science and Technology **2**, 044001 (2017).
- [32] L. Jiang, J. Taylor, and M. Lukin, *Fast and robust approach to long-distance quantum communication with atomic ensembles*, Physical Review A **76**, 012301 (2007).

- [33] X. Wu, A. Kolar, J. Chung, D. Jin, T. Zhong, R. Kettimuthu, and M. Suchara, *SeQUeNCe: A customizable discrete-event simulator of quantum networks*, *Quantum Science and Technology* **6**, 045027 (2021).
- [34] N. Kalb, P. C. Humphreys, J. Slim, and R. Hanson, *Dephasing mechanisms of diamond-based nuclear-spin memories for quantum networks*, *Physical Review A* **97**, 062330 (2018).
- [35] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, *Verifying bqp computations on noisy devices with minimal overhead*, *PRX Quantum* **2**, 040302 (2021).
- [36] J. F. Fitzsimons and E. Kashefi, *Unconditionally verifiable blind quantum computation*, *Physical Review A* **96**, 012303 (2017).
- [37] T. Morimae and K. Fujii, *Blind topological measurement-based quantum computation*, *Nature communications* **3**, 1036 (2012).
- [38] H.-L. Huang, Q. Zhao, X. Ma, C. Liu, Z.-E. Su, X.-L. Wang, L. Li, N.-L. Liu, B. C. Sanders, C.-Y. Lu, *et al.*, *Experimental blind quantum computing for a classical client*, *Physical review letters* **119**, 050503 (2017).
- [39] A. Gheorghiu, E. Kashefi, and P. Wallden, *Robustness and device independence of verifiable blind quantum computing*, *New Journal of Physics* **17**, 083040 (2015).
- [40] V. Dunjko, E. Kashefi, and A. Leverrier, *Blind quantum computing with weak coherent pulses*, *Physical review letters* **108**, 200502 (2012).
- [41] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal blind quantum computation*, in *2009 50th Annual IEEE Symposium on Foundations of Computer Science (IEEE, 2009)* pp. 517–526.
- [42] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, *Demonstration of blind quantum computing*, *science* **335**, 303 (2012).
- [43] C. H. Bennett, D. P. DiVincenzo, P. W. Shor, J. A. Smolin, B. M. Terhal, and W. K. Wootters, *Remote state preparation*, *Physical Review Letters* **87**, 077902 (2001).
- [44] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, *et al.*, *Heralded entanglement between solid-state qubits separated by three metres*, *Nature* **497**, 86 (2013).
- [45] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. Vermeulen, R. N. Schouten, C. Abellán, *et al.*, *Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres*, *Nature* **526**, 682 (2015).
- [46] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356**, 928 (2017).

- [47] P. C. Humphreys, N. Kalb, J. P. Morits, R. N. Schouten, R. F. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network*, *Nature* **558**, 268 (2018).
- [48] M. Pompili, S. L. Hermans, S. Baier, H. K. Beukers, P. C. Humphreys, R. N. Schouten, R. F. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, *et al.*, *Realization of a multinode quantum network of remote solid-state qubits*, *Science* **372**, 259 (2021).
- [49] S. Hermans, M. Pompili, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Qubit teleportation between non-neighbouring nodes in a quantum network*, *Nature* **605**, 663 (2022).
- [50] M. H. Abobeih, J. Cramer, M. A. Bakker, N. Kalb, M. Markham, D. J. Twitchen, and T. H. Taminiau, *One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment*, *Nature communications* **9**, 1 (2018).
- [51] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A ten-qubit solid-state spin register with quantum memory up to one minute*, *Physical Review X* **9**, 031045 (2019).
- [52] V. Krutyanskiy, M. Canteri, M. Meraner, J. Bate, V. Krcmarsky, J. Schupp, N. Sangouard, and B. P. Lanyon, *A telecom-wavelength quantum repeater node based on a trapped-ion processor*, (2022), arXiv:2210.05418 .
- [53] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, *Entanglement of trapped-ion qubits separated by 230 meters*, *Physical Review Letters* **130**, 050803 (2023).
- [54] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. Lanyon, *Light-matter entanglement over 50 km of optical fibre*, *npj Quantum Information* **5**, 1 (2019).
- [55] J. Schupp, V. Krcmarsky, V. Krutyanskiy, M. Meraner, T. Northup, and B. Lanyon, *Interface between Trapped-Ion Qubits and Traveling Photons with Close-to-Optimal Efficiency*, *PRX Quantum* **2**, 020331 (2021), publisher: American Physical Society.
- [56] V. Krutyanskiy, M. Meraner, J. Schupp, and B. P. Lanyon, *Polarisation-preserving photon frequency conversion from a trapped-ion-compatible wavelength to the telecom C-band*, *Applied Physics B* **123**, 228 (2017).
- [57] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas, *High-Fidelity Readout of Trapped-Ion Qubits*, *Physical Review Letters* **100**, 200502 (2008).
- [58] C. F. Roos, M. Chwalla, K. Kim, M. Riebe, and R. Blatt, *'Designer atoms' for quantum metrology*, *Nature* **443**, 316 (2006).
- [59] S. Baier, M. Galli, V. Krutyanskii, B. Lanyon, and T. Northup, private communications (2022).

- [60] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpędek, M. Skrzypczyk, *et al.*, *Netsquid, a network simulator for quantum information using discrete events*, Communications Physics **4**, 1 (2021).
- [61] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum repeaters: the role of imperfect local operations in quantum communication*, Physical Review Letters **81**, 5932 (1998).
- [62] M. Ruf, N. H. Wan, H. Choi, D. Englund, and R. Hanson, *Quantum networks based on color centers in diamond*, Journal of Applied Physics **130**, 070901 (2021).
- [63] M. Ruf, M. J. Weaver, S. B. van Dam, and R. Hanson, *Resonant excitation and purcell enhancement of coherent nitrogen-vacancy centers coupled to a fabry-perot microcavity*, Physical Review Applied **15**, 024049 (2021).
- [64] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt, *A quantum information processor with trapped ions*, New Journal of Physics **15**, 123012 (2013).
- [65] K. Mølmer and A. Sørensen, *Multiparticle entanglement of hot trapped ions*, Physical Review Letters **82**, 1835 (1999).
- [66] C. Cabrillo, J. I. Cirac, P. Garcia-Fernandez, and P. Zoller, *Creation of entangled states of distant atoms by interference*, Physical Review A **59**, 1025 (1999).
- [67] S. D. Barrett and P. Kok, *Efficient high-fidelity quantum computation using matter qubits and linear optics*, Physical Review A **71**, 060310 (2005).
- [68] G. Vardoyan, M. Skrzypczyk, and S. Wehner, *On the quantum performance evaluation of two distributed quantum architectures*, Performance Evaluation **153**, 102242 (2022).
- [69] M. Horodecki, P. Horodecki, and R. Horodecki, *General teleportation channel, singlet fraction, and quasidistillation*, Physical Review A **60**, 1888 (1999), publisher: American Physical Society.
- [70] L. Jiang, J. M. Taylor, N. Khaneja, and M. D. Lukin, *Optimal approach to quantum communication using dynamic programming*, Proceedings of the National Academy of Sciences **104**, 17291 (2007).
- [71] T. Coopmans, S. Brand, and D. Elkouss, *Improved analytical bounds on delivery times of long-distance entanglement*, Physical Review A **105**, 012608 (2022).
- [72] N. Sangouard, C. Simon, H. De Riedmatten, and N. Gisin, *Quantum repeaters based on atomic ensembles and linear optics*, Reviews of Modern Physics **83**, 33 (2011).
- [73] W. Dür and H. J. Briegel, *Entanglement purification and quantum error correction*, Reports on Progress in Physics **70**, 1381 (2007).

- [74] C. K. Hong, Z. Y. Ou, and L. Mandel, *Measurement of subpicosecond time intervals between two photons by interference*, Physical Review Letters **59**, 2044 (1987), publisher: American Physical Society.
- [75] F. Bouchard, A. Sit, Y. Zhang, R. Fickler, F. M. Miatto, Y. Yao, F. Sciarrino, and E. Karimi, *Two-photon interference: The Hong–Ou–Mandel effect*, Reports on Progress in Physics **84**, 012402 (2020).
- [76] W. Pfaff, B. J. Hensen, H. Bernien, S. B. van Dam, M. S. Blok, T. H. Taminiau, M. J. Tiggeleman, R. N. Schouten, M. Markham, D. J. Twitchen, *et al.*, *Unconditional quantum teleportation between distant solid-state quantum bits*, Science **345**, 532 (2014).
- [77] A. Stute, B. Casabone, P. Schindler, T. Monz, P. O. Schmidt, B. Brandstätter, T. E. Northup, and R. Blatt, *Tunable ion–photon entanglement in an optical cavity*, Nature **485**, 482 (2012).
- [78] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpędek, M. Pompili, A. Stolk, P. Pawelczak, R. Knegjens, J. de Oliveira Filho, R. Hanson, and S. Wehner, *A link layer protocol for quantum networks*, in *Proceedings of the ACM Special Interest Group on Data Communication, SIGCOMM '19* (Association for Computing Machinery, New York, NY, USA, 2019) pp. 159–173.
- [79] F. F. da Silva, A. Torres-Knoop, T. Coopmans, D. Maier, and S. Wehner, *Optimizing entanglement generation and distribution using genetic algorithms*, Quantum Science and Technology (2021).
- [80] J. Cramer, N. Kalb, M. A. Rol, B. Hensen, M. S. Blok, M. Markham, D. J. Twitchen, R. Hanson, and T. H. Taminiau, *Repeated quantum error correction on a continuously encoded qubit by real-time feedback*, Nature communications **7**, 1 (2016).
- [81] T. H. Taminiau, J. Cramer, T. van der Sar, V. V. Dobrovitski, and R. Hanson, *Universal control and error correction in multi-qubit spin registers in diamond*, Nature nanotechnology **9**, 171 (2014).
- [82] A. Reiserer, N. Kalb, M. S. Blok, K. J. van Bemmelen, T. H. Taminiau, R. Hanson, D. J. Twitchen, and M. Markham, *Robust quantum-network memory using decoherence-protected subspaces of nuclear spins*, Physical Review X **6**, 021040 (2016).
- [83] H. G. Barros, A. Stute, T. E. Northup, C. Russo, P. O. Schmidt, and R. Blatt, *Deterministic single-photon source from a single ion*, New Journal of Physics **11**, 103004 (2009).
- [84] B. Casabone, A. Stute, K. Friebe, B. Brandstätter, K. Schüppert, R. Blatt, and T. E. Northup, *Heralded Entanglement of Two Ions in an Optical Cavity*, Physical Review Letters **111**, 100505 (2013).
- [85] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther, *Continuous generation of single photons with controlled waveform in an ion-trap cavity system*, Nature **431**, 1075 (2004).

- [86] M. Meraner, A. Mazloom, V. Krutyanskiy, V. Krcmarsky, J. Schupp, D. Fioretto, P. Sekatski, T. E. Northup, N. Sangouard, and B. P. Lanyon, *Indistinguishable photons from a trapped-ion quantum network node*, Physical Review A **102**, 052614 (2020).
- [87] A. Stute, B. Casabone, B. Brandstätter, D. Habicher, H. G. Barros, P. O. Schmidt, T. E. Northup, and R. Blatt, *Toward an ion-photon quantum interface in an optical cavity*, Applied Physics B **107**, 1145 (2012).
- [88] T. Walker, S. V. Kashanian, T. Ward, and M. Keller, *Improving the indistinguishability of single photons from an ion-cavity system*, Physical Review A **102**, 032616 (2020).
- [89] T. Walker, K. Miyanishi, R. Ikuta, H. Takahashi, S. Vartabi Kashanian, Y. Tsujimoto, K. Hayasaka, T. Yamamoto, N. Imoto, and M. Keller, *Long-Distance Single Photon Transmission from a Trapped Ion via Quantum Frequency Conversion*, Physical Review Letters **120**, 203601 (2018).
- [90] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network*, Physical Review Letters **124**, 110501 (2020).
- [91] T. van Leent, M. Bock, F. Fertig, R. Garthoff, S. Eppelt, Y. Zhou, P. Malik, M. Seubert, T. Bauer, W. Rosenfeld, W. Zhang, C. Becher, and H. Weinfurter, *Entangling single atoms over 33 km telecom fibre*, Nature **607**, 69 (2022).
- [92] C. Crocker, M. Lichtman, K. Sosnova, A. Carter, S. Scarano, and C. Monroe, *High purity single photons entangled with an atomic qubit*, Optics Express **27**, 28143 (2019).
- [93] I. V. Inlek, C. Crocker, M. Lichtman, K. Sosnova, and C. Monroe, *Multispecies Trapped-Ion Node for Quantum Networking*, Physical Review Letters **118**, 250502 (2017).
- [94] D. P. Nadlinger, P. Drmota, B. C. Nichol, G. Araneda, D. Main, R. Srinivas, D. M. Lucas, C. J. Ballance, K. Ivanov, E. Y.-Z. Tan, P. Sekatski, R. L. Urbanke, R. Renner, N. Sangouard, and J.-D. Bancal, *Device-Independent Quantum Key Distribution*, Nature **607**, 682 (2022).
- [95] *Netsquid-Magic*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-magic> (2022).
- [96] J. Watrous, *The Theory of Quantum Information*, 1st ed. (Cambridge University Press).
- [97] C. H. Bennett, D. P. DiVincenzo, P. W. Shor, J. A. Smolin, B. M. Terhal, and W. K. Wootters, *Remote state preparation*, Phys. Rev. Lett. **87**, 077902 (2001).
- [98] J. L. W. V. Jensen, *Sur les fonctions convexes et les inégalités entre les valeurs Moyennes*, (1906).

- [99] M. D. Bowdrey, D. K. L. Oi, A. Short, K. Banaszek, and J. Jones, *Fidelity of single qubit maps*, *Physics Letters A* **294**, 258 (2002).
- [100] S. O. Hansson, *Do we need second-order probabilities?* *Dialectica* **62**, 525 (2008).
- [101] V. Bužek, M. Hillery, and R. F. Werner, *Optimal manipulations with qubits: Universal-NOT gate*, *Physical Review A* **60**, R2626 (1999).
- [102] A. Klappenecker and M. Rotteler, *Mutually unbiased bases are complex projective 2-designs*, in *Proceedings. International Symposium on Information Theory, 2005. ISIT 2005.* (2005) pp. 1740–1744.
- [103] S. D. Barrett and P. Kok, *Efficient high-fidelity quantum computation using matter qubits and linear optics*, **71**, 060310, publisher: American Physical Society.
- [104] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake, S. Vig, B. E. Granger, R. P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M. J. Curry, A. R. Terrel, v. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimrman, and A. Scopatz, *Sympy: symbolic computing in python*, *PeerJ Computer Science* **3**, e103 (2017).
- [105] G. Avis, F. Ferreira da Silva, T. Coopmans, A. Dahlberg, H. Jirovská, D. Maier, and J. Rabbie, *Simulation code for Requirements for a processing-node quantum repeater on a real-world fiber grid* · *GitLab*, <https://gitlab.com/softwarequtech/simulation-code-for-requirements-for-a-processing-node-quantum-repeater-on-a-real-world-fiber-grid>.
- [106] B. Kambs and C. Becher, *Limitations on the indistinguishability of photons from remote solid state sources*, **20**, 115003.
- [107] D. A. Fioretto, *Towards a flexible source for indistinguishable photons based on trapped ions and cavities*, .
- [108] *Smart-Stopos*, <https://gitlab.com/aritoka/smart-stopos>.
- [109] B. J. Jain, H. Pohlheim, and J. Wegener, *On termination criteria of evolutionary algorithms*, in *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation* (2001) pp. 768–768.
- [110] A. Labay Mora, *Genetic algorithm-based optimisation of entanglement distribution to minimise hardware cost*, (2021).
- [111] *Netsquid-NetConf*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-netconf> (2022).
- [112] *Netsquid-NV*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-nv> (2022).
- [113] *Netsquid-PhysLayer*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-physlayer> (2022).

- [114] *NetSquid-TrappedIons*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-trappedions> (2022).
- [115] *NetSquid-SimulationTools*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-simulationtools> (2022).
- [116] M. Pompili, C. Delle Donne, I. te Raa, B. van der Vecht, M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk, S. L. N. Hermans, P. Pawełczak, W. Kozłowski, R. Hanson, and S. Wehner, *Experimental demonstration of entanglement delivery using a quantum network stack*, npj Quantum Information **8**, 1 (2022).
- [117] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication*, Physical Review Letters **81**, 5932 (1998).
- [118] L.-M. Duan, M. Lukin, I. Cirac, and P. Zoller, *Long-distance quantum communication with atomic ensembles and linear optics*, Nature **414**, 413 (2001).
- [119] *NetSquid-AbstractModel*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-abstractmodel> (2022).

7

Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber

Francisco Ferreira da Silva¹, Guus Avis¹, Joshua A. Slater and Stephanie Wehner.

We perform a numerical study of the distribution of entanglement on a real-world fiber grid connecting the German cities of Bonn and Berlin. The connection is realized using a chain of processing-node quantum repeaters spanning roughly 900 kilometers. We investigate how minimal hardware requirements depend on the target application, as well as on the number of repeaters in the chain. We find that requirements for blind quantum computing are markedly different than those for quantum key distribution, with the required coherence time being around two and a half times larger for the former. Further, we observe a trade-off regarding how target secret-key rates are achieved when using different numbers of repeaters: comparatively low-quality entangled states generated at a high rate are preferred for higher numbers of repeaters, whereas comparatively high-quality states generated at a lower rate are favored for lower numbers of repeaters. To obtain our results we employ an extensive simulation framework implemented using NetSquid, a discrete-event simulator for quantum networks. These are combined with an optimization methodology based on genetic algorithms to determine minimal hardware requirements.

7

7.1 Introduction

In Chapter 6 we have determined hardware requirements for a single quantum repeater on a real-world fiber grid. In this chapter, we extend these results in a number of key

¹These authors contributed equally.

This chapter is based on the preprint arXiv:2303.03234.

ways. First, instead of only considering a single quantum repeater, we study chains of up to seven processing-node quantum repeaters. We do so using a fiber grid that will be used to construct a trusted-node network; upgrading such a network (which can be used for QKD albeit without end-to-end security) by replacing trusted nodes by repeaters may prove a particularly natural way of realizing early quantum-repeater networks [1]. Second, we investigate how the requirements on the quantum hardware change depending on how many repeaters are placed in the network. Finally, we also address the question whether the required hardware quality depends on the application that needs to be executed. Specifically, we consider two applications: QKD and VBQC, as discussed in Section 3.3.1.

7.1.1 Setup

We consider the quantum-network path depicted in Figure 7.1, with two end nodes situated in Bonn and Berlin separated by 917.1 km of optical fiber corresponding to 214.7 dB of attenuation (at a telecom wavelength of 1550 nm). There are a total of sixteen locations be-




Figure 7.1: Map of Germany overlaid with a depiction of the fiber path connecting the German cities of Bonn and Berlin that we investigate, provided by Deutsche Telekom (DT). The white circles represent locations where DT plans to install trusted nodes and where, when building a repeater chain, processing nodes or heralding stations could be placed. These locations are connected to each other through fiber drawn in black. The maximum number of repeaters that can be placed between Bonn and Berlin in this fiber network is seven. We consider all possible repeater placements, assuming that the heralding stations are placed as symmetrically as possible (there are 986 such placements). The distance between Bonn and Berlin is 917.1 km via fiber, and approximately 480 km as the crow flies. The reason for such a large difference between the two values is that other major German cities, such as Hannover and Dortmund, are connected through the fiber link as well.

tween the end nodes where equipment can be placed, namely repeater nodes and heralding stations. Throughout this paper we assume that such a heralding station must be placed between every pair or neighboring network nodes (i.e., end nodes or repeater nodes), as these are required when entanglement is generated between those nodes through the interference and measurement of entangled photons [2–9]. This data has been provided

to us by Deutsche Telekom (DT), Germany's largest telecommunications provider, which plans to install trusted nodes in the locations depicted in Figure 7.1.

We assume neighboring nodes perform heralded entanglement generation [10, 11]. That is, entanglement consists of a series of attempts, and at the end of each attempt the partaking nodes learn whether an entangled state was successfully created or not. Examples of protocols for heralded entanglement generation are the double-click protocol [3, 6–9, 12], where photons are interfered and measured at a heralding station and success is declared in case two detectors click, the single-click protocol [2, 4, 5, 13, 14], where photons are also interfered but success is only declared in case one detector clicks, and direct transmission of an entangled photon from one node to the next where it is stored in heralded quantum memory [15–17]. Here, we employ a simplified model for heralded entanglement generation. We do this so that the protocol and its interplay with other components of the repeater chain can be readily understood and our modelling is not overly platform specific. First of all, we assume that each node can perform heralded entanglement generation with two neighbours in parallel, which is not currently possible for all quantum-repeater platforms [18]. Second, we model the elementary-link states ρ that are created upon the completion of a successful attempt as depolarized Bell states, i.e.,

$$\rho = W|\phi^+\rangle\langle\phi^+| + \frac{1-W}{4}\mathbb{1}, \quad (7.1)$$

where W is related to the fidelity F to the ideal Bell state $|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ as $F = (1 + 3W)/4$ and $\mathbb{1}$ is the four-dimensional identity matrix. We note that real entangled states generated in quantum-repeater chains are often not depolarized states [19] (see also Chapter 6). Yet, as depolarized Bell states represent a worst-case type of noise [20], using a depolarizing model ensures that we will not find hardware requirements that are artificially low due to this simplification. Third, we take the time t_{attempt} required to perform one attempt to be given by

$$t_{\text{attempt}} = \frac{L}{c}, \quad (7.2)$$

where L is the fiber distance between the two nodes and $c = 2.14 \times 10^5$ km/s is the speed of light in fiber. That is, it corresponds to the communication time associated with sending photons to a heralding station that is exactly in the center between two nodes and then receiving a message with the measurement outcome. This is equivalent to the time required to directly transmit a photon from one node to the next. In reality it may be longer, as the attempt time could be further limited by the rate of the photon source, local operations or the synchronization of emission times [5, 21]. Finally, we take the success probability p_{el} of each attempt to be

$$p_{\text{el}} = p_{\text{det}} \times 10^{-\frac{\alpha_{\text{att}} L}{10}}. \quad (7.3)$$

Here, p_{det} is the probability that an emitted photon that is led through fiber to a detector is detected, given that it is not lost while travelling in fiber. This parameter combines multiple sources of loss, such as the detector's efficiency, the probability of emitting the photon in the right mode and the probability of successfully sending the photon into the fiber, but not the fiber's attenuation losses. α_{att} is the fiber's attenuation coefficient (in

dB/km). Therefore, the success probability corresponds to the success probability of directly transmitting a photon between the nodes and measuring it there. We note that for the double-click protocol the scaling with length would be the same, although the prefactor would be different (p_{det}^2 instead of p_{det} , as two photons must be detected). For the single-click protocol the scaling would be more gentle (roughly replacing L by $L/2$ in the exponential), and while the prefactor would be linear in p_{det} , there would also be a factor that depends on the device settings (specifically on the bright-state parameters chosen at both nodes, which tune a trade-off between success probability and state fidelity [5, 22]). Additionally, we allow also for the possibility of multiplexed heralded entanglement generation [23–25]. This essentially consists of performing multiple attempts of generating the same elementary-link state in parallel. Multiplexing can be done across multiple degrees of freedom, such as frequency, time or space. We remain agnostic regarding how the multiplexing is performed, including it in our model as one parameter corresponding to the number of multiplexing modes used, n . The probability of successfully generating an elementary link assuming the use of multiplexing is then the probability that at least one of the multiplexing modes succeeds:

$$p_{\text{multiple modes}} = 1 - (1 - p_{\text{el}})^n. \quad (7.4)$$

The nodes implement a swap-asap protocol [26, 27]. That is, as soon as a node holds two entangled states, one shared with each of its neighbours, it performs an entanglement swap in order to create an entangled state spanning a larger distance. We assume this swap is realized deterministically, since we are modelling processing nodes that can implement a swap using quantum gates and measurements on their processors. It may however introduce noise, which we model as depolarizing. We quantify how well the swap can be performed using the swap-quality parameter s_q . The d -dimensional depolarizing noise channel of parameter p acts on a state ρ as follows,

$$\rho \rightarrow p\rho + (1 - p)\frac{\mathbb{1}}{d}. \quad (7.5)$$

This means that, with probability p , ρ is left unchanged, and with probability $1 - p$ it is mapped to the maximally-mixed state, i.e., all information is lost. Then, we model entanglement swapping as a two-qubit depolarizing channel (i.e., $d = 4$) with parameter $p = s_q$ followed by a perfect entanglement-swapping operation (i.e., a measurement in the Bell basis [28]). We assume that the gates and measurements applied by the end nodes when executing QKD and VBQC are noiseless and instantaneous. States stored in memory undergo decoherence, which we model as exponential depolarizing noise, i.e.,

$$\rho \rightarrow e^{-t/T}\rho + (1 - e^{-t/T})\frac{\mathbb{1}}{d}, \quad (7.6)$$

where t is the time for which the state ρ has been held in memory and T is the memory's coherence time. To combat the effects of memory decoherence, entangled states are discarded after a local cut-off time. The cut-off time is defined as follows: a timer starts once a state is created in memory through the successful generation of an elementary link. If the timer reaches the local cut-off time, the state is discarded. That is, the qubit holding

the state is reset. Additionally, the node sends a classical message along the chain so that the qubit with which the first qubit was entangled can also be reset. As a result, a number of elementary links in the chain must be regenerated (with the exact number depending on how far away the entangled qubit was).

7.1.2 Applications

Having discussed our modelling of the entanglement generation process between Bonn and Berlin, we turn to the applications that will make use of the entanglement, QKD and VBQC. We investigate the BB84 QKD protocol [29] (in its entanglement-based form [30]) between the end nodes situated in Bonn and Berlin. We record the entanglement generation rate and estimate the quantum bit error rate (QBER) that would have been obtained when measuring the generated state in order to estimate the achievable asymptotic secret-key rate (SKR) as per the following equation [31]:

$$\text{SKR} = E_R \cdot \max\left(0, (1 - 2H(Q))\right), \quad (7.7)$$

where E_R is the entanglement-generation rate, $H(p) = -p \log_2(p) - (1-p) \log_2(1-p)$ is the binary entropy function and Q is the QBER. We note that all the noise sources we consider are depolarizing, hence the entangled states generated will be of the form of the state shown in Equation 7.5. Therefore, the QBER is the same irrespective of the measurement basis. The end nodes do not wait until end-to-end entanglement is established before measuring their qubits. Instead, they measure them as soon as they have established entanglement with their nearest neighbours, as this minimizes the amount of time states spend in memory, resulting in laxer hardware requirements.

We also investigate a two-qubit version of the VBQC protocol introduced in [32]. In such protocols, a client wishes to delegate a computation to a powerful remote server in a secure and verifiable fashion [33]. In particular, we investigate the repeated execution of test rounds of the protocol, which consist of the server performing a controlled-Z gate followed by a measurement. In these rounds the client knows the computation's expected outcome, and can therefore compare them to the observed outcomes. Under the assumption of an honest server, wrong outcomes are a result of noise. We call this the *BQC test protocol*. The fraction of successful BQC test protocol rounds is therefore a metric for the quality of the entanglement used for transmitting qubits. We define the success rate as the number of rounds of the protocol that can be executed with a successful result per time unit. More concretely, if p_s is the success probability of a test round and R_{rounds} is the rate at which rounds can be executed, the BQC-test-protocol success rate is given by:

$$R_{\text{BQC}} = R_{\text{rounds}} \cdot p_s. \quad (7.8)$$

While the BQC test protocol is in and of itself not an interesting application of a quantum network, it can be considered a benchmark for how well the network is suited to VBQC and possibly other multi-qubit applications. The fact that, in contrast with QKD, it requires the distribution of multiple entangled states and the storage of qubits between rounds makes it a more meaningful benchmark for quantum-network applications that require multiple live qubits contemporaneously. Further details on the BQC test protocol can be found in Section 7.10.

The two applications we have just introduced define our performance targets.

7.1.3 Minimal hardware requirements

We wish to find the *minimal hardware requirements* that are needed to realize different target SKRs and BQC-test-protocol success rates. These correspond to the minimal improvements over state-of-the-art hardware parameters that enable meeting the targets. We phrase the problem of finding minimal hardware requirements as a constrained optimization problem. Namely, we wish to minimize the hardware improvement while ensuring that the constraint induced by the performance target is met. This constraint is relaxed through a process known as scalarization [34, 35], resulting in a single-objective optimization problem, in which the quantity to be minimized is the sum of the cost associated to the hardware improvement and a penalty term for the rate target. The resulting cost function is given by:

$$\begin{aligned}
 C = & w_1 \left(1 + (R_{target} - R_{real})\right)^2 \\
 & \cdot \Theta(R_{target} - R_{real}) \\
 & + w_2 H_C(x_1, \dots, x_N),
 \end{aligned} \tag{7.9}$$

where H_C is the hardware improvement cost associated to parameter set $\{x_1, \dots, x_N\}$, w_i are the weights assigned to the objectives, Θ is the Heaviside step function, R_{target} is the rate target and R_{real} is the rate of application execution achieved by the parameter set. We note that R_{real} and R_{target} can be either a SKR or a BQC-test-protocol success rate. H_C maps sets of hardware parameters to a number, the cost, which represents how large of an improvement over the state-of-the-art they represent. In order to compute this cost consistently across different parameters, we use no-imperfection probabilities as done in Chapter 6. By no-imperfection probability, we mean the probability that there is no error or loss associated to a given parameter. For example, the no-error probability associated to a photon detection probability p_{det} (defined in Section 7.1.1) of 0.1 is 0.1. For the no-error probabilities associated to the other hardware parameters, see Section 7.9.1. We say that a parameter is improved by a factor of k if its no-imperfection probability becomes $\sqrt[k]{p_{ni}}$, with p_{ni} being the state-of-the-art no-imperfection probability. For example, improving the no-imperfection probability of 0.1 associated to $p_{det} = 0.1$ by a factor of 5, we get a no-imperfection probability of ≈ 0.63 , corresponding to $p_{det} \approx 0.63$. The hardware cost associated to a set of parameters is given by the sum of the improvement factors of the parameters. The weights w_i are chosen such that the term of the overall cost function corresponding to meeting the rate target is always larger than the one corresponding to the hardware cost, ensuring that even though we have relaxed the constraints by scalarizing, we are still effectively requiring that the minimal hardware requirements are such that the performance target is met. To ensure this, we picked $w_1, w_2 \gg w_3$, such that $w_1 (1 + (R_{target} - R_{real})^2) \Theta(R_{target} - R_{real}) \gg w_2 H_C(x_1, \dots, x_N)$. Specifically, we set $w_1 = 1 \times 10^{100}$ and $w_2 = 1$. No particular heuristic was used to select these numbers.

We note that the hardware cost is meant only to represent a measure of the hardness of improving the hardware to a certain level, and not any form of monetary cost. At present quantum repeater systems are research setups, with commercial solutions only starting to emerge. Therefore, assigning any specific commercial cost numbers would be too speculative at this point, and would require an in-depth study outside the scope of this project.

7.1.4 State-of-the-art parameters

Computing minimal hardware requirements as described in Section 7.1.3 is done with respect to a baseline over which we are improving. In this work, this baseline consists of parameters measured for color centers in diamond, as they are physical systems using which various quantum-networking primitives have been demonstrated. These include long-lived quantum memories [36], remote entanglement generation [6, 7], quantum teleportation [12], entanglement distillation [13], entanglement swapping [37] and a three-node network [5]. We do not impose that all parameters must have been demonstrated in the same experiment or even with the same color center. The parameters we consider are shown in Table 7.1. Details on how these parameters were determined can be found

Parameter	Value
Coherence time	1 s [38]
Number of multiplexing modes	1
Fidelity of elementary links	0.83 [7]
Photon detection probability	0.255 [15]
Swap quality	0.83 [13, 39]

Table 7.1: State-of-the-art color-center parameters. We note that not all of these parameter values have been realized in a single experiment. We have number of modes as 1 without reference because to the best of our knowledge multiplexed entanglement generation has not been demonstrated using color centers.

in Section 7.7.

7.1.5 Determining minimal hardware requirements

In order to determine minimal hardware requirements, we need to (i) be able to evaluate how a given set of hardware parameters performs and (ii) optimize over the parameter space to find the parameters that minimize the requirements while still performing adequately (i.e., the parameters that minimize the cost function defined in Equation (8.17)).

We evaluate the performance of hardware parameters using general processing-node repeater-chain simulations developed in NetSquid, a discrete-event based quantum-network simulator [26]. The simulations are general in the sense that they can be used to investigate swap-asap repeater chains of arbitrary size and spacing (i.e., nodes and heralding stations need not be equidistant). They take into account time-dependent noise, classical control communication and the constraints imposed by a real-world fiber network. The code for executing such simulations has been made publicly available at [40] and is largely based on the simulations presented in Chapter 6. Our code that utilizes these simulations to produce the results here presented can be found at [41] (and the corresponding data at [42]).

Given that we can evaluate the performance of any parameter set on the Bonn-Berlin path, we perform parameter optimization using a genetic algorithm [43] to minimize the cost function defined in Section 7.1.3 using the high-performance computing cluster Snelius. For further details, see Section 7.9.

7.2 Impact of number of repeaters on hardware requirements

In this section we answer the question of how hardware requirements are affected by the number of repeaters deployed in a quantum network. Specifically, we investigate the minimal hardware requirements for performing BB84 between the German cities of Bonn and Berlin at a key rate of 10 Hz. We assume the cities are connected by the network path shown in Figure 7.1. We determine what these minimal requirements are in two cases: (i) optimizing over the number of repeaters and (ii) restricting the number of repeaters to specific values. In both cases we optimize over the placement of the repeaters.

7.2.1 Absolute minimal number of multiplexing modes

Before determining minimal requirements, we aim to answer the question of what are the absolute minimal number of multiplexing modes required to perform QKD between the German cities of Bonn and Berlin at rates of 1, 10 and 100 Hz. By *absolute minimal* number of multiplexing modes, we mean the minimum number of multiplexing modes that is required if the only source of imperfection in the setup is fiber attenuation. This provides a lower bound on the number of multiplexing modes in the minimal hardware requirements, as the introduction of other hardware imperfections can only lead to more stringent demands on the number of modes. We emphasize that for the purposes of answering this question we are temporarily setting aside the real-world fiber path introduced in Figure 7.1. Instead, we are going to consider a *symmetrized* version of that path. By this we mean a path with the same total length and attenuation, but in which nodes and heralding stations are placed equidistantly, and where the attenuation is evenly distributed throughout the path, i.e., all elementary links have the same attenuation. The reason for doing so is that the minimal number of modes for this path is a lower bound for the same quantity on any other path with the same total length and attenuation. To see this, we note that it has been shown that repeater chains of the type studied here perform best when all nodes are positioned as symmetrically as possible [44]. This implies that such a chain will have less stringent hardware requirements to attain a given performance target in comparison to chains which are subject to real-world restrictions such as the ones imposed by the fiber path shown in Figure 7.1, and, therefore, also less stringent requirements on the number of multiplexing modes.

Determining the absolute minimal number of multiplexing modes serves two purposes. First, it allows us to limit the search space of the optimization we run for finding minimal hardware requirements. Second, it gives us a general idea of how many repeaters might be required to achieve the target with reasonable hardware demands. For example, if for a specific number of repeaters hundreds of thousands of multiplexing modes are required to meet the target without any noise sources, that may indicate that using that number of repeaters is not practically feasible.

In Figure 7.2 we show the absolute minimal number of modes required to distribute secret key at rates of 1, 10 and 100 Hz using BB84 in the symmetrized Bonn - Berlin path. We find that more multiplexing modes are required for higher rate targets, and that this number grows superexponentially as the number of repeaters decreases, so as to counteract the effects of photon loss in fiber. Further, we find that achieving a SKR of 10 Hz with




Figure 7.2: Minimal number of multiplexing modes required to achieve 1, 10 and 100 Hz of SKR over 917.1 km of fiber with a total of 214.7 dB of attenuation, corresponding to a symmetrized version of the path between Bonn and Berlin that we investigate. That is, for N repeaters, the symmetrized path has $N + 1$ elementary links, each of length $917.1/(N + 1)$ km and of attenuation $214.7/(N + 1)$ dB. We assume that there are no hardware imperfections, and that repeaters are uniformly spaced.

fewer than 3 repeaters requires hundreds of thousands of multiplexing modes even in the absence of any sources of noise. As the hardware cost (defined in Section 7.1.3) associated with so many multiplexing modes far outweighs typical values for the minimal total hardware cost found for three or more repeaters we limit the rest of our investigation to configurations with three or more repeaters.

7.2.2 Minimal hardware requirements for quantum-key distribution

We now turn our attention to the minimal hardware requirements for performing quantum-key distribution at a rate of 10 Hz using the BB84 protocol. In particular, we investigate them along the path connecting Bonn and Berlin depicted in Figure 7.1. As Figure 7.2 illustrates, the number of repeaters used can have a considerable impact on the hardware requirements. Further, it is expected that the same is true for the placement of repeaters and heralding stations (see Chapter 6 and [44]). With this in mind, we ask two questions: (i) what are the minimal hardware requirements when allowing for the placement of up to the largest number of repeaters that fits in the fiber path (seven) and (ii) what are the minimal hardware requirements when restricting the maximum number of repeaters to five. We expect that this will lead to different parameter regimes, illustrating two possible directions towards achieving the target performance.

In Figure 7.3 we show the directions along which hardware must be improved for distributing secret key at rates of 10 Hz using BB84 in the network path connecting Bonn and Berlin. The corresponding minimal hardware requirements can be found in Table 7.2. In each case we find that the hardware requirements are minimized when the number of repeaters used is maximized. That is, for seven repeaters in case (i) and five repeaters in

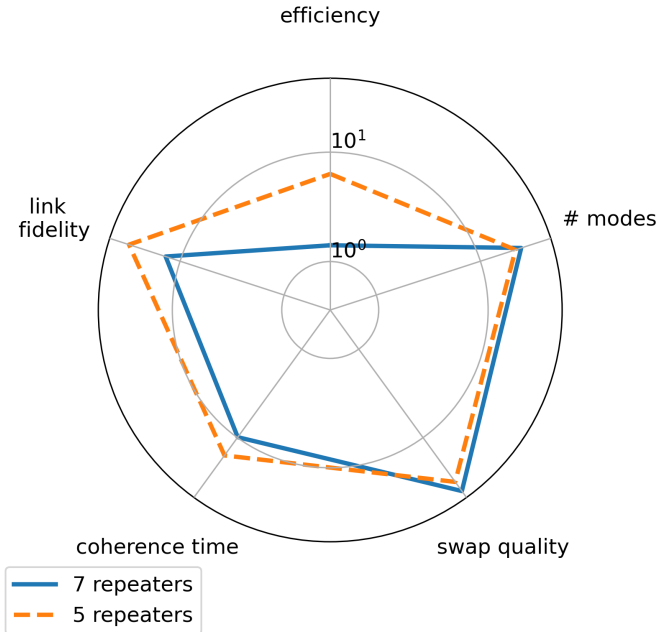


Figure 7.3: Directions along which hardware must be improved to enable attaining a secret-key rate of 10 Hz between the German cities of Bonn and Berlin. The blue (orange) line was obtained by performing an optimization in which the algorithm was allowed to use a maximum of seven (five) repeaters. The further away the line is from the center of the plot towards a given parameter, the more that parameter must be improved with respect to the current state-of-the-art. Improvement is depicted for the following parameters, clockwise from the top: overall photon detection probability excluding attenuation in fiber, number of multiplexing modes, fidelity of entanglement swap, coherence time of memory qubits and fidelity of elementary links. Note the use of a logarithmic scale.

7

Application	QKD				BQC
	1	10	100	10	
Rate (Hz)	1	10	100	10	
Number of repeaters	7	Max 5	Max 7	7	
Coherence time (s)	1.81	4.23	3.14	10.1	
Number of multiplexing modes	175	544	592	799	
Fidelity of elementary links	0.989	0.995	0.987	0.996	
Photon detection probability p_{det}	0.604	0.785	0.360	0.804	
Swap quality	0.996	0.996	0.997	0.997	

Table 7.2: Minimal hardware requirements to achieve 1, 10 and 100 Hz of secret-key rate and 10 Hz of blind quantum computing test protocol success rate between the German cities of Bonn and Berlin. The photon detection probability p_{det} is the probability of a photon being detected given that it is not lost in fiber. It combines multiple sources of loss, such as the detector's efficiency, the probability of emitting the photon in the right mode and the probability of successfully sending the photon into the fiber. More details can be found in Section 7.1.1 and 7.7.

case (ii). Hardware requirements are more stringent in case fewer repeaters are used. In particular, the overall photon detection probability excluding attenuation in fiber must be improved to a much larger degree (0.79 vs 0.36) if only five repeaters are used. This is needed to overcome the increased attenuation losses associated with the longer elementary links. The coherence time required when using five repeaters is also larger than the time required when using seven repeaters (4.2 s vs 3.1 s). This can be explained by the fact that keeping the entanglement-generation rate high is more costly in case of five repeaters. Therefore keeping the QBER small to extract as many secret bits as possible from each entangled state is more valuable. Furthermore, since the entanglement-generation rate is smaller for five repeaters, qubits are stored for longer times before they can be swapped and hence a larger coherence time is required to achieve the same QBER. We study this interplay further in Section 7.2.3. Finally, we notice that while the requirements on most hardware parameters are more stringent for five repeaters as compared to seven repeaters, this is not the case for the requirement on the swap quality. In fact, the requirement on the swap quality is even slightly looser for five repeaters (0.996 vs 0.997). This is explained by the fact that when there are more repeaters, there are more entanglement swaps associated with every end-to-end entangled state. Therefore, when there are more repeaters the final error rate is more sensitive to noise in the swaps, creating a larger incentive to improve the associated parameter in the seven-repeater case as compared to the five-repeater case.

7.2.3 Secret-key rate: quantum-bit error rate and entanglement generation rate

A specific value for the SKR can be obtained through many different pairs of values for the entanglement-generation rate and the QBER, as follows from Equation (7.7). This opens up a trade-off between the entanglement generation rate and the QBER, as briefly discussed in Section 7.2.2. Here, we investigate this trade-off more deeply by repeating our process for determining minimal hardware requirements to achieve an SKR of 10 Hz while keeping the number of repeaters a fixed parameter. We did this for 4, 5, 6 and 7 repeaters. For each case, we still optimize over all possible placements of the repeaters in the fiber grid. In Figure 7.4 we show the QBER and entanglement-generation rate achieved with the minimal hardware requirements for the best setup found by our optimization procedure for varying number of repeaters. We observe two different regimes. For 4 and 5 repeaters, which we name the ‘few-repeater’ regime, we find a low QBER (~ 5%) and an entanglement-generation rate of 20 - 30 Hz. On the other hand, for 6 and 7 repeaters, i.e., the ‘many-repeater’ regime, we find a comparatively higher QBER (~ 9%) and an entanglement generation rate of almost 80 Hz. In other words, in the many-repeater regime, distributing many entangled pairs of comparatively lower quality requires less hardware improvement. On the other hand, in the few-repeater regime it seems to be more feasible to distribute fewer pairs of comparatively higher quality. As the number of repeaters used decreases, it becomes harder to overcome the effect of fiber attenuation, which makes improving the quality of the entangled states delivered a more attractive option for increasing the SKR.

We finalize by remarking that, perhaps surprisingly, the variance in the time it takes to distribute one entangled state appears to grow as the number of repeaters in the chain increases (as shown by the increasing error bar on the rate in Figure 7.4). While interesting,

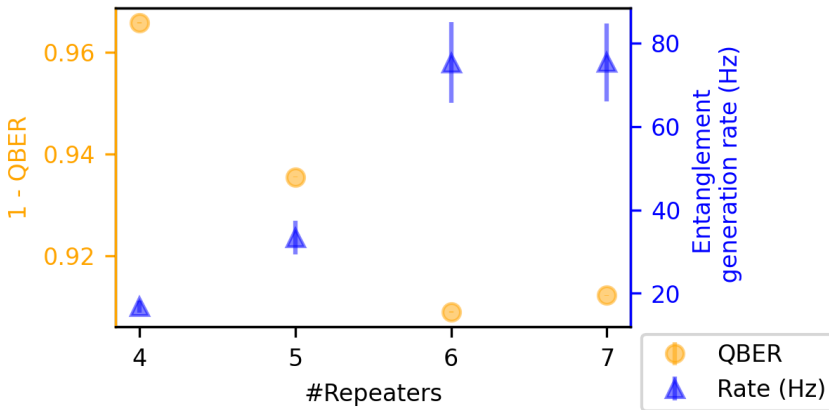


Figure 7.4: QBER and entanglement generation rate obtained with the minimal hardware requirements to achieve 10 Hz of SKR in the Bonn - Berlin setup with different numbers of repeaters, up to seven, the maximum allowed in the setup we study. The error bars are given by the standard error of the mean. Each data point corresponds to 2000 simulations of an entanglement-based BB84 protocol.

further investigation is beyond the scope of this work.

We show the repeater placement corresponding to the minimal hardware requirements found when optimizing over the number of repeaters and their placement in Section 7.8.

7

7.3 Impact of target on hardware requirements

We now turn our attention to the impact of the performance target on the hardware requirements. We approach this from two angles: (i) the impact of varying the SKR target and (ii) the impact of holding the required rate constant but changing the target quantum-network application. It is clear that, given the same repeater chain, increasing the target rate will lead to more stringent requirements. However, it is not *a priori* obvious if the relative importance of the different parameters will change as the target rate is increased. It is further not obvious how changing the target application impacts the hardware requirements. These are questions of practical relevance: given that one wishes to build a repeater chain capable of distributing entanglement to perform QKD at a rate of 100 Hz, it seems crucial to know whether building a repeater chain for performing QKD at a rate of 1 Hz is a step in the right direction. In other words, this investigation can shed light on whether the process of improving hardware for quantum-repeater chains should be approached incrementally. The same question holds for the different target applications. It is likely that quantum repeaters will initially be used for QKD as they begin to replace their trusted-node predecessors, and only progressively start to be used for applications that require multiple live qubits. We would then like to know whether the hardware improvements necessary to perform QKD using quantum repeaters are similar to the ones for multi-qubit applications.

7.3.1 Requirements for different secret-key-rate targets

In Figure 7.5 we show the directions along which hardware must be improved for distributing secret key at rates of 1, 10 and 100 Hz using BB84 in the network path connecting Bonn and Berlin. The corresponding minimal hardware requirements can be found in Table 7.2.




Figure 7.5: Directions along which hardware must be improved to enable attaining secret-key rates of 1 (blue, full), 10 (orange, dashed) and 100 Hz (green, dotted) between the German cities of Bonn and Berlin. The further away the line is from the center of the plot towards a given parameter, the more that parameter must be improved with respect to the current state-of-the-art. Improvement is depicted for the following parameters, clockwise from the top: overall photon detection probability excluding attenuation in fiber, number of multiplexing modes, fidelity of entanglement swap, coherence time of memory qubits and fidelity of elementary links. Note the use of a logarithmic scale.

The hardware requirements become more stringent as the SKR target grows. Further, the coherence time requires significantly less improvement in the 1 Hz case when compared to 10 and 100 Hz. This comes as something of a surprise, given that we expect qubits to spend less time in memory for higher SKR values, as these should correspond to higher entanglement-generation rates (and hence lower waiting times). In order to further investigate why this happens, we show in Figure 7.6 the QBER and entanglement generation rate achieved with the minimal hardware requirements for the best setup found by our optimization procedure for different SKR targets. We find that both the entanglement generation rate and $1 - \text{QBER}$ increase with the target SKR. We conjecture that the increase in coherence time observed for higher SKR targets is due to the necessary entanglement generation rate being very high. In fact, it is so high that it requires a huge number of multiplexing modes, which in turn imply a very high cost. This makes it comparatively less costly to extract more key from each entangled state than to generate states faster.




Figure 7.6: QBER and entanglement generation rate obtained with the minimal hardware requirements to achieve 1, 10 and 100 Hz of SKR in the Bonn - Berlin setup using the configuration found to be optimal for 10 Hz. The error bars are given by the standard error of the mean. Each data point corresponds to 2000 simulations of an entanglement-based BB84 protocol.

7.3.2 Requirements for secret-key and blind-quantum-computing success rates

In Figure 7.7 we show the directions along which hardware must be improved for performing QKD and BQC at a rate of 10 Hz in the network path connecting Bonn and Berlin. The corresponding minimal hardware requirements can be found in Table 7.2. It is plain to see that the two applications require improvements in distinct parameters. In particular, we emphasize the much larger coherence time required for BQC, corresponding to roughly a factor of 2.5 difference (7.99 vs 3.14 seconds). This can be explained by the fact that BQC, unlike QKD, requires two entangled pairs to be alive at the same time, implying that one entangled pair must be stored at the end nodes while the second one is generated. Further, the fact that the minimal coherence time required for BQC is high means that comparatively less noise will be caused by decoherence. This, in turn, means that in order to achieve the same state quality, the swap quality and the elementary link fidelity need not be as good.

We have also observed that there is a significant difference in the entanglement generation rates achieved by the parameter sets corresponding to the improvements shown in Figure 7.7. The minimal hardware requirements for QKD achieve an entanglement generation rate of almost 80 Hz, whereas the ones for the BQC-test-protocol result in an entanglement generation rate of around 20 Hz. In the same vein as what was discussed in Section 7.2.3, this is a result of the SKR and the BQC test protocol success rate being composite quantities, depending not only on the rate at which entangled states are delivered, but also on the quality of these states. We believe that the difference observed in entanglement generation rate between the two applications is due to the fact that there is a threshold state quality to obtain non-zero secret-key ($\sim 11\%$ QBER or equivalently ~ 0.84 fidelity, both under the assumption of depolarizing noise). Such a threshold does not exist for the BQC test protocol. This fundamental difference means that the state quality requirements are more stringent in the QKD case, making improving the entanglement




Figure 7.7: Directions along which hardware must be improved to enable attaining secret-key (QKD, blue) and blind quantum computing (BQC, orange) test protocol rates of 10 Hz between the German cities of Bonn and Berlin. The further away the line is from the center of the plot towards a given parameter, the more that parameter must be improved with respect to the current state-of-the-art. Improvement is depicted for the following parameters, clockwise from the top: overall photon detection probability excluding attenuation in fiber, number of multiplexing modes, fidelity of entanglement swap, coherence time of memory qubits and fidelity of elementary links. Note the use of a logarithmic scale.

generation rate a more attractive possibility. We do however note that even though the BQC test protocol does not impose a threshold on state quality, the complete VBQC protocol proposed in [32] does.

7.4 Conclusion

We have determined minimal hardware requirements for generating entanglement between two nodes separated by roughly 900 km of real-world optical fiber using a chain of processing-node quantum repeaters. We investigated both how such requirements depend on how many repeaters are employed and on the quantum-network application for which the entanglement is used. Notably, we have found that the hardware requirements for performing quantum key distribution and a simplified form of blind quantum computing are qualitatively different, with blind quantum computing requiring a coherence time which is roughly a factor of 2.5 larger for the same target rate in the setup we investigated. We further observed that given that most metrics one is interested in when evaluating quantum-network performance depend on both the rate at which entanglement is generated and its quality, there is room for trade-offs: for example, we found that when

employing a large number of repeaters to achieve a given secret-key rate in the setup we studied it is easier to generate many entangled pairs of comparatively lower quality, with the opposite being true if fewer repeaters are used.

The blind-quantum-computing requirements we determined were obtained for a simplified form of the protocol, which is useful as a benchmark for quantum-network performance but is not an interesting application in and of itself. It would be interesting to learn how the results presented would change if instead a complete verified blind quantum computing protocol such as the one introduced in [32] were studied.

7.5 Data availability

The data presented in this work have been made available at <https://doi.org/10.4121/22193539> [42].

7.6 Code availability

The code that was used to perform the simulations and generate the plots in this paper has been made available at <https://gitlab.com/softwarequtech/simulation-code-for-requirements-for-upgrading-trusted-nodes-to-a-repeater-chain-over-900-km-of-optical-fiber> [41].

7.7 Baseline parameters

Here we discuss how we determined the baseline hardware parameters shown in Table 7.1. We did so by following two steps: (i) finding state-of-the-art color-center hardware parameters in the literature and (ii) converting these to the hardware model we employ. In Table 7.3 we show the relevant state-of-the-art color center parameters we have identified and provide their respective references. We now discuss how these are converted to the

Parameter	State-of-the-art value
Number of modes	1
Carbon coherence time	1 s [38]
Elementary-link fidelity	0.83 [7]
Electron initialization fidelity	0.995 [5]
Carbon initialization fidelity	0.99 [38]
Electron-carbon two-qubit gate fidelity	0.97 [13]
Electron single-qubit gate fidelity	0.995 [5]
Carbon single-qubit gate fidelity	0.999 [39]
Electron readout fidelity	0.93(0) 0.995(1) [37]
Photonic interface efficiency	0.855 [15]
Frequency conversion efficiency	0.3 [45]

Table 7.3: State-of-the-art color center parameters. We have number of modes as 1 without reference because to the best of our knowledge multiplexed entanglement generation has not been demonstrated using color centers.

parameters shown in Table 7.1. The elementary-link fidelity and number of modes can be used directly without conversion. Color-center memories have both an electron qubit (also known as communication qubits due to their optical interface) and possibly multi-

ple carbon qubits (also known as memory qubits due to being long-lived). We assume a ‘best-of-both-worlds’ situation, in which the qubits in our model are both endowed with an optical interface that allows them to generate entanglement and a long (1s baseline) memory lifetime. This simplification allows us to treat all qubits in the nodes equally. As explained in Section 7.1.1 we combine all photon-related inefficiencies, with the exception of fiber attenuation, into one parameter, p_{det} . This is done as follows:

$$p_{\text{det}} = p_{\text{photon interface}} \cdot p_{\text{conv}}, \quad (7.10)$$

where $p_{\text{photon interface}}$ is the photonic interface efficiency and p_{conv} is the frequency-conversion efficiency. This results in the 0.255 number reported in Table 7.1. We note that the experiment reported in [15] does not consist of entanglement generation through a heralding station, as we assume in this paper. We have made a best guess of how the parameters reported there would translate to a scheme where entangled photons are interfered and measured at a heralding station. An entanglement swap in a color center (this concrete example was demonstrated using a nitrogen-vacancy center) consists of single-qubit gates on both carbon and electron, two-qubit gates and measurement and initialization of the electron (see Figure 17 in Supplementary Note 5 of [26] for an image of the circuit). We make the simplifying assumption that all errors are depolarizing. First, we convert each of the initialization and gate fidelities in Table 7.3 to depolarizing parameters (in accordance with Equation (7.5)), and then multiply the depolarizing parameters corresponding to all the operations in the circuit together to obtain the swap quality (which parametrizes a depolarizing channel as detailed in Section 7.1.1), i.e.,

$$s_q = (1 - p_{\text{carbon}})^2 \cdot (1 - p_{\text{electron-carbon}}) \cdot (1 - p_{\text{electron}})^2 \cdot (1 - p_{\text{electron init}}) \cdot (1 - p_{\text{electron meas}})^2 \cdot (1 - p_{\text{retrieve}}), \quad (7.11)$$

where p_{carbon} is the depolarizing parameter of the carbon single-qubit gate, $p_{\text{electron-carbon}}$ of the two-qubit gate, p_{electron} of the electron single-qubit gate, $p_{\text{electron init}}$ of the electron initialization, $p_{\text{electron meas}}$ of the electron measurement and p_{retrieve} of the retrieve operation (maps the carbon state to the electron, see Figure 17 (b) in Supplementary Note 5 of [26]).

7.8 Repeater placement chosen by optimization method

As described in Section 7.2.2, we determined minimal hardware requirements for performing QKD at a rate of 10 Hz over the network path depicted in Figure 7.1. In doing so, we optimized over the number of repeaters used and their placement. We then used the placement our optimization method found to perform best for determining minimal hardware requirements for other performance targets, as described in Section 7.3. In Figure 7.8 we show this placement. In Table 7.4 we show the lengths and attenuations of the elementary links defined by the repeater placement.

In order to optimize over the number of repeaters and their placement, we have first generated all the 986 possible ways repeaters can be placed in the network (such that there is still space for the required heralding stations between repeaters and end nodes). To each configuration we assigned a number r corresponding to the number of repeaters in



Figure 7.8: Map of Germany overlaid with a depiction of the fiber path connecting the German cities of Bonn and Berlin that we investigated. The white circles represent end nodes, in Bonn and Berlin, and repeater nodes elsewhere. This placement corresponds to the best found by our optimization method, in the sense that it allowed for minimization of hardware requirements for a target secret-key rate of 10 Hz.

Link	Length (km)	Attenuation (dB)
Bonn - Wuppertal	138.9	32.8
Wuppertal - Münster	133.2	31.4
Münster - Warmsen	126.2	29.6
Warmsen - Hannover	97.2	22.7
Hannover - Liebenburg	122.0	28.4
Liebenburg - Magdeburg	115.5	26.9
Magdeburg - Havel	103.9	24.3
Havel - Berlin	80.2	18.6

Table 7.4: Length and attenuation of elementary links depicted in Figure 7.8.

the network. Then, for each configuration we computed the chain asymmetry parameter defined as

$$\mathcal{A}_{\text{chain}} = \frac{1}{r} \sum_{i=1}^r \frac{|L_{\text{left},i} - L_{\text{right},i}|}{L_{\text{left},i} + L_{\text{right},i}}, \quad (7.12)$$

where $L_{\text{left},i}$ ($L_{\text{right},i}$) is the distance between repeater node i and its left- (right-) hand neighboring node. Next, we ordered all the configurations with the same value of r by their values of $\mathcal{A}_{\text{chain}}$, and label their position in this ordering as n . This number is then an identifier for how asymmetric (as quantified by the chain asymmetry parameter) a configuration is relative to the other configurations with the same number of repeaters. $n = 0$ corresponds to the most symmetric setup and $n = m_r - 1$ corresponds to the most asymmetric setup, where m_r is the number of configurations with r repeaters. All configurations are then stored in a table by their values for r and n .

Then, when we optimize over the hardware parameters, we also optimize over two additional parameters. These are r (the number of repeaters) and a , which is a number between zero and one. Given a pair (r, a) , the configuration that is used is chosen as

follows. First, the number a is mapped to a value of n using

$$n = \text{round}(a(m_r - 1)), \quad (7.13)$$

(where round denotes rounding to the closest integer) i.e., it uses $n = 0$ for $a = 0$ and $n = m_r - 1$ for $a = 1$. Second, the unique configuration defined by the values of r and n is taken from the table and used in the simulation. The reason why we optimize over a instead of over n directly is that a quantifies how asymmetric the chosen configuration is in a way that is independent of r (the range is always between 0 and 1, instead of between 0 and $m_r - 1$). This makes it easier to vary r and a independently compared to r and n .

7.9 Optimization method

In this section we provide more details regarding our optimization methodology. This methodology is based on genetic algorithms, which come in several different flavors. Our particular implementation is heavily based on the one introduced in Chapter 5 and used in Chapters 6 and 8. There are two things that we do discuss in this section. First, as the parameter set we use here is different than in other chapters, we explain in Section 7.9.1 how we define the probability of no imperfection for each of these, as required by the definition of the hardware cost function H_c in Section 7.1.3. Second, we have employed a simple local optimization performed on the best solution found by the genetic algorithm, which we explain in Section 7.9.2. Additionally, we also give the details of the machine used to perform the actual optimizations in Section 7.9.3. Finally we would like to remark that the code for our implementation, together with the tools required for integration with NetSquid simulations, is publicly accessible at [46].

7

7.9.1 No-imperfection probabilities

We show in Table 7.5 the probability of no-imperfection for all parameters considered in our hardware models.

Parameter	Probability of no-imperfection
Photon detection probability p_{det}	p_{det}
Coherence time T	$e^{-1/T}$
Swap quality s_q	s_q
Elementary link fidelity F_{el}	F_{el}
Number of multiplexing modes N	$1 - (1 - p_{\text{surv baseline}})^N$

Table 7.5: Probabilities of no-imperfection for hardware parameters we optimized over in this work. $p_{\text{surv baseline}}$ (defined in Equation (7.14)) is the probability of one photon (i.e., no multiplexing) surviving traveling an elementary link made up out of two times the average fiber segment in the fiber path we study (shown in Figure 7.1).

We start by defining the quantity $p_{\text{surv baseline}}$ that appears in this table more rigorously. It is computed as follows,

$$p_{\text{surv baseline}} = 10^{-\overline{\alpha_{\text{att}}}/10}, \quad (7.14)$$

with $\overline{\alpha_{\text{att}}}$ given by,

$$\overline{\alpha_{\text{att}}} = \frac{2}{N} \sum_{i=1}^N \alpha_{\text{att},i} L_i. \quad (7.15)$$

Here, L_i is the length of fiber segment i in the fiber path under consideration, $\alpha_{\text{att},i}$ is the attenuation coefficient of fiber segment i (i.e., the amount of attenuation per unit length), and N is the total number of fiber segments in the path. For the fiber path considered in this paper (i.e., the one depicted in Figure 7.1), $N = 17$. An elementary link between two neighboring nodes must consist of at least two fiber segments to allow for the installation of a heralding station. $\overline{\alpha_{\text{att}}}$ can then be thought of as the total amount of attenuation on an elementary link made up of two times the average fiber segment. This means $p_{\text{surv baseline}}$ is the probability of a photon surviving traveling through this average elementary link. The reason for constructing this quantity is that it provides a baseline for the photon survival probability in fiber, which can then be improved upon by increasing the number of multiplexing modes, thereby enabling us to associate a cost function.

The coherence time T represents a timescale for depolarization, with the probability of the state becoming maximally mixed over a period of time t being given by $1 - e^{-t/T}$, with the respective probability of no-imperfection then being $e^{-t/T}$. In this case, improving T by a factor of k is equivalent to multiplying it by k .

For the swap quality, s_q is the probability that the two-qubit state before the Bell-state measurement is not replaced with a maximally mixed state, and therefore s_q is the corresponding probability of no imperfection. Finally, for the elementary-link fidelity we take the fidelity itself to be the probability of no imperfection.

7

7.9.2 Local optimization

Genetic algorithms are derivative-free optimization algorithms that are particularly useful when applied to functions whose cost landscape is largely unknown but is assumed to have many local minima [47]. Through a balancing act of exploration (i.e., investigation of many different areas of parameter space) and exploitation (i.e., investigation of local optima) they often manage to avoid being trapped in local optima as gradient-based methods are wont to. Nevertheless, use of a genetic algorithm does not guarantee that one can find the global optimum. Further, one can not even be sure that one has maximally exploited the best optimum found. For this reason, we complement the exploration performed by the genetic algorithm with a deterministic local optimization method which we apply to the best parameter set found by the genetic algorithm. The algorithm used is a variation of an iterative local search algorithm [48]. It consists of iteratively making small changes on a parameter and evaluating the cost associated to the resulting parameter set. In case it has decreased, it is kept and we again make a small change on the same parameter. If the cost increases, we discard the change and move on to another parameter. This process is repeated for all parameters being optimized over. We must however emphasize that this still does not guarantee that the global optimum will be found.

More details on this method can be found in Chapter 4.2 of [49].

7.9.3 Performing the optimization

Each optimization run was executed on a thin node of the Snellius supercomputer [50]. Each of these nodes is endowed with 2 AMD Rome 7H12 CPUs (2.6 GHz), for a total of 128 cores and a total of 256 GiB of memory.

7.10 BQC test protocol

In this section we describe the BQC test protocol that is used as a performance metric in this paper. This protocol consists of repeated execution of test rounds as required by the VBQC protocol presented in [32]. In each round of the VBQC protocol, a server is tasked by a client to execute a quantum computation on qubits transmitted by the client and then send the classical result of that computation back to the client. In test rounds, the client has prepared the transmitted qubits in such a way that it knows the correct outcome of the computation.

Therefore, executing test rounds allows the client to verify whether the server is honest. However, noise in the quantum hardware can also lead to failed test rounds. The more often test rounds fail due to noise, the harder it is for the client to verify the server's honesty.

The BQC test protocol that we consider is not itself a VBQC protocol. In fact, its only purpose is to benchmark how suited a quantum network could be to perform BQC protocols (and perhaps other applications that require multiple live qubits simultaneously). The performance metric that we consider for this protocol is the success rate, defined as the average number of successful test rounds that can be executed per time unit (i.e., the product of the rate R and success probability p_s , as in Equation (7.8)). We specifically consider an entanglement-based two-qubit version of the protocol. In that case, the protocol is as follows:

1. The client chooses d and r uniformly at random from $\{0, 1\}$ and θ from $\{j\pi/4\}_{0 \leq j \leq 7}$, and then defines two quantum states, $|\text{dummy}\rangle = |i\rangle$ and $|\text{trap}\rangle = |+\theta\rangle$, where $|\pm_\phi\rangle \equiv \frac{1}{\sqrt{2}}(|0\rangle \pm e^{i\phi}|1\rangle)$. Additionally, it uniformly at random designates $|\phi_1\rangle$ to be $|\text{dummy}\rangle$ or $|\text{trap}\rangle$. $|\phi_2\rangle$ is designated to be the option that was not chosen.
2. When an entangled state shared between the client and server becomes available, the client uses quantum teleportation to transmit the state $|\phi_1\rangle$ to the server. The server stores the received state in quantum memory.
3. When a second entangled state becomes available, the client uses quantum teleportation to transmit the state $|\phi_2\rangle$ to the server.
4. The server performs a CZ gate between its two qubits.
5. The server measures the qubit that was used to receive the state $|\text{trap}\rangle$ in the basis $\{|+\theta+r\pi\rangle, |-\theta+r\pi\rangle\}$ and transmits the result back to the client.
6. The client declares the test round a success if it receives measurement result matching its expectation (i.e., if the outcome is equal to $d \oplus r$, where \oplus is addition modulo two), and a failure otherwise.

7. The client and server go back to Step 1 to start the next test round.

Alternatively, remote state preparation [51] could be used to prepare the required states at the server, which may be easier to execute on real hardware than quantum teleportation. In fact, we have proven in Chapter 6 that using remote state preparation for the VBQC protocol in [32] is equivalent to using quantum teleportation in case the client and server implement gates noiselessly. Therefore the success rate will, under these assumptions, be the same whether quantum teleportation or remote state preparation is used.

We here assume that classical communication between the client and the server happens instantaneously and that both the client and server are able to perform gates and measurements noiselessly and instantly. However we do not assume they are able to store qubits indefinitely; the first teleported state undergoes depolarizing noise as described in Equation (7.6), where the coherence time T is the same as the coherence time of the repeater nodes (i.e., it is varied by the optimizations performed in this paper). Under these assumptions, R_{rounds} is simply half the rate at which entanglement can be distributed when entanglement is being generated continuously, as one test round can be performed for every two entangled states that are produced. In order to calculate the success probability, we use the following result from Chapter 6:

$$1 - p_s = e^{-\frac{\Delta t}{T}} \left[F_{dummy}(1 - F_{trap}) + F_{trap}(1 - F_{dummy}) \right] + \frac{1}{2} \left(1 - e^{-\frac{\Delta t}{T}} \right). \quad (7.16)$$

Here, Δt is the time between the transmission of the first qubit and the second qubit. For the fidelities F_{trap} and F_{dummy} , let the density matrices for the state $|dummy\rangle$ after transmission to the server be ρ_{dummy} and ρ_{trap} for $|trap\rangle$. Then $F_{dummy} = \langle dummy | \rho_{dummy} | dummy \rangle$ and $F_{trap} = \langle trap | \rho_{trap} | trap \rangle$.

We then determine the success rate as follows. First, we simulate continuous entanglement generation between the end nodes of a repeater chain. Each time an end-to-end entangled state is generated it is removed from the simulation and stored as raw data, together with the time at which it was generated. Then, after the simulation has finished, we process the raw data to determine what the success rate would have been if the entangled states had been consumed by the BQC test protocol. To this end, we divide the data into single test rounds, each consisting of two entangled states that were generated in succession. We assign each test round a duration t , which is the amount of time between the start of the round and the end of the round (i.e., when the second state was generated), and a storage time Δt , which is the time between when the first entangled state and the second entangled state were generated. We furthermore calculate the p_s of that round using Equation (7.16), where we average over the two possible choices in the protocol for how $|\phi_1\rangle$ and $|\phi_2\rangle$ are designated (i.e., whether the first entangled state is used to transmit the dummy and the second to transmit the trap or vice versa). Then we calculate the rate as

$$R = \frac{1}{\langle t \rangle}, \quad (7.17)$$

where $\langle t \rangle$ is the average value of t over all the test rounds. Finally, we use R and the average value of p_s to calculate the success rate according to Equation (7.8). The processing code that realizes this calculation has been made publicly available at [52].

References

- [1] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead*, *Science* **362** (2018).
- [2] C. Cabrillo, J. I. Cirac, P. Garcia-Fernandez, and P. Zoller, *Creation of entangled states of distant atoms by interference*, *Physical Review A* **59**, 1025 (1999).
- [3] S. D. Barrett and P. Kok, *Efficient high-fidelity quantum computation using matter qubits and linear optics*, **71**, 060310, publisher: American Physical Society.
- [4] P. C. Humphreys, N. Kalb, J. P. Morits, R. N. Schouten, R. F. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network*, *Nature* **558**, 268 (2018).
- [5] M. Pompili, S. L. Hermans, S. Baier, H. K. Beukers, P. C. Humphreys, R. N. Schouten, R. F. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, *et al.*, *Realization of a multinode quantum network of remote solid-state qubits*, *Science* **372**, 259 (2021).
- [6] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, *et al.*, *Heralded entanglement between solid-state qubits separated by three metres*, *Nature* **497**, 86 (2013).
- [7] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. Vermeulen, R. N. Schouten, C. Abellán, *et al.*, *Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres*, *Nature* **526**, 682 (2015).
- [8] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network*, *Physical Review Letters* **124**, 110501 (2020).
- [9] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, *Entanglement of trapped-ion qubits separated by 230 meters*, *Phys. Rev. Lett.* **130**, 050803 (2023).
- [10] K. Azuma, S. E. Economou, D. Elkouss, P. Hilaire, L. Jiang, H.-K. Lo, and I. Tzitrin, *Quantum repeaters: From quantum networks to the quantum internet*, (2022), arXiv:2212.10820 .
- [11] T. E. Northup and R. Blatt, *Quantum information transfer using photons*, *Nature Photon* **8**, 356 (2014), arxiv:1708.00424 .
- [12] W. Pfaff, B. J. Hensen, H. Bernien, S. B. van Dam, M. S. Blok, T. H. Taminiau, M. J. Tiggelman, R. N. Schouten, M. Markham, D. J. Twitchen, *et al.*, *Unconditional quantum teleportation between distant solid-state quantum bits*, *Science* **345**, 532 (2014).
- [13] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356**, 928 (2017).

- [14] L. Slodička, G. Hétet, N. Röck, P. Schindler, M. Hennrich, and R. Blatt, *Atom-Atom Entanglement by Single-Photon Detection*, Phys. Rev. Lett. **110**, 083603 (2013).
- [15] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, *et al.*, *Experimental demonstration of memory-enhanced quantum communication*, Nature **580**, 60 (2020).
- [16] S. Langenfeld, S. Welte, L. Hartung, S. Daiss, P. Thomas, O. Morin, E. Distante, and G. Rempe, *Quantum Teleportation between Remote Qubit Memories with Only a Single Photon as a Resource*, Phys. Rev. Lett. **126**, 130502 (2021), 2105.04338 .
- [17] G. W. Lin, X. B. Zou, X. M. Lin, and G. C. Guo, *Heralded quantum memory for single-photon polarization qubits*, EPL **86**, 30006 (2009).
- [18] M. Ruf, N. H. Wan, H. Choi, D. Englund, and R. Hanson, *Quantum networks based on color centers in diamond*, Journal of Applied Physics **130**, 070901 (2021).
- [19] S. Hermans, M. Pompili, L. dos Santos Martins, A. Rodriguez-Pardo Montblanch, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study*, New Journal of Physics (2023).
- [20] M. Horodecki, P. Horodecki, and R. Horodecki, *General teleportation channel, singlet fraction, and quasidistillation*, Physical Review A **60**, 1888 (1999), publisher: American Physical Society.
- [21] M. Pompili, C. Delle Donne, I. te Raa, B. van der Vecht, M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk, S. L. N. Hermans, P. Pawełczak, W. Kozłowski, R. Hanson, and S. Wehner, *Experimental demonstration of entanglement delivery using a quantum network stack*, npj Quantum Inf **8**, 1 (2022).
- [22] L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, *Fault-tolerant quantum repeaters with minimal physical resources and implementations based on single-photon emitters*, Phys. Rev. A **72**, 052330 (2005).
- [23] M. F. Askarani, K. Chakraborty, and G. C. Do Amaral, *Entanglement distribution in multi-platform buffered-router-assisted frequency-multiplexed automated repeater chains*, New Journal of Physics **23**, 063078 (2021).
- [24] S. B. van Dam, P. C. Humphreys, F. Rozpędek, S. Wehner, and R. Hanson, *Multiplexed entanglement generation over quantum networks using multi-qubit nodes*, Quantum Science and Technology **2**, 034002 (2017).
- [25] N. Sinclair, E. Saglamyurek, H. Mallahzadeh, J. A. Slater, M. George, R. Ricken, M. P. Hedges, D. Oblak, C. Simon, W. Sohler, *et al.*, *Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control*, Physical review letters **113**, 053603 (2014).

- [26] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpędek, M. Skrzypczyk, *et al.*, *Netsquid, a network simulator for quantum information using discrete events*, Communications Physics **4**, 1 (2021).
- [27] Á. G. Iñesta, G. Vardoyan, L. Scavuzzo, and S. Wehner, *Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs*, (2022), arxiv:arXiv:2207.06533 .
- [28] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels*, Physical review letters **70**, 1895 (1993).
- [29] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Theoretical Computer Science Theoretical Aspects of Quantum Cryptography – Celebrating 30 Years of BB84, **560**, 7 (2014).
- [30] C. H. Bennett, G. Brassard, and N. D. Mermin, *Quantum cryptography without bell’s theorem*, Physical review letters **68**, 557 (1992).
- [31] P. W. Shor and J. Preskill, *Simple proof of security of the bb84 quantum key distribution protocol*, Physical review letters **85**, 441 (2000).
- [32] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, *Verifying bqp computations on noisy devices with minimal overhead*, PRX Quantum **2**, 040302 (2021).
- [33] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal blind quantum computation*, in *2009 50th Annual IEEE Symposium on Foundations of Computer Science (IEEE, 2009)* pp. 517–526.
- [34] A. Pascoletti and P. Serafini, *Scalarizing vector optimization problems*, Journal of Optimization Theory and Applications **42**, 499 (1984).
- [35] J. D. Schaffer, *Some experiments in machine learning using vector evaluated genetic algorithms*, Tech. Rep. (Vanderbilt Univ., Nashville, TN (USA), 1985).
- [36] C. Bradley, S. de Bone, P. Möller, S. Baier, M. Degen, S. Loenen, H. Bartling, M. Markham, D. Twitchen, R. Hanson, *et al.*, *Robust quantum-network memory based on spin qubits in isotopically engineered diamond*, npj Quantum Information **8**, 122 (2022).
- [37] S. Hermans, M. Pompili, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Qubit teleportation between non-neighbouring nodes in a quantum network*, Nature **605**, 663 (2022).
- [38] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A ten-qubit solid-state spin register with quantum memory up to one minute*, Physical Review X **9**, 031045 (2019).

- [39] T. H. Taminiou, J. Cramer, T. van der Sar, V. V. Dobrovitski, and R. Hanson, *Universal control and error correction in multi-qubit spin registers in diamond*, *Nature nanotechnology* **9**, 171 (2014).
- [40] *NetSquid-QRepChain*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-qrepchain> (2023).
- [41] G. Avis and F. Ferreira da Silva, *Simulation code for Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber*, <https://gitlab.com/softwarequtech/simulation-code-for-requirements-for-upgrading-trusted-nodes-to-a-repeater-chain-over-900-km-of-optical-fiber>.
- [42] F. Ferreira da Silva, G. Avis, and S. Wehner, *Replication data for: Requirements for upgrading trusted-nodes to a repeater chain over 900 km of optical fiber*, <https://doi.org/10.4121/22193539> (2023).
- [43] F. F. da Silva, A. Torres-Knoop, T. Coopmans, D. Maier, and S. Wehner, *Optimizing entanglement generation and distribution using genetic algorithms*, *Quantum Science and Technology* (2021).
- [44] G. Avis, R. Knegjens, A. S. Sørensen, and S. Wehner, *Asymmetric node placement in fiber-based quantum networks*, arXiv preprint arXiv:2305.09635 (2023).
- [45] S. Zaske, A. Lenhard, C. A. Keßler, J. Kettler, C. Hepp, C. Arend, R. Albrecht, W.-M. Schulz, M. Jetter, P. Michler, *et al.*, *Visible-to-telecom quantum frequency conversion of light from a single quantum emitter*, *Physical review letters* **109**, 147404 (2012).
- [46] A. Torres-Knoop, T. Coopmans, D. Maier, and F. Silva, *smart-stopos*, <https://gitlab.com/aritoka/smart-stopos> (2020).
- [47] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, 1st ed. (Addison-Wesley Longman Publishing Co., Inc., USA, 1989).
- [48] S. Luke, *Essentials of metaheuristics*, Vol. 2 (Lulu Raleigh, 2013).
- [49] A. Labay Mora, *Genetic algorithm-based optimisation of entanglement distribution to minimise hardware cost*, (2021).
- [50] *Snellius*, <https://www.surf.nl/en/dutch-national-supercomputer-snellius>.
- [51] C. H. Bennett, D. P. DiVincenzo, P. W. Shor, J. A. Smolin, B. M. Terhal, and W. K. Wootters, *Remote state preparation*, *Physical Review Letters* **87**, 077902 (2001).
- [52] *NetSquid-SimulationTools*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-simulationtools> (2022).

8

Reducing entanglement-distribution hardware requirements via joint hardware-protocol optimization

Adrià Labay Mora, Francisco Ferreira da Silva and Stephanie Wehner.

We conduct a numerical investigation of fiber-based entanglement distribution over distances of up to 1600 km using a chain of processing-node quantum repeaters. We determine minimal hardware requirements while simultaneously optimizing over protocols for entanglement generation and entanglement purification, as well as over strategies for entanglement swapping. Notably, we discover that through an adequate choice of protocols the hardware improvement cost scales linearly with the distance covered. Our results highlight the crucial role of good protocol choices, such as employing purification to meet high-fidelity targets and adopting a SWAP-ASAP policy for faster rates, in significantly reducing hardware requirements. To carry out this analysis, we employ an extensive simulation framework implemented with NetSquid, a discrete-event-based quantum-network simulator, and a genetic-algorithm-based optimization methodology to determine minimal hardware requirements.

8.1 Introduction

Part of the challenge in developing quantum repeaters is that hardware requirements are not fully known. There have been many investigations of such requirements (see Chapter 6

This chapter is based on the preprint arXiv:2309.11448.

and references therein). These requirements depend on multiple factors, such as which protocols are employed to generate and distribute entanglement, the placement of the repeaters (which will likely be constrained by existing fiber infrastructure, as argued in Chapter 7 and [1]), and how many repeaters are used. To the best of our knowledge, there has yet to be a study of hardware requirements that simultaneously investigates the protocols executed by the nodes in the repeater chain and the nodes' hardware in the presence of time-dependent noise. The aim of this chapter is to address this gap.

Different protocol choices likely result in different hardware requirements, as they put emphasis on different hardware properties. For example, entanglement-purification protocols can be used to enhance the quality of shared links but necessitate higher-quality gates, as they require that more operations are performed. Many such questions arise when one considers all the building blocks required to generate entanglement over long distances using quantum repeaters.

In this chapter, we investigate minimal hardware requirements for quantum-repeater chains spanning up to 1600 km. We consider requirements for achieving (a) a fidelity of 0.8 and a rate of 1 Hz and (b) a fidelity of 0.9 and a rate of 0.1 Hz. We remark that a secret key can be distilled from states satisfying each of these fidelity targets using the BB84 quantum key distribution protocol [2] in its entanglement-based version [3] (with two-way communication being required for states fulfilling only the lower-fidelity target [4]). This is shown in Section 8.5. We expect that by picking these two targets we will be able to probe two different parameter regimes in terms of which protocols perform best and what are the corresponding hardware requirements. We determine minimal hardware requirements while optimizing over the entanglement generation protocol used for establishing nearest-neighbor links, the purification protocol, and the global network protocol that controls the sequence of actions in the chain [5, 6]. We combine a simulation-based approach that allows us to accurately account for the effects of time-dependent noise with an optimization methodology based on Genetic Algorithms (GAs) to determine minimal hardware requirements [7, 8].

8

We find that meeting the performance targets we set at distances of over 200 km is only possible using quantum repeaters, with a spacing of roughly 100 km performing best. Further, we find that an adequate choice of protocols minimizes the required hardware improvement over the experimental state-of-the-art.

8.2 Methodology

In this section, we introduce our approach to finding minimal requirements for quantum-repeater hardware. We elaborate on how we model hardware, the repeater protocols we consider, and the optimization methodology used. A visual summary of the contents of this section can be seen in Figure 8.1. We note that similar methodologies have been employed earlier, as described in Chapters 5, 6 and 7.

8.2.1 Hardware model

Multiple different physical systems are being investigated as possible hardware platforms for quantum repeaters. These include color centers in diamond [9], trapped ions [10, 11], neutral atoms [11, 12] and quantum dots [13, 14]. Despite impressive recent develop-




Figure 8.1: Building blocks of the repeater chain and optimization method we consider. The chain consists of $2^n + 1$ ($n \in \mathbb{N}$) identical quantum devices with 2 end nodes and $2^n - 1$ repeaters which are equally spaced in a line with total distance d . Each node hardware is parameterized in terms of noise parameters as detailed in Section 8.2.1. Entangled states can be generated between neighboring nodes using single-click or double-click, both heralded entanglement generation protocols. These links can be purified using one of the two studied protocols: EPL or DEJMPS. Both consume two elementary links to probabilistically yield a higher fidelity one. Finally, the network protocols orchestrate the global sequence of actions in the chain. We distinguish between SWAP-ASAP (no entanglement purification) and BDCZ (a nesting level strategy with purification). The network is simulated using NetSquid from which we extract the fidelity and rate of the entangled states shared between the two end nodes. The genetic algorithm searches in the space of all hardware and protocol parameters for the solution with the lowest improvement over state-of-the-art parameters which still satisfy the target values.

ments [15, 16] a scalable quantum repeater has yet to be demonstrated. Modeling of quantum-repeater hardware can be useful in understanding how hardware limitations impact the repeater's performance, shedding light on, for example, which hardware parameters require the most improvement.

We consider a simplified platform-agnostic model for quantum nodes that aims to capture the most relevant noise sources common to all processing-node repeaters, i.e., repeaters that can not only store quantum information but also perform quantum gates. This renders our results relevant to all of them. We assume that nodes have a quantum memory of N_{qb} fully-connected qubits (i.e., two-qubit gates can be executed between any two qubits in the memory) which decohere with characteristic relaxation time T_1 and dephasing time T_2 . This means that if a state ρ is stored in memory, it undergoes amplitude damping:

$$\rho \rightarrow E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger, \quad (8.1)$$

where the Kraus operators E_0 and E_1 are given by

$$E_0 = |0\rangle\langle 0| + \sqrt{1 - p_{T_1}} |1\rangle\langle 1|, E_1 = \sqrt{p_{T_1}} |0\rangle\langle 1|, \quad (8.2)$$

and the probability p_{T_1} for the amplitude damping process is given by:

$$p_{T_1} = 1 - e^{-t/T_1}. \quad (8.3)$$

States stored in memory also undergo dephasing:

$$\rho \rightarrow (1 - p_{T_2})\rho + p_{T_2}Z\rho Z, \quad (8.4)$$

with the probability p_{T_2} for the dephasing process being given by

$$p_{T_2} = \frac{1}{2} (1 - e^{t/T_2} e^{t/2T_1}). \quad (8.5)$$

All qubits are assumed to have the same T_1 and T_2 .

Arbitrary single-qubit rotations can be performed on every qubit and are subject to depolarizing noise with probability p_1 . Similarly, two-qubit gates are subject to depolarizing noise within probability p_2 . We model gate noise by first applying the operation perfectly and then applying the noise channel. The d -dimensional depolarizing noise channel acts on a state ρ as follows,

$$\rho \rightarrow p\rho + (1 - p)\frac{\mathbb{I}}{d}, \quad (8.6)$$

where \mathbb{I} is the d -dimensional identity matrix. The single-qubit (two-qubit) case then corresponds to $d = 2$ ($d = 4$). All qubits can be measured in the Z basis, with a bit-flip error probability ξ_0 and ξ_1 of obtaining the wrong outcome. We assume that repeaters can only attempt to generate entanglement with one neighbor at a time, as is the case for many proposed quantum-repeater platforms (see Chapter 6 as well as [9]).

8.2.2 Protocols for end-to-end entanglement generation

We aim to generate entanglement between two distant end nodes of the type described in Section 8.2.1 which are connected by a chain of quantum repeaters realized with the same type of nodes. This task is broken down into the generation of entanglement between neighboring nodes, the connection of two of these short links into a longer one, and link purification. In this section, we elaborate on protocols for each of these tasks.

8

Nearest-neighbor entanglement generation

We consider heralded entanglement generation protocols [17]. These rely on the existence of a heralding station which we assume to be placed equidistantly between two neighboring nodes. The station consists of a protocol-dependent combination of beam splitters and photon detectors. In particular, we investigate single [18] and double-click [19] entanglement-generation protocols, whose success is heralded by the detection of one and two photons, respectively. In both protocols, entanglement generation is done in successive attempts, with the participating nodes learning if they have been successful in generating entanglement at the end of each attempt. We assume that each entanglement generation attempt takes time L/c where L is the fiber distance between the two nodes and $c \sim 2.14 \times 10^5$ km/s is approximately the speed of light in fiber. The probability p_{det} of a photon emitted by a node being successfully detected at the midpoint station is given by

$$p_{\text{det}} = p_{\text{emd}} \times 10^{-(\alpha_{\text{att}}/10)(L/2)}, \quad (8.7)$$

where $\alpha_{\text{att}} = 0.2 \text{ dB km}^{-1}$ is the fiber's attenuation coefficient and p_{emd} is the probability that an emitted photon is detected, given that it was not lost in fiber. This parameter then

combines multiple loss sources, such as the probability of emitting the photon in the right mode, the probability of collecting it into the fiber, and the probability of detecting it given that it arrived at the detector.

To first approximation, the entangled states ρ_{sc} generated with the Single-Click (SC) protocol are of the type [20, 21]

$$\rho_{sc} = (1 - \alpha)|\psi^\pm\rangle\langle\psi^\pm| + \alpha|11\rangle\langle 11|, \quad (8.8)$$

where α is a tunable parameter and $|\psi^\pm\rangle = (|01\rangle \pm |10\rangle)/\sqrt{2}$ is the ideal Bell state. The fidelity of this state to the ideal Bell state is $F = 1 - \alpha$, and the probability of successfully generating entanglement with the single-click protocol is given by $2p_{det}\alpha$. This has two important consequences: (i) in this protocol, fidelity can be traded off against success probability (and hence entanglement generation rate) and (ii) states generated with this protocol are imperfect (i.e., they have non-unit fidelity to the ideal Bell state) even under the assumption of perfect hardware.

We consider multiple noise sources that reduce the fidelity to the state in Equation (8.8). Specifically, the probability of emitting two photons instead of one p_{double} , the phase uncertainty σ_ϕ acquired while traveling through the fiber [22] and the Hong-Ou-Mandel visibility V of the interfering photons [23]. The visibility is defined as $V = 1 - C_{min}/C_{max}$ [24], where C_{min} is the probability that two photons being interfered at a 50:50 beamsplitter are detected at two different detectors when indistinguishability is optimized and C_{max} is the same probability when photons are made distinguishable (e.g., by having them arrive at different times).

To account for these effects we define the state efficiency

$$\eta_f = \frac{1 + \sqrt{V}}{2}(1 - p_{ph}), \quad (8.9)$$

where p_{ph} can be derived from the hardware parameters p_{double} and $p_\phi = [1 - \exp(-\sigma_\phi^2/2)]/2$ as

$$p_{ph} = (1 - p_\phi)p_d(1 - p_{double}) + p_\phi[p_{double}^2(1 - p_{double})^2]. \quad (8.10)$$

Hence, the fidelity of the elementary link generated using the single-click protocol is

$$f_{SC} = (1 - \alpha)\eta_f. \quad (8.11)$$

States generated with the Double-Click (DC) protocol are, under the assumption of perfect hardware, ideal Bell states. Therefore, in this protocol, there is no inherent limitation on the achievable fidelity. We do however model two noise sources, namely the Hong-Ou-Mandel visibility V and the light-matter-interface fidelity f_{lm} , i.e., we allow for the possibility of depolarizing noise in the light-matter state generated at the nodes. The state generated then looks like this:

$$\begin{aligned} \rho_{DC} = & \frac{f_{lm}}{2} [(1 \pm V)|\Phi_{01}\rangle\langle\Phi_{01}| + (1 \mp V)|\Phi_{11}\rangle\langle\Phi_{11}|] \\ & + \frac{1 - f_{lm}}{2} [|00\rangle\langle 00| + |11\rangle\langle 11|]. \end{aligned} \quad (8.12)$$

Its fidelity is $f_{DC} = f_{lm}(1 + V)/2$. The double-click protocol succeeds with probability $1/2p_{\text{det}}^2$. The scaling of the success probability with p_{det} , and consequently with the fiber length, is thus less favorable for the double-click protocol when compared with the single-click protocol.

Entanglement swapping

Entanglement swapping is a protocol based on quantum teleportation [25] that effectively generates longer entangled links by consuming shorter ones. Imperfections in the links and the gates used in the swap circuit cause the fidelity to decay exponentially with the number of swaps [5]. We assume the entanglement swap circuit is implemented through a Hadamard gate, a CNOT gate, and measurements in the computational basis. As discussed in Section 8.2.1, the gates suffer from depolarizing noise, and the measurements from bit-flip errors.

Entanglement purification

Entanglement purification protocols probabilistically generate fewer higher-fidelity entangled pairs from many lower-fidelity ones. We focus in particular on DEJMPS [26] and Extreme Photon Loss (EPL) [22, 27], which are both two-to-one protocols.

Concretely, the EPL protocol consists of applying CNOT gates between the entangled pairs (using one of the pairs as controls and the other pair as targets), measuring the target qubits, and keeping the entangled pair corresponding to the control qubits if the outcomes are both 1 in the Z basis. For any other combination of measurement outcomes, the remaining entangled pair is discarded. This protocol yields maximally entangled states when applied to states of the form ρ_{sc} (see Equation (8.8)), succeeding with probability $\frac{1}{2}(1 - \alpha)^2$. In fact, EPL has been shown to be optimal for such states, in the sense that (i) no other purification protocol achieves a higher fidelity, and (ii) no other protocol achieves the same fidelity with higher success probability [21].

The DEJMPS protocol starts with Alice and Bob applying the unitaries U_A and U_B , respectively, to each of their qubits. U_A is defined as

$$|0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle), \quad |1\rangle \rightarrow \frac{1}{\sqrt{2}}(|1\rangle - i|0\rangle), \quad (8.13)$$

and U_B is defined as

$$|0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle), \quad |1\rangle \rightarrow \frac{1}{\sqrt{2}}(|1\rangle + i|0\rangle). \quad (8.14)$$

They then apply CNOTs and perform measurements in the Z basis just as in the EPL protocol, but in this case accept if both measurement outcomes are equal (i.e., also in the 00 case). The output fidelity of the protocol for an input state ρ is

$$F = \frac{A^2 + B^2}{p_{\text{succ}}}, \quad (8.15)$$

with a probability of success

$$p_{\text{succ}} = (A + B)^2 + (C + D)^2, \quad (8.16)$$




Figure 8.2: Diagram of the BDCZ protocol on a 5-node network. The height of each node determines the nesting level where entanglement swap is performed. Qubits at the nodes are represented as circles which are filled if they store an entangled state. The network sequentially creates two pairs of entangled states between neighboring quantum nodes. The fidelity of such links is small (depicted with the color intensity, lower fidelity corresponds to grayer colors) and suffers from time dephasing while in memory (Equations (8.3) and (8.5)). Two links shared between the same pair of nodes are purified, leading to a high-fidelity state that can be swapped. Afterwards, the nesting level increases and the nodes with lower height are no longer needed in the process. This process is repeated until end-to-end entanglement has been established.

where

$$\begin{aligned} A &= \langle \Phi_+ | \rho | \Phi_+ \rangle, & B &= \langle \Phi_- | \rho | \Phi_- \rangle, \\ C &= \langle \Psi_+ | \rho | \Psi_+ \rangle, & D &= \langle \Psi_- | \rho | \Psi_- \rangle, \end{aligned}$$

and $|\Phi_{\pm}\rangle = \frac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle)$ and $|\Psi_{\pm}\rangle = (\mathbb{I} \otimes X)|\Phi_{\pm}\rangle$. DEJMPS has been shown to be optimal for Bell-diagonal states of rank up to 3 [21], for the same definition of optimality given for EPL concerning states of the form ρ_{sc} .

Repeater chain protocols

Repeater chain protocols orchestrate the subprotocols described above in order to generate end-to-end entanglement across a repeater chain. We investigate two such protocols, SWAP-ASAP [6, 28] and BDCZ [5]. The first one consists of a swap-as-soon-as-possible strategy, in which repeater nodes perform an entanglement swap whenever they hold two entangled pairs. The second one is a nested strategy that combines entanglement generation and entanglement purification. In this protocol, nodes are assigned a height that depends on their relative position in the chain and determines when they are allowed to perform an entanglement swap (see Figure 8.2). Whenever two nodes of the same height are connected by an entangled pair they can either decide they are ready to swap or generate more entangled pairs and then perform purification. Swapping results in nodes of larger height being connected, for which purification is a high-stake endeavor, as failure implies regenerating entangled pairs at the initial height.

Both the hardware modeling introduced in Section 8.2.1 and the protocols introduced in this section are simulated using the discrete-event-based quantum-network-simulator NetSquid, building on previous work presented in Chapter 5 and [6]. These simulations allow us to evaluate the performance achieved by different sets of hardware parameters.

8.2.3 Optimization algorithm

In Chapter 5, we restated the question of finding minimal hardware requirements as an optimization problem. We defined minimal hardware requirements as the hardware parameters requiring the smallest improvement over state-of-the-art parameters that fulfill given performance targets. This is quantified by mapping sets of hardware parameters to an improvement cost via a cost function that can then be minimized. Here, we expand on this work by also optimizing over protocol parameters. A common theme among the protocols discussed in Section 8.2.2 is rate-fidelity trade-offs. For example, employing a single-click instead of a double-click entanglement generation protocol will typically result in higher entanglement generation rates, at the expense of lower fidelities. It should then be expected that different protocol choices result in different minimal hardware requirements, and therefore that optimizing hardware and protocol parameters simultaneously should lead to less stringent hardware requirements.

The cost function that encodes the question we aim to answer is

$$T_C(\vec{x}; \vec{y}) = \mathcal{C}(\vec{x}) + A\mathcal{P}(\vec{y}), \quad (8.17)$$

which contains two terms. The first term,

$$\mathcal{C}(\vec{x}) = \sum_j [\log_{x_{\text{base}}^j} (x^j)]^{-1}, \quad (8.18)$$

is the cost of improving a set of hardware parameters from a given baseline \vec{x}_{base} to \vec{x} . This is the *hardware cost*. Here, $\vec{x} = (x_1, x_2, \dots)$ is a vector containing the value of the optimization parameters. The second term,

$$\mathcal{P}(\vec{y}) = \sum_k [1 + (y_{\text{target}}^k - y^k)^2] \Theta(y_{\text{target}}^k - y^k), \quad (8.19)$$

is the penalty assigned for not meeting the performance targets \vec{y}_{target} . It vanishes if all targets are met and is larger as the gap between the targets and the achieved performance grows. Finally, the hyperparameter A is chosen such that $C(\vec{x}) < AP(\vec{y})$ for any allowed parameter set \vec{x} , effectively ensuring that the performance targets are hard constraints. The goal of the optimization procedure is to find the set of parameters that minimize the hardware cost while satisfying the performance targets. This set of parameters is the minimal hardware requirements.

To find minimal hardware requirements, we must solve the optimization problem we just defined. We do so by employing a genetic-algorithm-based optimization methodology. These algorithms have advantages over deterministic methods like gradient descent when little is known about the landscape of the function to be optimized or if it is non-differentiable [29]. Moreover, local methods tend to converge to the local minimum closest

to the starting point, thus they often fail to find the global optimum in problems with multiple minima. GAs can avoid this through a strategy combining exploration and exploitation. The approach we employ here is represented pictorially in Figure 8.1. A collection of 120 random optimization parameters is generated $\{\vec{x}_a\}$ defining the initial population. Then, the cost function Equation (8.17) is evaluated for each parameter set \vec{x}_a , where the target values $\vec{y}_a = (F, R)$ are evaluated from an average over 100 realizations of the Net-squid repeater-chain simulation. Averaging was done due to the stochastic nature of the simulations. We empirically found that 100 realizations struck a good balance between accuracy and computation time. Parameter sets achieving low cost are selected for the next generation of the GA together with new sets of parameters created using crossover and mutation genetic processes. In the final round, after 500 iterations, the parameter set \vec{x}_{\min} with minimum cost function is selected as the optimal combination of hardware and protocol parameters with the lowest hardware improvement over state-of-the-art hardware parameters. This a standard way to terminate an evolutionary algorithm [30].

To the previous procedure, we add an extra step which consists of an iterative local search optimization over the best solution found by the GA to ensure exploitation of the minimum found [31]. This deterministic optimization algorithm is also a gradient-free method that with few function evaluations can, if possible, reduce the total parameter cost. We note that it only searches for possible reduced minima on a region around the optimal hardware parameters without changing the protocols. That is, we assume the GA has found the combination of protocols leading to the lowest hardware cost, but has not been able to reach it.

The complete optimization procedure – including 500 iterations of the GA with 120 individuals per generation and 100 repetitions of the network simulation per individual – takes about 30min and 3 days for two- and nine-node networks on a high-performance supercomputer [32]. The code used for the simulations is open-source and can be found in Ref. [33]. The data extracted with the best individuals for each network scenario can also be found in Ref. [34].

For more details on the optimization methodology, see Section 8.7.

8.2.4 State-of-the-art parameters

Determining minimal hardware requirements as described in Section 8.2.3 is done with respect to a baseline. We used experimentally realized parameters in color center experiments to determine the parameters defining this baseline. This choice was made as color centers have been used to demonstrate several quantum-networking primitives, such as remote entanglement generation [20], long-lived quantum memories [35], entanglement purification [22] and entanglement swapping in a three-node network [36]. The parameter values used are shown in Section 8.2.4. Details on how they were determined can be found in Section 8.8. We then optimize over all hardware parameters shown in Section 8.2.4 as well as over entanglement generation, whether SWAP-ASAP or BDCZ is used, which purification protocol is employed, and how many times purification is done. In case single-click is employed we optimize over the bright-state parameter α as well.

Parameter	Value
p_1 (Single-qubit gate error)	(4/3)0.001%
p_2 (Two-qubit gate error)	0.02%
ξ_0, ξ_1 (Readout error)	0.05%, 0.005%
p_{init} (Initialisation error)	0.02%
T_1 (Relaxation time)	1 h
T_2 (Dephasing time)	1 s
p_{emd} (Photonic efficiency excluding fiber)	0.46%
η_f (SC) (State efficiency)	91.96%
f_{elem} (DC) (Elementary link fidelity)	92%

Table 8.1: State-of-the-art color-center-based hardware parameters involved in the optimization. All gate-based errors (p_1 , p_2 , $\xi_{0/1}$ and p_{init}) are improved simultaneously with the same cost. The last two parameters η_f and f_{elem} apply only to single-click (SC) and double-click (DC) respectively as explained in Section 8.2.2. The value for the parameters has been obtained from state-of-the-art color-center experiments as explained in Ref. [6].

8.3 Results

We choose fidelity F_t and entanglement generation rate R_t as our performance metrics. Concretely, we pick two pairs of target values: (a) $F_t = 0.8$ and $R_t = 1$ Hz and (b) $F_t = 0.9$ and $R_t = 0.1$ Hz. As discussed in Section 8.2, any choice of protocols for entanglement generation and purification implies trade-offs. For example, the success probability of SC scales more favorably compared to DC, but this comes at the expense of inherently lower fidelity. Therefore, we expect that by picking different targets we will be able to probe different regimes in terms of which protocols and hardware improvements are found to be optimal.

Hence, in each simulation for a particular distance and number of repeaters, the genetic algorithm explores the full space of hardware parameters and protocols introduced in Section 8.2.1 and Section 8.2.2 respectively. We here assume that the number of quantum repeaters used is a quantity to be optimized and to which no cost is assigned. As the BDCZ protocol is only well defined for numbers of nodes $N = 2^n - 1$ due to its hierarchical swap structure, we restrict ourselves to configurations verifying this condition.

8.3.1 Optimal hardware cost

We show in Figure 8.3 the cost, as defined in Equation (8.17), of distributing entanglement satisfying targets (a) and (b) over distances ranging from 200 km to 1600 km. This corresponds to the best solutions found by the GA after the local iterative minimization of the hardware parameters.

We initially observe that the total hardware cost follows a linear scaling behavior in relation to the covered distance, as depicted by the linear fits depicted in the figure. This finding is unexpected, considering that multiple quantities associated with quantum repeater chains typically exhibit exponential scaling with respect to the distance or the number of repeaters employed. It is the case for, among others, the photon transmission probability with distance and the fidelity decay with the number of repeaters [5]. We must however note that the cost function we employed does not explicitly assign a cost to placing




Figure 8.3: Total hardware cost (see Equation (8.17)) resulting from the best combination of hardware parameters, network protocols, and number of repeaters for a specific total distance, as found by our method. (a) The target values are $F_t = 0.8$ and $R_t = 1\text{Hz}$ and (b) $F_t = 0.9$ and $R_t = 0.1\text{Hz}$. In both panels, the optimal protocol strategies are denoted by the color, shape, and filling of the markers. The color coding indicates the number of repeaters used ($N_{nodes} = N_{qr} + 2$), while the shape represents the type of purification and network protocol applied (circles for SWAP-ASAP, squares for BDCZ with EPL, up-triangles for BDCZ with one round of DEJMPS and down-triangles for BDCZ with two rounds of DEJMPS in the first nesting level). Filled and unfilled markers indicate the use of DC and SC, respectively. The black line represents a linear fit of the cost values with the fit parameters shown in the inset.

quantum repeaters. It does so implicitly, as using different numbers of repeaters incurs different hardware requirements and hence different costs. Nevertheless, this provides evidence that optimizing the choice of protocols and the number of repeaters employed allows for more favorable scaling of hardware requirements. This highlights the importance of addressing protocol and hardware optimization in tandem. We further note that (i) the slope of the linear fit for target (b) is larger than for target (a) and (ii) that the overall cost is also higher for target (b), implying that achieving high fidelities requires more improvement over the state of the art than achieving high rates.

SWAP-ASAP is used for all distances for target (a). This is likely because meeting the more demanding rate target is challenging when employing purification. Purification incurs a significant time penalty, as it (i) requires generating multiple entangled pairs and (ii) succeeds only probabilistically, with states having to be regenerated in case of failure. Avoiding purification might accelerate the end-to-end entanglement generation process but results in lower-quality states. For target (a) this is counteracted by employing the double-click protocol for all distances except 200 km, as this protocol has no inherent limitations on the fidelity of the elementary link states generated, and can therefore enable higher fidelities than the single-click protocol.

Similar trade-offs can be observed for target (b). The EPL protocol, which succeeds with a lower probability than DEJMPS, is employed at shorter distances. However, at distances of 1200 km and above (which employ seven or more repeaters), DEJMPS is preferred. This is because as more repeaters are used, more links must be purified, effectively resulting in more potential points of failure. Therefore DEJMPS is a more appropriate protocol for this scenario, given that it typically succeeds with a higher probability than EPL. A similar argument applies to the entanglement generation protocol change between 1200 km and 1600 km from double-click to single-click, respectively. In other words, as the distance to be covered increases, one must use the protocols with the highest success probabilities i.e., DEJMPS instead of EPL and single-click instead of double-click.

We also note that the same number of repeaters is used for both targets at any given distance. Furthermore, this number increases with the overall distance to cover. Consequently, the internode distance stays approximately constant at around 100 km (see Section 8.6 for further details), which seems to indicate that this is the ideal spacing between repeaters in this particular scenario. The exception is the case where the end-to-end distance is 200 km, where no repeaters are used.

8

8.3.2 Optimal hardware parameters

We now turn our attention to the improvement required for each hardware parameter. Concretely, we do so in Figure 8.4 for end-to-end distances between 200 km and 1600 km. In each panel, we show in the radial axis the hardware cost of the hardware parameters needed to achieve the corresponding solution in Figure 8.3. Starting from the top, the first hardware parameter shown is the elementary link fidelity f_{elem} or state efficiency η_f for DC and SC respectively. We recall that the SC elementary link fidelity can be recovered using Equation (8.11) with the corresponding value of the bright-state population α being shown in Section 8.3.2. Proceeding clockwise, the remaining parameters are the probability that an emitted photon is detected, given that it was not lost in fiber p_{emd} , the efficiency of two-qubit gates η_2 , and the coherence time T_2 . For simplicity, we have excluded the

(a) $F_t = 0.8$ and $R_t = 1$ Hz

Distance	Repeaters	$f_{elem} (\alpha)$	$p_{emd} (%)$	$p_2 (%)$ (Cost)	T_1 (h)	T_2 (s)	Protocols
200	0	0.8022 (0.16)	39.55	2% (1/5)	1	1	SWAP-ASAP + SC
400	3	0.9891 (-)	66.09	2% (1/5)	1	12.78	SWAP-ASAP + DC
600	7	0.9810 (-)	54.08	0.36% (5.63)	1	7.04	SWAP-ASAP + DC
800	7	0.9893 (-)	77.51	0.35% (5.751)	1	10.47	SWAP-ASAP + DC
1200	15	0.9961 (-)	69.24	0.58% (3.47)	1	8.41	SWAP-ASAP + DC
1600	15	0.9920 (-)	85.28	0.037% (57.28)	1.47	22.84	SWAP-ASAP + DC

(b) $F_t = 0.9$ and $R_t = 0.1$ Hz

Distance	Repeaters	$f_{elem} (\alpha)$	$p_{emd} (%)$	$p_2 (%)$ (Cost)	T_1 (h)	T_2 (s)	Protocols
200	0	0.9441 (0.034)	16.68	2% (1/5)	1	1	SWAP-ASAP + SC
400	3	0.7182 (0.26)	49.56	0.41% (3.22)	1	9.68	BDCZ + EPL + SC
600	7	0.6795 (0.29)	54.65	0.12% (16.86)	1	15.29	BDCZ + EPL + SC
800	7	0.9963 (-)	60.99	0.13% (15.68)	1	28.84	SWAP-ASAP + DC
1200	15	0.9893 (-)	78.58	0.12% (19.98)	1	97.17	BDCZ + DEJMPS (1) + DC
1600	15	0.7368 (0.074)	60.67	0.04% (50.50)	1	112	BDCZ + DEJMPS (2) + SC

Table 8.2: Minimal hardware requirements and corresponding protocols per total distance. The elementary link fidelity is shown instead of the state efficiency η_f for single click to allow for better comparison with double click. We also remark that all gate error p_1 , p_2 , ξ_0 , ξ_1 and p_{init} have been improved during the optimisation with the same cost shown in parenthesis. For simplicity, we have only included the final value of the two-qubit gate error. Finally, in the last column, we specify the protocols used as a combination of network protocol (SWAP-ASAP or BDCZ) plus entanglement purification (DEJMPS or EPL) if any plus entanglement generation (SC or DC). In the case of DEJMPS, we also specify the number of times a link is purified in the first nesting level with a newly generated link. We recall that entanglement purification is only applied to the first nesting level of the BDCZ protocol.




Figure 8.4: Parameter cost for the best solutions found at different total distances. The further from the center, the lines are towards a given parameter, the more improvement that parameter requires (logarithmic scale). The four parameters around the circle that we encounter in clockwise order: the elementary link fidelity f_{elem} (DC) or state efficiency η_f (SC), probability that an emitted photon is detected, given that it was not lost in fiber p_{emd} , two-qubit gate efficiency η_2 , and coherence time T_2 . The absolute value of the parameters can be found in Section 8.3.2 together with the protocols employed. Note the logarithmic scale.

relaxation time T_1 from the figures as it is generally not improved beyond its state-of-the-art value (indicating that it is effectively already good enough). All parameter values can be found in Section 8.3.2.

As already seen in Section 8.3.1, the best solutions for the lower-fidelity target (purple dashed line in Figure 8.4) do not make use of purification, as opposed to those for the higher-fidelity target (yellow dotted line in Figure 8.4). As a result, the best solutions for the lower-fidelity target have comparatively worse two-qubit gates and memory dephasing times. This is to be expected, as the use of purification implies that (i) more gates will be executed and (ii) that states will spend a long time in memory as more pairs need to be generated. In Section 8.3.2 (a), we see an extreme example of this: the baseline gate quality is sufficient to span 400 km using SWAP-ASAP for $F_t = 0.8$ whereas two-qubit gates with error probability $\sim 0.4\%$ are required for $F_t = 0.9$ while employing EPL over the same distance. On the other hand, the lower-fidelity-target solutions do require higher-quality elementary links and light-matter interfaces, to be able to (i) meet the more demanding rate target while also using DC and (ii) achieve a relatively high fidelity without purification.

Overall, the probability that an emitted photon is detected, given that it was not lost in fiber, and T_2 are the parameters requiring the most improvement. The first one is improved from a baseline value of $\sim 0.5\%$ to $> 50\%$ for all distances over 200 km, and T_2 also requires 1-2 order-of-magnitude improvements over the 1 s baseline for all distances over 200 km. We note also that the hardware requirements for 200 km are significantly different than those for other distances. This is to be expected, as for 200 km no repeaters are employed, which naturally results in significantly different hardware demands. For example, memory quality becomes unimportant given that states are not stored for a significant amount of time, and gates do not require improvement either given that no operations must be performed. In that case, only the probability that an emitted photon is detected, given that it was not lost in fiber requires significant improvement.

We must note that although we are only showing the two-qubit gate error in Figure 8.4, all single-qubit gates and measurements have been improved by the same factors. This was

done because even a perfect two-qubit gate was not sufficient to achieve $F_t = 0.8$ with > 15 repeaters due to errors in other operations in the entanglement swap.

8.4 Conclusion

We have determined hardware requirements for chains of processing-node quantum repeaters spanning up to 1600 km while optimizing over protocols for entanglement generation, purification, and swapping strategy. We have found that the overall hardware cost grows linearly with distance assuming a suitable choice of protocols is made. This is surprising as various quantum-repeater-related quantities scale exponentially, such as photon loss in fiber and noise introduced in swapping. A suitable protocol choice can then combat the exponential growth of hardware cost with distance that one would naively expect. Such choices include, for example, employing purification to enable achieving higher-fidelity targets and a SWAP-ASAP strategy to attain higher rates.

The state-of-the-art parameters we considered were based on NV-center experiments. However, the hardware model we used is abstract enough that the results we found are broadly applicable to any form of processing-node repeater. Furthermore, the methodology we employed is fully general and can easily be adapted for use with different performance targets, hardware models, protocols, or network topologies.

In this chapter we considered an idealized scenario in which all nodes are equally spaced with uniform fiber attenuation. However, real-world deployment of quantum networks will likely make use of existing fiber infrastructure (see Chapter 7 and [1]), for which this does not hold. In fact, we saw in Chapters 6 and 7 that taking the constraints imposed by real-world fiber networks into account significantly affects hardware requirements. A natural extension would then be to apply our methods to such a network to investigate whether the linear scaling of hardware requirements we observed still holds. Furthermore, one could also consider more intricate protocol choices. In real-world fiber networks, some links suffer from more attenuation than others. It might then be that it is wise to adopt a different purification strategy per link, as for example, it is more costly to regenerate a link in which attenuation is very high. Other examples of possible protocol choices that we have not explored are the use of cut-off timers and different swapping strategies. Allowing the genetic algorithm to also optimize over such protocols could enable further reductions in hardware requirements, although the associated growth of the parameter space could pose challenges.

8.5 Fidelity requirements for quantum key distribution

In this appendix, we show that states that meet the fidelity targets we set are also good enough to generate secret-key through the BB84 protocol. To do so, we will establish a connection between the maximum quantum-bit error rate (QBER) that can be tolerated in this protocol and the fidelity of the states. We assume that the states generated are depolarized Bell states, i.e.,

$$\rho = W|\phi^+\rangle\langle\phi^+| + \frac{1-W}{4}\mathbb{I}, \quad (8.20)$$

where W is related to the fidelity F to the ideal Bell state $|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ as $F = (1 + 3W)/4$ and \mathbb{I} is the four-dimensional identity matrix. This is a simplification, as the states we generate are not of this form. However, any entangled state of fidelity F to the ideal Bell state $|\phi^+\rangle\langle\phi^+|$ can be brought to a state of this form with the same fidelity by application of random unitaries, a process known as twirling [37]. While one would never want to do this in reality, it is useful from a theory standpoint, as we can be sure that if a depolarized Bell state of fidelity F is good enough to distill secret key, so is any other state of the same fidelity.

The QBER for a depolarized Bell state of parameter W is given by $(1 - W)/2$. This is due to the maximally-mixed component of the state, which has a weight of $1 - W$. The probability of getting different outcomes when measuring two qubits in a maximally mixed state is 0.5, regardless of the measurement basis, resulting in a QBER of 0.5 for this state. Using the relation between the parameter W and the QBER, and between W and the fidelity F , we can derive the following relation between F and the QBER, which we from here on out denote as Q :

$$F = 1 - \frac{3Q}{2}. \quad (8.21)$$

The secret-key rate SKR of the BB84 protocol is computed as (assuming that Q is identical in both measurement bases, as is the case for depolarized Bell states) [3]:

$$\text{SKR} = R \cdot \max\{0, 1 - 2H(Q)\}, \quad (8.22)$$

where R is the entanglement generation rate and $H(p) = -p \log(p) - (1 - p) \log(1 - p)$ is the binary entropy function.

Using this expression, we find that the maximum QBER that still allows for a non-zero SKR is ~ 0.11 , which for a depolarized Bell state corresponds to a fidelity of ~ 0.84 . We further note that through two-way communication, the QBER threshold of the BB84 protocol can be raised to 0.2 [4], which corresponds to a depolarized Bell state fidelity of 0.7.

Therefore, we conclude that states satisfying both of the fidelity targets we considered could also be used to distill secret key using the BB84 protocol.

8.6 Waiting Time and maximum internode distance

Here we compute the expected time required to distribute entanglement under different assumptions. This is not strictly necessary as one can perform simulations to estimate the time required for entanglement distribution. However, using such analytical results allows for saving computational resources. For example, one might determine analytically that a target rate cannot be attained by a given repeater chain under the assumption that the only source of loss is attenuation in fiber. In this case, performing the simulation is not necessary, as introducing more imperfections can only negatively affect the rate of entanglement generation.

Without entanglement purification The waiting time T_{gen} is dominated by entanglement generation. This can be modeled with a geometric distribution, which determines

the probability that the t -th attempt succeeds after $t - 1$ failed runs, [38]

$$P[T_{gen} = t] = p_{gen}(1 - p_{gen})^{t-1} \quad (8.23)$$

where p_{gen} is the success probability of the entanglement generation protocol. Then, the average waiting time to create a link is

$$E[T_{gen}] = \frac{T_{cycle}^* + 2T_{com}}{p_{gen}} \quad (8.24)$$

where $T_{com} = (L/2)/c$ is the time needed for a photon to reach the detector placed equidistantly between two nodes separated by a distance L . For long distances, $T_{com} \gg T_{cycle}^*$ leads to an entanglement generation rate that decreases proportionally to $\exp(-L)/L$.

With entanglement purification To decrease the time taken by the purification process, we adopt a strategy that converts spatial resources to temporal resources. Specifically, at level l , instead of waiting for the generation of all M links, we perform sequential purification as soon as two links are available. As a result, a link is purified M times with newly created pairs that have suffered from less decoherence. Using this strategy, the expected waiting time T^l at the l -th level needed to purify a pair $d = M - 1$ times can be calculated using the iteration formula [39]

$$T_{k+1}^l = \frac{T_k^l + T_0^l}{p_{suc}(\rho_k, \rho_0)}, \quad T_0^l = T^{l-1}, \quad (8.25)$$

where $p_{suc}(\rho_k, \rho_0)$ is the probability of successful purification between a state ρ_k that has been purified k times and the new pair ρ_0 . Here $1/p_{suc}$ takes into account the number of repetitions until a single step succeeds to which we have to add the necessary time to create all previous pairs, T_k^l , as well as the new pair T_0^l . Concretely, for $l = 0$, T_0^l corresponds to the entanglement generation waiting time in Equation (8.24).

Solving the recursion, one finds that the average waiting time to purify a pair d times is

$$E[T_d^l] = \sum_{j=1}^d \left[\prod_{k=j}^d T_0^l \frac{1 + \delta_{k1}}{p_{suc}(\rho_k, \rho_0)} \right] \approx \sum_{j=1}^d \left[\frac{2T_0^l}{p_{suc}(\rho_0, \rho_0)} \right]^{d-j} \quad (8.26)$$

where the approximation is valid when the success probabilities at the different k steps are similar.

Maximum internode distance From the expected waiting time, it is possible to calculate the maximum distance for which the creation of an entangled state fulfilling the target rate and fidelity is possible. We remark that it is typically the rate target that cannot be met, due to the exponential decrease in the elementary link success probability, as well as the growth of the communication time with distance. It is possible to give an upper bound on this distance D_{total} by considering the ideal case scenario of perfect quantum memory and photon detectors. Knowing the waiting time to create a single link Equation (8.24), the total waiting time to generate the end-to-end link with repeaters is at least $2E[T_{gen}]$

because they can only perform one action at a time, as explained in Section 8.2.1. Then, by finding the root of the function

$$R(L_{node}) = \frac{1}{2E[T_{gen}(L_{node})]} - R_t = \frac{1}{2} \frac{\eta_{fiber}(L_{node}/2)}{T_{cycle}^* + L_{node}/c} - R_t, \quad (8.27)$$

where $L_{node} = D_{total}/(N_{node} - 1)$ and we assumed $\alpha = 0.5$, we can give an upper bound on the maximum distance.

The above does not consider the target fidelity, as it depends on the strategy used, but for SWAP-ASAP we can give an upper bound using that

$$F_L = \frac{1}{4} + \frac{3}{4} \left[\frac{(1-p_1)^2(1-p_2)(3+4(\xi_0\xi_1 - \xi_0 - \xi_1))}{3} \right]^{L-1} \left(\frac{4F-1}{3} \right)^L$$

where F_L is the fidelity after L swaps and F is the elementary-link fidelity. Hence, we can find the roots of the system of equations

$$\frac{\alpha \eta_{fiber}(L_{node}/2)}{T_{cycle}^* + L_{node}/c} - R_t = 0 \quad (8.28a)$$

$$\frac{1}{4} + \frac{3}{4} \left[\frac{4(1-\alpha) - 1}{3} \right]^{N_{node}-1} - F_t = 0 \quad (8.28b)$$

in terms of α and D_{total} . This method will give a much tighter bound because the trade-off between fidelity and rate is taken into account. Even tighter bounds on the waiting time have been studied [38, 40], but these simple cases are enough for our purposes.

8

(a) $F_t = 0.8$ and $R_t = 1$ Hz			
Number of repeaters	Rate only ($\alpha = 0.5$)	SWAP-ASAP	
	Distance (km)	Distance (km)	α
0	263	231	0.2
1	479	377	0.10774
3	958	669	0.05596
7	1917	1168	0.02852
(b) $F_t = 0.9$ and $R_t = 0.1$ Hz			
0	343	285	0.1
1	638	481	0.05179
3	1277	872	0.02636
7	2554	1563	0.01330

Table 8.3: Upper bounds on the maximum end-to-end distance for which an entangled state meeting the target metrics can be generated assuming entanglement generation is done using the single-click protocol. Only photon loss is considered, assuming perfect quantum memories. The column on the left (rate only) corresponds to considering the rate bound only, whereas the one on the right takes the fidelity in the SWAP-ASAP case into consideration as well. The solutions correspond to the roots of Equation (8.27) and Equation (8.28).

The roots for the two methods are found using numerical methods, concretely the hybrid algorithm implemented in the `scipy` python library. These are shown in Section 8.6 for the two pairs of target values studied. Essentially, the first method gives the maximum possible distance while the second gives the maximum distance considering the trade-off typical of SC, but only applicable to SWAP-ASAP, this second distance is always smaller than the former.

One can note that the important parameter in the rate-only case is not the total distance but the internode length. In fact, dividing the distances in Section 8.6(a) by $N_{qr} + 1$ gives a maximum internode distance of 240 km and 319 km for the two pairs of targets, respectively.

8.7 Optimization Algorithm

In this appendix, we go into more detail regarding the optimization method we employed and the choices we made.

Unfortunately, we cannot explore the whole space of parameters for computational reasons. Therefore, we optimized only over the parameters that have a larger impact on the performance. All the parameters over which we optimized can be seen in Section 8.7. We also wanted to compare hardware requirements with different numbers of repeaters for a given distance. Thus, we performed different optimization runs for each distance, number of repeaters, and entanglement generation protocol.

The procedure for finding the optimal set of protocol and hardware parameters works as follows. The GA is initiated with a population of 120 individuals where each one of them consists of a vector \vec{x} containing random values of the parameters in Section 8.7. Then, we use NetSquid to simulate the generation of an end-to-end link for each individual. From this simulation, we extract the fidelity of the final link F and the time needed to generate it (with the entanglement generation rate R being its inverse). However, since link generation is probabilistic, we do not use these values to compute the cost function. Instead, we repeat the simulation 200 times for 2 and 3 nodes and 100 times for a higher number of nodes and compute the mean fidelity \bar{F} and mean rate \bar{R} . These values were chosen for practical purposes as a balance between computational time and accuracy. Hence, the combination of \vec{x} , \bar{F} and \bar{R} is used to evaluate the cost function in Equation (8.17). We note that the order is important because first calculating the cost and then averaging over all realizations might assign a very high-cost value to an optimal solution just because in one realization the performance targets were not reached.

We now move to the GA. This part of the process is divided into three stages that can be seen in Figure 8.1. First, we select the 24 individuals $\mathcal{S} = \{\vec{x}_a\}$ with lower cost. Second, 72 new individuals are created by crossing the parameters of two individuals randomly selected from among the 24. This step splits two individuals $\vec{x}_a, \vec{x}_b \in \mathcal{S}$ in two parts $\vec{x}_{a,b}^{<k}$ and $\vec{x}_{a,b}^{\geq k}$ at a random position k ranging from one to the number of parameters minus one. Then, the new individual is created by combining the first part of a and the second part of b , i.e. $\vec{x}' = (\vec{x}_a^{<k}, \vec{x}_b^{\geq k})$. Finally, the last 24 individuals missing to recover a population of 120 genes are created by choosing individuals from \mathcal{S} and mutating a random parameter over a region close to the previous value. The process is then repeated 500 times, which was found numerically to allow the convergence of all the situations studied.

(a) Single click		
Parameter	Baseline	Range
α	–	$[\epsilon, 0.5]$
η_f	0.9196	$[0.9196, 1)$
p_{emd}	0.0046	$[0.0046, 1)$
k_{gates}	1	$[1, 10^4]$
T_1	1 h	$[1 \text{ h}, 1 \times 10^3 \text{ h}]$
T_2	1 s	$[1 \text{ s}, 1 \times 10^5 \text{ s}]$
Strategy	–	SWAP-ASAP, EPL, DEJMPS ($n = 1, 2, 3$ iterations)
(b) Double click		
Parameter	Baseline	Range
f_{elem}	0.92	$[0.92, 1)$
p_{emd}	0.0046	$[0.0046, 1)$
k_{gates}	1	$[1, 10^4]$
T_1	1 h	$[1 \text{ h}, 1 \times 10^3 \text{ h}]$
T_2	1 s	$[1 \text{ s}, 1 \times 10^5 \text{ s}]$
Strategy	–	SWAP-ASAP, EPL, DEJMPS ($n = 1, 2, 3$ iterations)

Table 8.4: Parameters over which we optimized for (a) single- and (b) double-click entanglement generation protocols. The baseline value and the possible range of improvement are also shown for the hardware parameters. In terms of the protocol parameters, no baseline value is used and the range corresponds to the possible values it can take.

8

After the final iteration, the individual \vec{x}_{\min} with the lowest cost (8.17). However, there is no assurance of finding the global optimum, and there may be room for exploiting the minimum found. Therefore, we added an additional optimization step to attempt to further reduce the cost. Concretely, we use a Hill climbing algorithm [31] that searches for minimum solutions in a region around \vec{x}_{\min} . This algorithm is a gradient-free method and it involves iteratively introducing minor modifications to a hardware parameter and assessing the resulting cost. If the cost decreases, the modified parameter is retained, and the process is repeated by making further adjustments to the same parameter. However, if the cost increases, the modification is discarded, and the algorithm proceeds to explore other parameters. The extra minimization procedure allows us to exploit the minima found by the GA assuming the optimal combination of protocols has already been found. Hence, we only try to minimize the hardware parameters. The output of the local search algorithm is what we call in the paper the optimal solution. We must note, however, that it is not possible to guarantee that the global optimum is found.

The simulations were executed on the High-Performance Computing facility in the Netherlands. The super-computer used is the Cartesius system which consists of nodes between 16 and 64 CPUs [32] in which we can parallelize the fitness evaluation of the in-

dividuals within a generation. The most common node contains 2×12 -core 2.4 GHz Intel Xeon E5-2695 v2 (Ivy Bridge) CPUs/node with 64 GB/node. For this reason, the population size will be a multiple of 12 (number of cores) to take advantage of the parallelization of the cost evaluation, concretely 120 individuals. The computational time increases exponentially with the number of nodes in the repeater chain. A single execution of the simulation for 2 nodes takes 0.5 ms, increasing to 100 ms for 9 nodes. Implying a total running time from 30 min for 2 nodes to 3 days for 9 nodes.

8.8 Validation

In this work, we used an abstract model for repeater nodes with the goal of approximating the behavior of all types of processing nodes. This makes results more broadly applicable and easier to interpret but comes at the cost of accuracy. To get a sense for how well our abstract model performs, we compare results obtained using it with those obtained running a hardware-specific model simulating an NV-center repeater chain in [6]. There, the fidelity and entanglement generation rate was measured for a linear repeater chain with 0 (direct connection) and 3 repeaters. Entanglement generation was done using SC and the network protocol used was SWAP-ASAP. A star topology was considered, with the center qubit being optically active and used as a communication qubit. All other qubits were used exclusively as memory qubits. The different types of qubits have different coherence times [6]. Furthermore, induced dephasing noise was considered. This was modeled by a dephasing channel which is applied to memory qubits whenever the communication qubit is used to attempt entanglement generation.

We, on the other hand, assume all qubits to be identical. The properties of the qubits, namely coherence times and gate errors, were assumed to be given by the worst between memory and communication qubits in the NV case.

We assume the states generated have fidelity $(1 - \alpha)\eta_f$ to ρ_{sc} (see Eq. 8.8), where we have condensed all the parameters that reduce the fidelity into $\eta_f = \frac{1 + \sqrt{V}}{2}(1 - p_{ph})$. This accounts for visibility, the dephasing introduced due to double photonic excitation and phase uncertainty. For comparison, we consider here also a simpler approximation, proposed in [20] which assumes perfect state efficiency, $\eta_f = 1$. In both cases, we disregard dark counts and the bright-state population is set to $\alpha = 0.1$. The three models are compared for two parameter sets. The first consists of near-term hardware values, and the second of an improved set of parameters. The exact parameters used are shown in Section 8.8.

The fidelity and rate obtained with near-term and improved parameters for our abstract model and the hardware-specific NV model of [6] can be seen in Figure 8.5. We remark that the two abstract models only differ in the value of η_f , which only affects the elementary link fidelity, so no difference in rate is expected.

Starting with the no-repeater scenario, we see good agreement in the rate when $\eta_{trans} > p_{dc}$. After this point, the rate in the abstract model continues to decrease exponentially, whereas the NV model stabilizes due to the presence of dark counts. The most important difference occurs in the fidelity, even in this regime. The NV model always gives a lower value, as expected, but for near-term hardware, the abstract model with perfect state efficiency deviates by more than 10%. The more realistic model ($\eta_f < 1$) does give a better

Parameter	NV equivalent	Near-term	Improved x10
$p_{1,gate}$	Carbon single qubit gate error	(4/3)0.001	(4/3)0.0001
$p_{2,gate}$	Electron-Carbon (EC) controlled R_X gate error	0.02	0.002
ξ_0, ξ_1	Electron readout error	0.05, 0.005	0.005, 0.0005
p_{init}	Electron initialisation error	0.02	0.002
T_1	Electron relaxation time	1 h	10 h
T_2	Carbon dephasing time	1 s	10 s
T_{cycle}^*	Photon emission delay		$3.8\mu\text{s}$
$t_{1,gate}$	Carbon single qubit gate duration		$20\mu\text{s}$
$t_{2,gate}$	EC controlled R_X gate duration		$500\mu\text{s}$
t_{init}	Carbon initialisation duration		$310\mu\text{s}$
t_{meas}	Electron read-out duration		$3.7\mu\text{s}$
N_{qubit}	Number of qubits per node		4
γ	Transmission loss		0.2 dB/km
c	Speed of light in fiber		2.14×10^5 km/s
p_{emd}	Photonic efficiency excluding fiber	0.0046	0.58
V	Photon visibility	0.9	0.99
p_d	Probability of double excitation	0.06	0.003
p_ϕ	Interferometric phase uncertainty	0.35 rad	0.11 rad

Table 8.5: Parameters used in the validation plots shown in Figure 8.5, same as those in [6]. In case there were differences in parameters between the memory and communication qubits, the most pessimistic value was chosen.

8

approximation when $\eta_{trans} \ll p_{dc}$, but we can see that the fidelity decreases significantly after $\eta_{trans} \approx p_{dc}$, showing that dark counts have a larger effect on the fidelity than on the rate. Despite that, for distances where the achieved rate is higher than 0.1 Hz, the difference between the expected fidelity is $< 3\%$ with imperfect state efficiency. We note that this is the regime we investigate in this work. With improved hardware, both abstract models give accurate values for the target metrics in the regime where dark counts can be neglected.

In the three-repeaters case, we see that the agreement in rate is good for all distances considered because the internode distance is never large enough for dark counts to become relevant. On the contrary, similar to the previous scenario, the fidelity achieved is slightly higher than in the NV model, with the abstract model with $\eta_f < 1$ being the one that reaches a closer value. Nevertheless, the difference is much smaller than in the previous case due to the higher amount of operations and storage time needed, which take a much more important role than the errors introduced during the creation of elementary links. However, there is one region with improved hardware where the fidelity in both abstract models falls below the NV one. The same region shows the largest deviation from the expected rate. This can be due to the choice of parameters made, i.e., the fact that we considered the most pessimistic properties of electron and carbon qubits, but also the un-




Figure 8.5: Comparison of NV-center model in [6] (cyan) with the abstract model using two approximations for the elementary link fidelity: in orange, with imperfect state efficiency; and in green, with perfect state efficiency. The distance is defined as the total distance between the end nodes. The vertical black lines denote the distance at which the transmission efficiency is comparable to the dark count probability. For 5 nodes, this occurs at a distance larger than the ones considered.

restricted topology. At such short distances, the time spent mapping electrons to carbon states becomes important, resulting in a lower rate. This is neglected in the abstract models, which means that fewer operations are performed there. However, as we chose the worst parameters between memory and communication qubits when mapping from NV to abstract, more noise will be introduced in each operation. Part of the disagreement in the fidelity can also be due to neglecting the induced dephasing in the carbon qubits. Nevertheless, the agreement between the NV and the most elaborate abstract model is below 3% when the rate is above 0.1 Hz.

All in all, it is possible to conclude that the abstract model does give an accurate description of the rate in the regime $\eta_{trans} > p_{dc}$, showing that it is possible to disregard any restriction on the topology. The fidelity is better approximated with $\eta_f < 1$, although the difference between the two abstract models is reduced if three repeaters are used.

References

- [1] J. Rabbie, K. Chakraborty, G. Avis, and S. Wehner, *Designing quantum networks using preexisting infrastructure*, npj Quantum Information **8**, 5 (2022).

- [2] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Theoretical Computer Science **560**, 7 (2014).
- [3] C. H. Bennett, G. Brassard, and N. D. Mermin, *Quantum cryptography without bell's theorem*, Physical review letters **68**, 557 (1992).
- [4] D. Gottesman and H.-K. Lo, *Proof of security of quantum key distribution with two-way classical communications*, IEEE Transactions on Information Theory **49**, 457 (2003).
- [5] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum repeaters: the role of imperfect local operations in quantum communication*, Physical Review Letters **81**, 5932 (1998).
- [6] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpedek, M. Skrzypczyk, *et al.*, *Netsquid, a network simulator for quantum information using discrete events*, Communications Physics **4**, 1 (2021).
- [7] F. Ferreira da Silva, A. Torres-Knoop, T. Coopmans, D. Maier, and S. Wehner, *Optimizing entanglement generation and distribution using genetic algorithms*, Quantum Science and Technology (2021).
- [8] M. Chehimi, S. Pouryousef, N. K. Panigrahy, D. Towsley, and W. Saad, *Scaling limits of quantum repeater networks*, (2023), arXiv:2305.08696 [cs.NI].
- [9] M. Ruf, N. H. Wan, H. Choi, D. Englund, and R. Hanson, *Quantum networks based on color centers in diamond*, Journal of Applied Physics **130**, 070901 (2021).
- [10] L.-M. Duan and C. Monroe, *Colloquium: Quantum networks with trapped ions*, Reviews of Modern Physics **82**, 1209 (2010).
- [11] A. Reiserer and G. Rempe, *Cavity-based quantum networks with single atoms and optical photons*, Reviews of Modern Physics **87**, 1379 (2015).
- [12] S. Langenfeld, P. Thomas, O. Morin, and G. Rempe, *Quantum repeater node demonstrating unconditionally secure key distribution*, Physical review letters **126**, 230506 (2021).
- [13] W. Gao, P. Fallahi, E. Togan, J. Miguel-Sánchez, and A. Imamoglu, *Observation of entanglement between a quantum dot spin and a single photon*, Nature **491**, 426 (2012).
- [14] T. Li, G.-J. Yang, and F.-G. Deng, *Heralded quantum repeater for a quantum communication network based on quantum dots embedded in optical microcavities*, Physical Review A **93**, 012302 (2016).
- [15] M. Pompili, S. L. Hermans, S. Baier, H. K. Beukers, P. C. Humphreys, R. N. Schouten, R. F. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, *et al.*, *Realization of a multinode quantum network of remote solid-state qubits*, Science **372**, 259 (2021).
- [16] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, *et al.*, *Experimental demonstration of memory-enhanced quantum communication*, Nature **580**, 60 (2020).

- [17] K. Azuma, S. E. Economou, D. Elkouss, P. Hilaire, L. Jiang, H.-K. Lo, and I. Tzitrin, *Quantum repeaters: From quantum networks to the quantum internet*, arXiv preprint arXiv:2212.10820 (2022).
- [18] C. Cabrillo, J. I. Cirac, P. Garcia-Fernandez, and P. Zoller, *Creation of entangled states of distant atoms by interference*, *Physical Review A* **59**, 1025 (1999).
- [19] S. D. Barrett and P. Kok, *Efficient high-fidelity quantum computation using matter qubits and linear optics*, *Physical Review A* **71**, 060310 (2005).
- [20] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network*, *Nature* **558**, 268 (2018).
- [21] F. Rozpedek, T. Schiet, L. P. Thinh, D. Elkouss, A. C. Doherty, and S. Wehner, *Optimizing practical entanglement distillation*, *Physical Review A* **97**, 062333 (2018).
- [22] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356**, 928 (2017).
- [23] C.-K. Hong, Z.-Y. Ou, and L. Mandel, *Measurement of subpicosecond time intervals between two photons by interference*, *Physical review letters* **59**, 2044 (1987).
- [24] F. Bouchard, A. Sit, Y. Zhang, R. Fickler, F. M. Miatto, Y. Yao, F. Sciarrino, and E. Karimi, *Two-photon interference: the hong-ou-mandel effect*, *Reports on Progress in Physics* **84**, 012402 (2020).
- [25] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels*, *Phys. Rev. Lett.* **70**, 1895 (1993).
- [26] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, *Quantum privacy amplification and the security of quantum cryptography over noisy channels*, *Phys. Rev. Lett.* **77**, 2818 (1996).
- [27] S. B. v. Dam, P. C. Humphreys, F. Rozpedek, S. Wehner, and R. Hanson, *Multiplexed entanglement generation over quantum networks using multi-qubit nodes*, *Quantum Science and Technology* **2**, 034002 (2017).
- [28] Á. G. Iñesta, G. Vardoyan, L. Scavuzzo, and S. Wehner, *Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs*, arXiv preprint arXiv:2207.06533 (2022).
- [29] D. Goldberg and D. Edward, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Artificial Intelligence (Addison-Wesley Publishing Company, 1989).
- [30] B. J. Jain, H. Pohlheim, and J. Wegener, *On termination criteria of evolutionary algorithms*, in *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation* (2001) pp. 768–768.

- [31] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis, *How easy is local search?* *Journal of computer and system sciences* **37**, 79 (1988).
- [32] SURFsara, *Cartesius: the dutch supercomputer*, (2021), accessed April 11, 2021.
- [33] A. Labay Mora, *Simulation code for reducing hardware requirements for entanglement distribution via joint hardware-protocol optimization* · *GitLab*, <https://gitlab.com/1abay11/NetSquid-SimplifiedRepChain>.
- [34] A. Labay Mora, F. Ferreira da Silva, and S. Wehner, *Replication data for: Reducing hardware requirements for entanglement distribution via joint hardware-protocol optimization*, <https://doi.org/10.4121/0c6f100c-cf16-4fe1-81fe-afcafabc7ca.v1> (2023).
- [35] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A ten-qubit solid-state spin register with quantum memory up to one minute*, *Physical Review X* **9**, 031045 (2019).
- [36] S. Hermans, M. Pompili, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Qubit teleportation between non-neighbouring nodes in a quantum network*, *Nature* **605**, 663 (2022).
- [37] M. Horodecki, P. Horodecki, and R. Horodecki, *General teleportation channel, singlet fraction, and quasidistillation*, *Physical Review A* **60**, 1888 (1999).
- [38] S. Brand, T. Coopmans, and D. Elkouss, *Efficient computation of the waiting time and fidelity in quantum repeater chains*, *IEEE Journal on Selected Areas in Communications* **38**, 619 (2020).
- [39] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller, *Quantum repeaters based on entanglement purification*, *Physical Review A* **59**, 169 (1999).
- [40] T. Coopmans, S. Brand, and D. Elkouss, *Improved analytical bounds on delivery times of long-distance entanglement*, *Physical Review A* **105**, 012608 (2022).

9

Protocols for generation of high-quality entanglement using reinforcement learning

Francisco Ferreira da Silva and Stephanie Wehner.

This chapter introduces a reinforcement learning approach to the problem of finding protocols for generating high-quality entanglement. By conceptualizing a quantum network as an environment and simulating it using NetSquid, a discrete-event-based simulator for quantum networks, we develop a framework for finding optimal protocols for the generation of high-quality entangled states between two remote parties. We introduce two applications of our approach: purification, in which the goal is to find the sequence of actions resulting in the best purification protocol and what we call purify, entangle and discard, in which the goal is to find the sequence of these actions that allows for the quickest path to high-quality entanglement between two nodes with imperfect memories. The tools we introduce can be of use in the design of quantum communication protocols tailored to imperfect quantum hardware.

9

9.1 Introduction

In previous chapters of this thesis we have mostly concerned ourselves with the nitty-gritty details of how quantum repeaters can be used to generate high-quality entanglement. Here we will take a higher-level view and consider instead two remote end nodes that wish to share high-quality entanglement, abstracting away the physical details of how they achieve this.

High-quality entangled states are an indispensable resource for many quantum-network applications [1]. Therefore, the generation of such states over long distances is a prerequisite for large-scale quantum networks. It is also a formidable challenge (see Chapter 2 and references therein). Many protocols have been proposed to tackle this challenge. These

include repeater-chain protocols, which dictate how repeater nodes should behave so as to distributed entanglement over distances longer than those that would be feasible through direct transmission (see Chapter 6 and references therein). Another example is that of purification protocols, which allow for the probabilistic transformation of many lower-quality entangled states into fewer higher-quality ones (see Chapter 3). Predicting how well such protocols perform is tractable in simplified scenarios, in which one for example considers that all nodes are identical and equally spaced, and ignores the effects of time-dependent noise. In such settings, it is even sometimes possible to make statements regarding the optimality of given protocols [2, 3]. However, it is very likely the case that real-world scenarios do not neatly fit into such simplifying assumptions. Given that employing hardware-tailored protocols can significantly boost achievable performance (see Chapter 8 for examples), numerical tools can play an important role in evaluating and designing entanglement generation and distribution protocols for near-term quantum hardware to be deployed in the real world. Quantum-network simulators such as NetSquid [4] can be employed to evaluate the performance of given quantum-network protocols, as exemplified in Chapters 5 to 8 of this thesis. It has also been shown that reinforcement learning (RL) can be a useful tool for the exploration and discovery of novel quantum-network protocols [3, 5].

9.2 Preliminaries

In this chapter we will cast the problem of designing quantum-networking protocols as a RL task. This approach is a direct extension of the one introduced in [5]. A RL task can, in simple terms, be seen as an interaction between an agent and an environment. At each step of this interaction the agent performs an action which impacts the environment's state. The agent then receives information about how the environment's state was changed, and possibly a reward. The goal is for the agent to find the state-action mapping, commonly known as a policy, which maximizes the reward obtained. In our particular case, the environment will be a two-node quantum network (although this framework is not restricted to two-node networks), whose state is simulated using NetSquid. The agent can perform actions that affect the environment's state, such as performing a quantum gate on a qubit stored on one of the node's memories or attempting to generate entanglement. A depiction of this process is shown in Figure 9.1. For a more detailed introduction to the topic see, for example, [6]. The environment we use is task-dependent. Nevertheless, the goal we concern ourselves with is always the generation of high-quality entanglement between remote nodes. Some basic concepts and operations must then be introduced. Nodes might wish to attempt entanglement generation. We assume that this process succeeds probabilistically, with the attempts being independent. This gives rise to a geometric distribution for the number of attempts required for success. We assume also that nodes are endowed with imperfect quantum memories. By this we mean that the states stored in the memories undergo decoherence. In particular, we will consider a depolarizing noise model for the memories (even though we emphasize that this framework is completely general in this regard and that different noise models can be used). In such a model, a single-qubit state ρ stored in a memory of coherence time T evolves as follows:

$$\rho \rightarrow \exp^{-t/T} \rho + (1 - \exp^{-t/T}) \frac{\mathbb{1}_2}{2}, \quad (9.1)$$




Figure 9.1: Conceptualization of quantum-network protocol discovery as a reinforcement learning (RL) task. An agent acts on an environment, which is a quantum network. The network's state changes in response, with this evolution being tracked with NetSquid. The agent receives some information about the state of the system and possibly a reward. The goal of the process is for the agent to maximize the reward obtained. The reward function should then be designed in such a way that good protocols correspond to high rewards.

where t is the time the state has spent in memory and $\mathbb{1}_2$ is the identity matrix of dimension 2, which corresponds to a maximally-mixed single-qubit state. A possible tool for combatting decoherence is the use of cut-offs [7]. The simplest form thereof is defining a local time T_c after which a state is discarded. In other words, whenever a state is created in a node's memory, a timer t starts. When $t = T_c$, the state is discarded. This is an imperfect strategy, as it for example ignores information such as whether the state was purified or with which fidelity it was created. Techniques such as fidelity tracking can result in the design of better cut-off protocols. Finding optimal cut-off times, even for a simple strategy as the one described, is in general non-trivial [8]. Finally, we consider also purification. Purification protocols consume n entangled states to probabilistically output k higher-fidelity entangled states, with $n > k$. We consider $2 \rightarrow 1$ purification protocols, i.e., protocols that consume two entangled pairs and output one. We focus in particular on DEJMPS [9] and EPL [10, 11]. Concretely, the EPL protocol consists in applying CNOT gates between the entangled pairs (using one of the pairs as controls and the other pair as targets), measuring the target qubits and keeping the entangled pair corresponding to the control qubits if the outcomes are both 1 in the Z basis. For any other combination of measurement outcomes, the remaining entangled pair is discarded. This protocol yields maximally entangled states when applied to states of the form ρ_{sc} (see equation (8.8)), succeeding with probability $\frac{1}{2}(1 - \alpha)^2$. In fact, EPL has been shown to be optimal for such states, in the sense that (i) no other purification protocol achieves a higher fidelity, and (ii) no other protocol achieves the same fidelity with higher success probability [2].

The DEJMPS protocol starts with Alice and Bob applying the unitaries U_A and U_B , respectively, to each of their qubits. U_A is defined as

$$|0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle), \quad |1\rangle \rightarrow \frac{1}{\sqrt{2}}(|1\rangle - i|0\rangle), \quad (9.2)$$

and U_B is defined as

$$|0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle), \quad |1\rangle \rightarrow \frac{1}{\sqrt{2}}(|1\rangle + i|0\rangle). \quad (9.3)$$

They then apply CNOTs and perform measurements in the Z basis just as in the EPL protocol, but in this case accept if both measurement outcomes are equal (i.e., also in the 00 case). The output fidelity of the protocol for an input state ρ is

$$F = \frac{A^2 + B^2}{p_{\text{succ}}}, \quad (9.4)$$

with probability of success

$$p_{\text{succ}} = (A+B)^2 + (C+D)^2, \quad (9.5)$$

where

$$\begin{aligned} A &= \langle \Phi_+ | \rho | \Phi_+ \rangle, \\ B &= \langle \Phi_- | \rho | \Phi_- \rangle, \\ C &= \langle \Psi_+ | \rho | \Psi_+ \rangle, \\ D &= \langle \Psi_- | \rho | \Psi_- \rangle, \end{aligned}$$

and $|\Phi^\pm\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |\pm 11\rangle)$. DEJMPS has been shown to be optimal (in the same sense as EPL is for states of the form ρ_{sc}) for Bell-diagonal states of rank up to 3 [2].

9.3 Prior work

In [5], the authors frame the problem of finding quantum-network protocols as a RL task. Let us look at a concrete example of one of the scenarios they consider, which we call the purification scenario. We consider two nodes with a perfect quantum memory of two qubits each. The nodes are also endowed with a universal gate set and it is assumed that all gates are noiseless. The starting state of the environment has the two nodes share two Werner states with non-unit fidelity F_{init} to an ideal Bell state. The goal is for the agent to find a protocol that probabilistically outputs one entangled state with fidelity $F > F_{\text{init}}$, i.e., to find a suitable $2 \rightarrow 1$ purification protocol. To this end, the agent can perform the following actions: single-qubit gates on each of the qubits (namely a Hadamard gate H and $P_x = HPH$), CNOT gates between qubits in the same quantum memory, Z -basis measurement on each of the qubits. Given the stochastic nature of quantum mechanics, the Z -measurements cause the environment to branch out, with each branch being defined by a particular measurement outcome. The agent can therefore also choose to accept or reject each branch, as is typical in purification protocols. The reward function R is defined as:

$$R = \max\left(0, \sqrt[10]{\pi_{i-1}^1 0 p_i \Delta F}\right), \quad (9.6)$$

where p_i is the success probability of the i th step (assuming the protocol found by the agent is applied recursively) and ΔF is the increase in fidelity after ten steps. Relating this to the framework we showed in Figure 9.1, an iteration of the process corresponds to the agent choosing one of the actions listed above (gate, measurement or acceptance/rejection), the state of the environment being updated and the agent learning about this. In this case, a reward is only given upon acceptance. Using this framework, the authors manage to find

the DEJMPS protocol, which is indeed the best-known $2 \rightarrow 1$ purification protocol for depolarized Bell states as measured by the reward function R .

NetSquid is a discrete-event-based quantum-network simulator [4]. It is particularly well-suited to the analysis of quantum networks with imperfect components whose performance is affected by time-dependent noise. It has been used to, for example, model repeater chains (see Chapters 5 to 8) and quantum switches [4, 12]. Modelling the quantum-networking environments that the agent will explore using NetSquid can then allow us to probe regimes and design protocols for situations that would be challenging or impossible to study analytically.

9.4 Our tool

We have developed tools for performing quantum-network RL experiments using NetSquid. We have prepared environments for two particular scenarios: a $2 \rightarrow 1$ purification environment, identical to the one described above, and what we call a purify, entangle and discard (PED) environment, in which two nodes with imperfect quantum memories aim to share entangled states of a target fidelity, which they do by (i) generate entanglement, (ii) perform purification and (iii) discard entanglement. This code is available at [13]. It directly inherits some desirable properties of NetSquid: besides accurately accounting for time-dependent noise, it is also completely modular. This means that changing, for example, the size of the quantum memories is trivial (which would easily allow for studying $n \rightarrow k$ purification protocols), as is investigating different noise models. Using NetSquid for this purpose is then a step towards the vision of designing hardware-aware protocols.

We will now describe in more detail the scenarios we have investigated using this framework by elaborating on the environments we have defined and the results we have obtained.

9.4.1 Purification

The scenario we consider here is identical to the one investigated in [5] and described in Section 9.3. Implementing this environment with NetSquid was done with two purposes in mind: (i) it allows for validation of this approach, given that there are some known optimality results and that we can compare the performance of the agent in our environment with what is described in [5] and (ii) using NetSquid makes future expansions to more complex purification scenarios (e.g., $n \rightarrow k$ rather than $2 \rightarrow 1$) simple. We consider two nodes, each endowed with a quantum memory of two qubits, which share two entangled pairs of fidelity F_{init} . The quantum memories are arbitrary in the sense that they can have any noise model (including no noise model), any coherence time and any form of noise on the gates. The configuration of the quantum memory is determined through a user-defined human-readable configuration file, as described in Chapter 6.

Let us now look at two particular questions that can be answered using this environment.

Application to Werner states

DEJMPS has been shown to be optimal for Bell-diagonal states of rank up to three [2]. We recall that the definition of optimal in this statement is (i) no other purification protocol achieves a higher fidelity, and (ii) no other protocol achieves the same fidelity with higher

success probability. Depolarized Bell states are defined as

$$\rho_{\text{db}}(p) = p|\phi\rangle\langle\phi| + (1-p)\frac{1}{4}, \quad (9.7)$$

where p is related to the state's fidelity F to the Bell state $\phi = 1/2(|00\rangle + |11\rangle)$ as $p = (4F - 1)/3$ and 1 is the 4-dimensional identity matrix. A depolarized Bell state is an example of a Bell-diagonal state of rank four. Numerical results suggest that DEJMPS might be the best $2 \rightarrow 1$ purification protocol for such states [2, 5], but there are no known optimality results. Using our framework to try to find purification protocols for such states then seems like an attractive proposition as it: (i) can serve as validation for our approach, as if it is the case that DEJMPS is the best protocol found we will be reproducing the results of [5] and (ii) finding a protocol that outperforms DEJMPS would be a novel result. With this in mind, we set the environment up as follows: two nodes share two depolarized Bell states with $F = 0.7$. Both nodes can perform CNOT gates with qubit 0 as control and qubit 1 as target (the qubit indexing is arbitrary) and Hadamard and $\sqrt{-iX}$ gates on both of their qubits. Further, the nodes can perform Z -basis measurements on both qubits. The $\sqrt{-iX}$ gate is defined as

$$\sqrt{-iX} = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & 1-i \\ 1-i & 0 \end{bmatrix}. \quad (9.8)$$

We assume that the nodes' memories are perfect (i.e., there is no decoherence) and that all gates can be executed noiselessly. The goal is then for the agent to find a sequence of actions that results in an increase in fidelity. The actions that the agent can take are as follows:

- $\sqrt{-iX}$ on each qubit (4 actions),
- Hadamard gate on each qubit (4 actions),
- CNOT gates between qubits 0 and 1 at the same node (2 actions),
- Z -basis measurements on each qubit (4 actions),
- Accept/Reject (2 actions).

This adds up to a total of 16 actions. In order to reduce the search space to be explored by the agent, we also forbid sequences of actions that would result in no entanglement being shared between the two nodes. Namely, the agent is not allowed to measure both qubits of either of the nodes. The observations available to the agent are the previous actions taken in the current trial. Different measurement outcomes lead to different branches and hence to different observations. The reward function R we employed was

$$R = \max(0, p\Delta F), \quad (9.9)$$

where p is the probability of success (which corresponds to the combined probability of the accepted branches) and ΔF is the difference between the average fidelity of the accepted branches and the input fidelity.

Running 50 agents on this environment for 500000 trials, the best protocol found corresponds to the known BBPSSW protocol [14]. It is also equivalent to DEJMPS for only one round of application, given that the initial states we consider are already in the depolarized Bell state form.

Application to R-states

The EPL protocol has been shown to be optimal for states of the form

$$\rho = p|\psi\rangle\langle\psi| + (1-p)|11\rangle\langle 11|, \quad (9.10)$$

where $|\psi\rangle = 1/2(|01\rangle + |10\rangle)$. We call these states R-states, following the terminology introduced in [2]. Applying our framework to a scenario in which the goal is to purify states of this form can then serve as validation of our approach. With this in mind, we set the environment up as follows: two nodes share two R-states with $F = 0.8$. Both nodes can perform CNOT gates with qubit 0 as control and qubit 1 as target (the qubit indexing is arbitrary but the CNOT can only be applied in one direction) and Hadamard and $\sqrt{-iX}$ gates on both of their qubits. The nodes can also perform Z-basis measurements on both qubits. The $\sqrt{-iX}$ gate is defined as

$$\sqrt{-iX} = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & 1-i \\ 1-i & 0 \end{bmatrix}. \quad (9.11)$$

We assume that the nodes' memories are perfect (i.e., there is no decoherence) and that all gates can be executed noiselessly. The goal is then for the agent to find a sequence of actions that results in an increase in fidelity. The actions that the agent can take are as follows:

- T-gate on each qubit (4 actions),
- Hadamard gate on each qubit (4 actions),
- CNOT gates between qubits 0 and 1 at the same node (2 actions),
- Z-basis measurements on each qubit (4 actions),
- Accept/Reject (2 actions).

The T-gate is defined as

$$T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \quad (9.12)$$

This adds up to a total of 16 actions. In order to reduce the search space to be explored by the agent, we also forbid sequences of actions that would result in no entanglement being shared between the two nodes. Namely, the agent is not allowed to measure both qubits of either of the nodes. The observations available to the agent are the previous actions taken in the current trial. Different measurement outcomes lead to different branches and hence to different observations. The reward function R we employed was

$$R = \max(0, p\Delta F), \quad (9.13)$$

where p is the probability of success (which corresponds to the combined probability of the accepted branches) and ΔF is the difference between the average fidelity of the accepted branches and the input fidelity.

Running 50 agents on this environment for 250000 trials, we surprisingly found that the best protocol found was not EPL. Although EPL was also found by some of the agents,

it was not the protocol accruing the highest reward value. Informally, the EPL protocol consists of applying CNOT gates on both nodes' qubits locally, measuring the target qubits and keeping the remaining the entangled pair if the measurement outcomes are 11. The protocol found to be best is identical, except for the fact that we also accept the case when the measurement outcomes are 00. It is easy to see that this protocol, which we call EPL+, results in a higher value of the reward function than EPL. It performs identically for the 11 branch, and by definition of the reward function the contribution of any given branch must be non-negative, so accepting more branches cannot reduce the amount of reward obtained. We note the following about the EPL+ protocol: (i) just as the EPL protocol it always outputs an R-state, but it does so with higher probability and (ii) for initial fidelities above $2/3$ it also outputs an R-state with positive ΔF even for the 00 branch. With this in mind, it is not only clear that the EPL+ protocol achieves a higher reward than EPL, but it is also a more useful protocol from a practical perspective. The state obtained in the 00 branch is at the very least an R-state, even if one of lower fidelity than initial. But any R-state is better than no state, given that an R-state can be purified with some probability irrespective of its initial fidelity. We finally note that this result does not contradict the optimality result of [2], as EPL+ does not achieve a higher output fidelity than EPL.

9.4.2 Purify, entangle and discard

We now introduce another possible application of our methodology. The purification scenario we introduced in Section 9.4.1 was useful for validating our approach and even resulted in an unexpected result, with EPL+ outperforming EPL according to the reward function we defined. However, there is no time-dependency in that scenario, which means that it does not take full advantage of the use of a tool like NetSquid.

With this in mind, we introduce the purify, entangle and discard (PED) scenario. Here, two remote nodes with imperfect two-qubit memories wish to share high-quality entanglement of fidelity at least F_t . To do so, they can perform three different operations: attempt entanglement generation, $2 \rightarrow 1$ purification and discard entanglement. Entanglement generation succeeds with probability p and in case of success the two nodes will share a depolarized Bell state of fidelity $F_i < F_t$, so as not to render the problem trivial. States stored in memory undergo depolarizing noise characterized by a coherence time T_c as per Equation 9.1. The expected time for entanglement generation to succeed is given by $T_e = 1/p$. We note that the ratio between the coherence time and the expected time for entanglement generation to succeed $N_c = T_c/T_e$ fully captures the time dynamics of the system. We assume purification corresponds to one round of DEJMPS. The time-dependency of this problem arises due to the probabilistic nature of entanglement generation, which directly affects how long states spend decohering in memory. This renders it hard to study analytically.

Let us now give a detailed description of how this scenario can be conceptualized as a RL task. The goal is for the agent to find a sequence of actions that results in the nodes sharing an entangled pair of fidelity at least F_t . The actions that the agent can take are as follows:

- Generate entanglement (1 action),

- Discard entanglement (2 actions),
- Purify (2 actions).

There are two actions for discarding, one for each memory position. That is, the agent can decide to discard the entangled pair in position 1 or in position 2. Similarly, there are two actions for purifying as the agent can decide to keep the pair in position 1 or in position 2, and discard the other. Note that action availability is dependent on the state of the environment. Concretely, the generate action is only available when at least one memory position is free, the discard action for a given position is only available when that position is taken and the purify actions are only available if both positions are taken.

The information about the environment exposed to the agent via observations is in the form of a nine-dimensional vector. The first four entries of the vector are the coefficients corresponding to each of the Bell states for the entangled state stored in position 1 expressed in the Bell basis. The same is true for the following four entries with regard to the entangled state stored in position 2. In case there is no state in a given position, the corresponding four entries all have value 0. The ninth and final entry of the vector corresponds to the simulation time that has elapsed since the trial began.

We now briefly discuss this choice for the form of the observation vector. A more naive choice could be to expose the fidelities of the entangled states stored in memory, rather than their coefficients in the Bell basis. This results in a much simpler state, but provides only incomplete information. This is because the performance (i.e., output state and probability of success) of DEJMPS, the purification protocol employed, depends not only on the fidelity of the input states but instead on all their coefficients in the Bell basis [9, 15]. Further, there is also some redundancy in the information exposed to the agent. The coefficients of a state in the Bell basis must add up to 1 in order for the state's density matrix to be valid. Hence, one of the coefficients can be derived from the other three and including it is redundant. However, we observed that the agent performed best when this redundancy was present. Finally, the simulation time elapsed was included due to the stochastic nature of entanglement generation. Without this, there would be no way for the agent to know how much time had elapsed (in opposition, without stochasticity the number of actions taken could be directly mapped to the elapsed time) and hence information would be incomplete.

A trial concludes whenever the nodes share an entangled pair of fidelity at least F_t (success) or a maximum number of actions N_{\max} (failure) has been taken. The reward awarded to the agent is

$$R = \begin{cases} 1 & \text{if success,} \\ 0 & \text{otherwise.} \end{cases} \quad (9.14)$$

While this is a relatively simple environment in terms of how many actions it consists of there are a couple of factors that make it challenging. Namely, it is highly stochastic and rewards are sparse. We note that this is the simplest possible form of this environment, and that some extensions can be done to render it more complex. For example, we assume that if the entanglement generation action is taken it is always done until success. Making available actions where the agent can decide to attempt entanglement generation for a given period of time only (corresponding to some expected success probability) could allow for more sophisticated strategies.

To the best of our knowledge there are no known optimality results for this problem.

Results

We will now present some results obtained for this environment. We start by introducing a heuristic policy we designed with the intent of comparing its performance with that of the policies found through a RL approach, which we call the purify-only policy.

Purify-only As the name indicates, the purify-only policy consists of the following steps:

1. If there are two entangled pairs in memory, perform purification and keep the highest-fidelity pair,
2. If not, generate entanglement.

The discard actions are never used. This is not an optimal policy. For one, always keeping the highest-fidelity pair when purifying is not an optimal choice, as it might be that the lower-fidelity pair is "easier" to purify given its distribution of coefficients in the Bell basis [15]. It is also certainly the case that discarding is sometimes a better choice than attempting purification, as attempting to purify while consuming a pair of too-low quality will just decrease the quality of the pair that is kept, even in case of success. Nevertheless, it is likely not a terrible heuristic. While the fidelity is an imperfect metric for how useful an entangled state is for purification, it will likely often be the case that the highest-fidelity pair is in fact the one that should be kept. Further, assuming one remains in a high-fidelity regime, discarding is likely not often useful. We note that designing heuristic policies that make use of the discard actions is non-trivial, mostly due to the fact that the performance of DEJMPS on all the Bell-basis coefficients of the entangled states, rendering simple policies of the type 'discard-if-fidelity-below-threshold' useless.

The metric we use for evaluating a policy's performance is the simulation time until success has been achieved, averaged over 1000 executions. The choice of the number of executions strikes a balance between statistical significance and computational feasibility. We use Proximal Policy Optimization (PPO) [16], in particular a maskable version that allows for varying action spaces (i.e., it accounts for the fact that the actions made available to the agent depend on the environment's state) [17]. The implementation we use is freely available at Stable Baselines3 [18]. We employ Optuna for hyperparameter optimization [19]. We perform the learning process for ten different agents using the best hyperparameters found. We show in Table 9.1 the performance of the purify-only policy compared to the best and worst policies found through our RL approach. We see that both the best and the worst (and hence all) of the policies found using PPO outperform the purify-only policy.

We briefly note that this environment can easily be extended to have a connection with a quantum-networking application, rather than simply having an arbitrary fidelity target. In particular, one can extend the environment to quantum key distribution (QKD) [20, 21] in its entanglement-based form [22] by adding measurement to the set of possible operations. In this case, the set of actions available to the agent would be:

- Generate entanglement (1 action),

Policy	Time until success
Purify-only	174.0 \pm 5.5
Worst found	163.1 \pm 4.9
Best found	145.1 \pm 4.6

Table 9.1: Performance of purify-only and best and worst policies found through reinforcement learning approach as measured by average time required until success. The uncertainty is given by the standard error of the mean, and the average is taken over 1000 executions. The environment considered had a ratio between the coherence time of the memories and expected time for entanglement generation $N_c = 5$, fidelity of generated entangled states $F_i = 0.88$ and target fidelity $F_t = 0.9$. The time unit is chosen such that entanglement generation takes on average 10 time steps. This choice is arbitrary.

- Discard entanglement (2 actions),
- Purify (2 actions),
- Measure (2 actions).

There is one measure action for each of the entangled states, with all other actions being the same as for the previous version of the environment. The termination condition would also require changing. One possibility would be to determine that an episode consists of a T_{episode} time steps, with the goal being to generate the maximum amount of key bits K (see Chapter 3 for details) in that time. An appropriate reward would then be:

$$R = K. \quad (9.15)$$

Code implementing this variation of the environment is also available at [13].

9.5 Conclusions

We have combined a reinforcement learning approach for protocol finding with the use of the quantum-network simulator NetSquid. Further, we have given two example applications of this approach and made our code freely available at [13]. We see two main lines of work arising from this chapter. First, having introduced and validated this approach, it can be extended to more complex scenarios that might be of real-world interest. For example, an interesting approach would be to define an environment corresponding to a repeater chain as done in [5] and allow for heterogeneous nodes with asymmetric placement. The NetSquid simulation framework introduced in Chapters 6 and 7 could be of great use for this purpose. We believe this would result in the design of end-to-end entanglement generation protocols tailored to a particular repeater chain that could outperform traditional ones [4, 23–25]. Secondly, we believe the PED scenario we have introduced is interesting in its own right. We have already seen that this reinforcement learning approach can find policies outperforming a simple heuristic. One could further investigate if the policies found using reinforcement learning can shed light on how to design better heuristics. It would also be interesting to see how the policies change when changing the goal of the process (i.e., going from a target fidelity to secret-key rate). Another point we briefly touched upon and that merits further study is that of how much freedom the agent gets in the entanglement generation process. Here we have assumed that when the

agent chooses to generate entanglement, it can only do so by committing to performing attempts until it succeeds. This does not have to be the case: it could choose instead to attempt entanglement generation up to only n times, for any integer n . This would likely enable better policies, albeit at the cost of more computational resources. The question of whether policies found for a given set of environment parameters (i.e., the ratio between coherence time and time for entanglement generation, the target fidelity and the initial fidelity) perform well for other environment parameters also strikes us as interesting.

References

- [1] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead*, *Science* **362**, eaam9288 (2018).
- [2] F. Rozpedek, T. Schiet, L. P. Thinh, D. Elkouss, A. C. Doherty, and S. Wehner, *Optimizing practical entanglement distillation*, *Physical Review A* **97**, 062333 (2018).
- [3] Á. G. Iñesta, G. Vardoyan, L. Scavuzzo, and S. Wehner, *Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs*, *npj Quantum Information* **9**, 46 (2023).
- [4] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpedek, M. Skrzypczyk, *et al.*, *Netsquid, a network simulator for quantum information using discrete events*, *Communications Physics* **4**, 164 (2021).
- [5] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, *Machine learning for long-distance quantum communication*, *PRX Quantum* **1**, 010301 (2020).
- [6] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction* (MIT press, 2018).
- [7] F. Rozpedek, K. Goodenough, J. Ribeiro, N. Kalb, V. C. Vivoli, A. Reiserer, R. Hanson, S. Wehner, and D. Elkouss, *Parameter regimes for a single sequential quantum repeater*, *Quantum Science and Technology* **3**, 034002 (2018).
- [8] B. Li, T. Coopmans, and D. Elkouss, *Efficient optimization of cutoffs in quantum repeater chains*, *IEEE Transactions on Quantum Engineering* **2**, 1 (2021).
- [9] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, *Quantum privacy amplification and the security of quantum cryptography over noisy channels*, *Phys. Rev. Lett.* **77**, 2818 (1996).
- [10] S. B. v. Dam, P. C. Humphreys, F. Rozpedek, S. Wehner, and R. Hanson, *Multiplexed entanglement generation over quantum networks using multi-qubit nodes*, *Quantum Science and Technology* **2**, 034002 (2017).
- [11] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356**, 928 (2017).

- [12] G. Avis, F. Rozpędek, and S. Wehner, *Analysis of multipartite entanglement distribution using a central quantum-network node*, *Physical Review A* **107**, 012609 (2023).
- [13] *NetSquid-ReinforcementLearning*, <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-reinforcementlearning> (2023).
- [14] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, *Purification of noisy entanglement and faithful teleportation via noisy channels*, *Physical review letters* **76**, 722 (1996).
- [15] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller, *Quantum repeaters based on entanglement purification*, *Physical Review A* **59**, 169 (1999).
- [16] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, *Proximal policy optimization algorithms*, arXiv preprint arXiv:1707.06347 (2017).
- [17] S. Huang and S. Ontañón, *A closer look at invalid action masking in policy gradient algorithms*, arXiv preprint arXiv:2006.14171 (2020).
- [18] A. Raffin, A. Hill, M. Ernestus, A. Gleave, A. Kanervisto, and N. Dormann, *Stable baselines3*, <https://github.com/DLR-RM/stable-baselines3> (2019).
- [19] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, *Optuna: A next-generation hyperparameter optimization framework*, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2019).
- [20] A. K. Ekert, *Quantum cryptography based on bell's theorem*, *Physical review letters* **67**, 661 (1991).
- [21] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, *Theoretical computer science* **560**, 7 (2014).
- [22] C. H. Bennett, G. Brassard, and N. D. Mermin, *Quantum cryptography without bell's theorem*, *Physical review letters* **68**, 557 (1992).
- [23] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum repeaters: the role of imperfect local operations in quantum communication*, *Physical Review Letters* **81**, 5932 (1998).
- [24] N. Sangouard, C. Simon, H. De Riedmatten, and N. Gisin, *Quantum repeaters based on atomic ensembles and linear optics*, *Reviews of Modern Physics* **83**, 33 (2011).
- [25] L. Kamin, E. Shchukin, F. Schmidt, and P. van Loock, *Exact rate analysis for quantum repeaters with imperfect memories and entanglement swapping as soon as possible*, *Physical Review Research* **5**, 023086 (2023).

10

Outlook

In this chapter we give a brief summary of results presented in this dissertation, as well as of possible future lines of research.

10.1 Summary of results

Minimal hardware requirements. Many of the results of this dissertation concern themselves with hardware requirements for quantum repeaters. By this we mean the smallest improvement over state-of-the-art hardware that allows for entanglement distribution using quantum repeaters meeting certain performance metrics. To determine them, we started by introducing an open-source, general and modular optimization methodology integrating genetic algorithms and NetSquid simulations of quantum-repeater networks. We then applied this methodology to determine minimal hardware requirements for different scenarios:

- A simple form of verifiable blind quantum computing (VBQC) between a client and a server separated by roughly 230 km of real-world optical fiber using a color-center or trapped-ion-based quantum repeater.
- Quantum key distribution and VBQC test rounds over more than 900 km of real-world optical fiber.
- Entanglement distribution meeting certain fidelity and rate targets over distances up to 1600 km while simultaneously optimizing over the number of repeaters placed and the protocols they execute.

Maximizing hardware performance. We have also concerned ourselves with the question of how we can make the most out of imperfect hardware through smart protocol choices and repeater placement. We have used the aforementioned methodology to simultaneously optimize over hardware and protocol parameters, showing that good protocol choices can result in significantly lower hardware requirements. We have also developed tools for performing reinforcement learning experiments using NetSquid, which allows for the exploration of protocols for high-quality entanglement generation and distribution tailored to specific, imperfect hardware.

Reducing the impact of constraints imposed by fiber networks. We have given formulas for determining quantum-repeater placements in preexisting locations (a cost-effective way of deploying quantum networks [1] at the cost of some performance loss [2]) if midpoint stations are required. We have determined hardware requirements under the assumption of constrained placement of network nodes, and explicitly compared how much larger the requirements are on a real-life fiber grid when compared to an idealized scenario where all nodes are equally spaced.

10.2 Future work

We have in Chapter 6 introduced the concept of absolute minimal hardware requirements, i.e., the value of a parameter that is required to achieve a certain level of performance if every other parameter (except for fiber attenuation) is perfect. One possible approach experimentally would be to try to get a setup to first meet all absolute minimal hardware requirements, as they are anyway a necessity. Determining minimal hardware requirements from that point in parameter space would be interesting, as they could be different from the minimal hardware requirements we have determined. This would result in another avenue to functional quantum-repeater networks.

Large parts of this dissertation focused on determining minimal hardware requirements for first-generation quantum repeaters. However, as discussed in Chapter 3, other quantum-repeater architectures exist. It would be interesting to perform a comparison of the hardware cost for generating entanglement between distant nodes using different generations of quantum repeaters. For example, one could quantify whether it is easier to achieve high-rate high-quality entanglement distribution using one chain of third-generation quantum repeaters, which require that repeaters be placed very closely together but promises high rates, or multiple parallel chains of first-generation quantum repeaters. The optimization approach introduced in Chapter 5 could lend itself to well to answering this question.

In real-world fiber networks some links suffer from more attenuation than others. It might then be that it is wise to adopt a different purification strategy per link, as for example it is more costly to regenerate a link in which attenuation is very high. Other examples of possible protocol choices that we have not explored are the use of cut-off timers and different swapping strategies. Optimizing over such protocols could enable further reductions in hardware requirements.

In Chapter 6 we showed that the VBQC protocol introduced in [3] is secure if the average probability of error can be bounded. However, we still assumed that the error probabilities in different rounds of the protocol were independent and identically distributed. If a quantum-repeater chain is used to distribute the entanglement that the client and server will consume to perform VBQC, this assumption will likely not be true. It would then be interesting to see if this assumption can be relaxed, for example by bounding the correlation between error probabilities between different rounds. Such a bound could then potentially be related back to properties of the entanglement distributed by the chain.

In Chapter 4 we gave formulas for counting the different configurations in which quantum repeaters and midpoints can be placed along a line. These results can likely be extended to arbitrary networks.

In Chapter 9 we introduced the purify, entangle and discard (PED) scenario as an example application of our tool for quantum-network protocol exploration. Studying this scenario further would be of interest. Namely, it is unclear whether there are provably-optimal policies and what they are.

The tools introduced in 9 could likely be applied to discover quantum-repeater chain protocols, in the vein of what was done in [4]. We envision this being useful when considering a very specific scenario: a particular real-world fiber path and given hardware quality, as one likely would when deploying quantum-repeater hardware in the real world. This approach could then be used to determine the protocol that would result in the best possible performance for this particular scenario.

References

- [1] J. Rabbie, K. Chakraborty, G. Avis, and S. Wehner, *Designing Quantum Networks Using Preexisting Infrastructure*, npj Quantum Information **8**, 5 (2022), 2005.14715 .
- [2] G. Avis, R. Knegjens, A. S. Sørensen, and S. Wehner, *Asymmetric node placement in fiber-based quantum networks*, arXiv preprint arXiv:2305.09635 (2023).
- [3] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, *Verifying bqp computations on noisy devices with minimal overhead*, PRX Quantum **2**, 040302 (2021).
- [4] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, *Machine learning for long-distance quantum communication*, PRX Quantum **1**, 010301 (2020).

Acknowledgments

This dissertation is the longest project I have completed. It makes me very happy to reflect on and express my gratitude towards everything and everyone that helped me along the way. First of all, I am grateful for having been born into circumstances that allowed me to develop into a person that had a chance of even being selected for a project like this. I am not sure who to thank for this: God, the Universe or my lucky stars. In any case, I am grateful for how things seem to work out as long as I try my best.

Which brings me to **Stephanie**. You took a chance on me, and I am very thankful for it. I still remember how crazy it felt that you would fly me halfway across Europe to meet you and your group based on a ten-minute chat! I am very happy you did. You are one of the most analytical and inquisitive people I know. I would like to think that these characteristics, so valuable in research, have to some extent rubbed off on me. My PhD coincided with some major events, both globally and personally: a worldwide pandemic, your loss and regaining of eyesight and the birth of Sara. This meant that the road felt at times bumpy, but I am proud of how we navigated it. If I came out the other side looking something like a researcher, it is in large part thanks to you.

I would also like to extend my thanks to the people who agreed to put on fancy outfits and discuss my work with me for an hour: thank you **Ronald, Gary, Przemek, Julien** and **Elham** for being on my committee.

Most of my PhD was spent, along with several others, developing code that was to be the basis for a "blueprint for the quantum internet". The degree of our success is arguable, and at times I felt like a kid roleplaying as a software engineer. Nevertheless, I had a blast doing it. This is in large part due to **Guus**. You were my closest collaborator, as evidenced by the fact that half of our theses overlap, and I don't think I could have asked for someone better for that role. Whether we were at the office, each sitting at home in front of our respective screens, or in one of the seventeen Starbucks we visited for work purposes in New York, you always remained curious, kind and integrous. You are my model for how an academic should behave and Stefan is lucky to have you. I realize that I will here again contribute to our theses overlapping even further, but I really do appreciate how you are always genuinely interested in my ramblings on etymology (they seem to not be such a hit elsewhere). Having recently successfully completed your defense (in which you, by the way, set a very high standard for me) I am very happy to start thinking of you as "only" a friend. I wish you and Eva all the best on your new adventure.

On the topic of the blueprint, I must also mention **David**. I remember when I first met you I was somewhat intimidated by how cool you looked, which seems very silly in retrospect because you also turned out to be one of the nicest people I know. The periods in which we shared an office were the most fun of my PhD, even if not the most productive (but as you very well know, there's much more to life than productivity). While my wallet definitely does not thank you for having gotten me into Magic: The Gathering, I do: playing with you in the office (exclusively outside working hours, of course) was

great fun. I am very happy you have found something that works better for you at SURF and I am sure that whatever your next step turns out to be, your combination of technical skills and people skills will be greatly appreciated.

Besides the "full-time" members of the blueprint team, plenty of part-time members have come and gone. **Tim**, I have never left a conversation with you feeling worse about myself: you are incredibly kind. It seems to me you are well on your way to becoming a PI, and I can think of no one better for such a role. **Julian**, I find the way in which you have managed to carve out a niche for yourself at TNO very impressive. I hope our paths keep crossing, because whether it was in front of a whiteboard or with a beer in hand, it's always been fun. **Axel**, I am in awe of how you managed to do about seven PhDs worth of work while still being a great colleague and an amazing climber on the side. **Hana**, despite having joined us while you were still a bachelor's student you proved to be a valuable addition to the team: it was nice to have someone around that could actually code. The fact that you were always down for a beer was also much appreciated, of course. **Ariana**, you were always a pleasure to talk to. Thank you for all your help and patience with the many times I came to you crying about how a bash script had hurt me. **Rob, Loek, Luc, Martijn** and **Flors**, it was great having you on the team and I wish you all the best.

During the course of my PhD I have had the privilege of supervising three master students. These were all very different experiences, but I both enjoyed and learned a lot from all of them. **Adrià**, you had it particularly rough because (i) your MEP was done almost completely remotely and (ii) you were my first student so I didn't really know what I was doing. You nevertheless did amazingly well (as seen in Chapter 8 of this dissertation) and transformed yourself into an independent researcher in a flash. It is no surprise to me how well you are doing in your own PhD. **Yu**, the language barrier made it tough at times, but I am proud of how we got through it and of what you achieved. **Wesley**, you find yourself in the somewhat awkward position of defending after your supervisor does. But given how easygoing, relaxed and hardworking you have shown yourself to be, I'm sure that won't be a problem for you.

Besides the people I got to directly work with over the course of my PhD, I also had the privilege of being a part of large, social group of smart and pleasant people in the form of the Wehner group (and its offshoots). Given their transient nature, research groups lend themselves well to splitting people into two categories: those who were there before, and those who came after. I'll start with the former. **Carlo**, if I had to associate one word with you, it'd be acceptance. You seem to always take people as they are, and this makes it very hard not to love you. I have very much enjoyed our time together, whether it was on random trips to France, doing any one of a multitude of sports or becoming experts on Korean cinema. I am very happy that you are now done with what I imagine must have been a very frustrating experiment (experience?) and that a new chapter can begin! **Sébastien**, het was een genoegen om gelijktijdig onze PhD's te doen. Of we nu over sport of over wetenschap spraken, het was altijd gezellig bij jou! Het is een beetje jammer dat onze ideeën over purificatie en optimalisatie nergens toe hebben geleid, maar ik vond het in ieder geval leuk. Ik waardeer je geduld en vriendelijkheid zeer toen ik begon te proberen in mijn soort-van-Nederlands tegen jou te praten. En je hebt zeker de beste DnD karakters! **Kenneth**, I very much appreciate how easy it is to talk to you, be

it about noise rap, relationships or the little nuggets of math you constantly manage to find in everyday life. I am very happy for you that you managed to not let bad academic experiences crush your curiosity. Having missed the latest Death Grips concert (for good reason), I hope we get to go to one together soon! **Kaushik**, thanks for all the dinners and nights out in those seemingly-distant pre-pandemic times! It's a shame that trip to Scotland never materialized, but maybe there will now be time for one to Singapore. **Matt**, thanks for showing (and driving) me around Chicago! **Álvaro**, I hope to one day make a slide deck that looks half as good as one of yours. **David E.**, your calm, encouraging presence was very much appreciated in my early days with the blueprint. **Filip**, your curiosity (including but not limited to research) is inspiring! **Bas, Jérémy, Mark, Victoria, Leon**, we didn't get the chance to properly get to know each other, in part due to the pandemic. You were nevertheless a very welcoming presence when I first joined the group. Thank you!

Ravi, it's great that you managed to transition to a role that suits you better without giving up on your research. Hopefully it also means that we will have plenty of time for more Kendrick concerts and stand-up shows (but maybe no more Oppenheimer screenings). **Janice**, your kindness and spontaneity are a real breath of fresh air in a group of mildly-repressed nerds. **Bart**, you come up with the best package names! **Scarlett**, I admire the extent to which you abide by your priorities. **Bethany**, you were great company during our US roadtrip and great opposition during our doubles matches! **Soubhadra**, it's great to see you picking up where Guus and I left off. I'm sure you'll do great! **Lars**, I hope we still have some more tight doubles matches now that you've become a master! **Emlyn**, you exude good vibes. Every room feels friendlier with you in it. **Mick**, I hope we still get to play some tennis! I feel like you owe me a match after snubbing me for your MEP ;) **Thomas**, it was a pleasure sharing a DnD table with you. Your encyclopedic knowledge of the rules made my first forays into the world of DMing that much less intimidating. **Luise**, it has been a pleasure getting to know you (and Ella) in the early days of your PhD! **Hemant**, the same, minus Ella, holds for you. **Jeroen**, I find it extremely impressive how only a few weeks into your PhD you already behave so much like a proper, independent researcher. **Eric**, thanks for all the brownies! **Siddhant**, I'm not sure how you manage to do so many cool things on the side of your PhD, but it's great that you do. **Gayane**, your arrival to the group brought along very different (and valuable!) perspectives, scientific and otherwise. It's great to see how well you're already doing as a new PI. **André**, *foste demasiado cedo. Foi um prazer ter em ti um bocadinho de casa em Delft antes de ires.*

One of the things that drew me to QuTech when I came for my interview some four and half years ago was the sense of community I felt, which was amazing and very different from my previous experience with research. This was definitely a correct impression, and I have gotten to meet some great people just by virtue of sitting in the same building as them. **Vicky**, some people's laughter is said to fill the room; yours fills up the whole building. I don't think I have once been bored in your presence. Thank you for being there and for being the best rave mom I could've asked for. **Ravi**, you've been a great workout buddy and you could also be a great chess buddy if only you weren't much too good for me. **Matt**, we've slowed down on the bouldering lately, but I had a lot of fun getting into it with you! Thanks for being so chill about having a bunch of nerds invade your lovely home to play pretend. **Anta**, your calm and content demeanor is contagious. **Jorge**, *muito fixe ver como*

o teu ténis tem evoluído! A este ritmo não há-de faltar muito para me começares a dar problemas ;) Another privilege of doing a PhD at QuTech as a theorist is to be surrounded by brilliant experimentalists who will gladly help you understand their setup, critique your modelling and review your papers. I'd like to thank all the members of the Hanson group who have done this for me: **Alejandro, Arian, Kian, Mariagrazia, Matteo** and **Sophie**, among others. Together with the **QINC social team** I had the pleasure of organizing virtual pub quizzes and (rainy) beach days. Thank you to **Max, Joe, Jake, Sjoerd, Vicky, Julius, Yannik** and **Maurice** for making the process that much more fun!

What was initially meant to be a short stint turned, partially due to the pandemic, into a longer stay at **Piet Heinstraat 82**. This was made very enjoyable by **Alex, Duncan, Fernando, Merlijn, Rui, Vitto** and **Vlad**. Thank you!

After a move- and pandemic-induced break, I picked up my rackets again and got back into tennis. This happens to be a hard thing to do without someone on the other side of the net, so there a few people I have to thank for making that possible. Ten eerste, de mannen van **Obvius Heren 1: Jippe, Kop, Misha, Tromp, Vito** en **Youri**. Ik weet niet zeker of ik jullie moet bedanken voor mij kennis te laten maken met daltons, maar het was echt leuk om samen met jullie weer wedstrijdtennis te spelen. **Kanta**, the way we first got to playing was very random, but I'm glad it happened. You're to blame for any improvements in my serve: seeing you annihilate my second serves was so frustrating that I just had to work on them. Here's hoping that we get to win many more doubles tournaments together! **Câmara**, aprecio muito que viajes 1h30 para vir jogar a Delft! Fico à espera que voltes a treinar para podermos repetir as batalhas de antigamente como deve ser ;) **Sjoerd**, thanks for all the early-morning sessions, and of course also for putting me in touch with **Vid!** Vid, I really appreciate how open you are, and while your goals are lofty you definitely have the passion to make them happen.

Mudar-me para Delft passados quase seis anos muito felizes em Lisboa não foi de todo uma decisão fácil, mas o facto de continuar a sentir que tinha um lugar para onde voltar em Lisboa ajudou: mostrou-me que nenhum adeus é para sempre. Por isto tenho muita gente a quem agradecer: **Bernardo, Condesso, David, Diogo, Filipe, Fred, Gabriel, Gonçalo, Henrique, Inês, João A., João O., Moreira, Nelson, Pardal, Patrícia, Rodrigo, Senart, Tiago, Tomás**, et al. Obrigado! **Ana**, obrigado por todas as vezes em que me ajudaste a ver as coisas de uma perspectiva diferente e frequentemente mais razoável. Faz-me muito feliz que continuemos a ser amigos e uma parte importante da vida um do outro passado tanto tempo e depois de tudo pelo que passámos juntos. **Cris**, foste uma presença importante durante a primeira fase da minha vida na Holanda e apesar de nos termos apercebido que não éramos a melhor escolha um para o outro, fico contente que tenhamos "feito as pazes". Desejo-te (e à Suri) tudo de bom.

Doing a PhD may 'only' take four years, but getting to the point where you're ready to embark on one takes a lot longer. My family deserves a lot of the credit for getting me to that point. **Pai**, obrigado por me ensinares o valor do trabalho e de fazer o que está certo mesmo quando o certo e o agradável não coincidem. **Mãe**, obrigado por me ouvires sempre e por me fazeres sentir que o que penso tem valor. **Teresa**, os últimos tempos não têm sido fáceis para ti, mas espero e acredito que o teu futuro com o Luís será risonho. **Carolina**, na minha cabeça de irmão mais velho o meu instinto será sempre um pouco pensar em ti como uma criança, mas a verdade é que já és adulta e licenciada! Espero que

saibas que podes sempre contar comigo para o que quer que seja.

muito amor para todos vós

Curriculum Vitæ

Francisco Ferreira da Silva

- 2019-2023 PhD in Quantum Networks
QuTech, Delft University of Technology, the Netherlands
Dissertation: Hardware and Protocol Optimization in Quantum-
Repeater Networks
Promotor: Prof. dr. S.D.C Wehner
Promotor: Prof. dr. ir. R. Hanson
- 2016-2018 M.Sc. in Engineering Physics
Instituto Superior Técnico, University of Lisbon, Portugal
Thesis: Quantum perceptrons
Supervisor: Dr. Yasser Omar
Supervisor: Dr. João Seixas
- 2013-2016 B.Sc. in Engineering Physics
Instituto Superior Técnico, University of Lisbon, Portugal
- 1995/08/14 Born in Vila Nova de Gaia, Portugal

List of Publications

9. **Francisco Ferreira da Silva**, Stephanie Wehner, *Protocols for generation of high-quality entanglement using reinforcement learning*, in preparation.
This article is included in this dissertation as Chapter 9.
8. Guus Avis*, **Francisco Ferreira da Silva***, Kenneth Goodenough*, Stephanie Wehner, *Counting quantum-repeater configurations*, in preparation.
This article is included in this dissertation as Chapter 4.
7. Adrià Labay Mora, **Francisco Ferreira da Silva**, Stephanie Wehner, *Reducing entanglement-distribution hardware requirements via joint hardware-protocol optimization*, preprint arXiv:2309.11448 (2023).
This article is included in this dissertation as Chapter 8.
6. **Francisco Ferreira da Silva***, Guus Avis*, Joshua A. Slater, Stephanie Wehner, *Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber*, preprint arXiv:2303.03234 (2023).
This article is included in this dissertation as Chapter 7.
5. Guus Avis*, **Francisco Ferreira da Silva***, Tim Coopmans, Axel Dahlberg, Hana Jirovská, David Maier, Julian Rabbie, Ariana Torres-Knoop, Stephanie Wehner, *Requirements for a processing-node quantum repeater on a real-world fiber grid*, npj Quantum Inf 9, 100 (2023).
This article is included in this dissertation as Chapter 6.
4. João P. Moutinho, Marco Pezzuto, Sagar Pratapsi, **Francisco Ferreira da Silva**, Silvano de Franceschi, Sougato Bose, António T. Costa, Yasser Omar, *Quantum dynamics for energetic advantage in a charge-based classical full-adder*, PRX Energy 2, 033002 (2023).
3. Chin-Te Liao, Sima Bahrani, **Francisco Ferreira da Silva**, Elham Kashefi, *Benchmarking of quantum protocols*, Scientific Reports 12, 1 5298 (2022).
2. **Francisco Ferreira da Silva**, Ariana Torres-Knoop, Tim Coopmans, David Maier, Stephanie Wehner, *Optimizing Entanglement Generation and Distribution Using Genetic Algorithms*, Quantum Science and Technology 6, 3 (2021)
This article is included in this dissertation as Chapter 5.
1. **Francisco Silva**, Mikel Sanz, João Seixas, Enrique Solano, Yasser Omar, *Perceptrons from Memristors*, Neural Networks 122, 273-278 (2020).

* These authors contributed equally to this work.