Distributing entanglement in quantum networks

Goodenough, K.D.

Publication date
2022

Document Version
Final published version

Citation (APA)
Goodenough, K. D. (2022). *Distributing entanglement in quantum networks*. [Dissertation (TU Delft), Delft University of Technology].

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

DISTRIBUTING ENTANGLEMENT IN QUANTUM 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 1 februari 2022 om 10:00 uur

door

Kenneth DUDLEY GOODENOUGH

Master of Science in Applied Physics,
Technische Universiteit Delft, Nederland
geboren te Kempton Park, Zuid-Afrika.

Dit proefschrift is goedgekeurd door de

promotor: Prof. dr. S.D.C. Wehner
copromotor: Dr. D. Elkouss Coronas

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof. dr. S.D.C. Wehner	Technische Universiteit Delft, promotor
Dr. D. Elkouss Coronas	Technische Universiteit Delft, copromotor

Onafhankelijke leden:	
Prof. dr. L. DiCarlo	Technische Universiteit Delft
Dr. ir. M. Veldhorst	Technische Universiteit Delft
Prof. dr. M. Razavi	University of Leeds
Dr. J. Borregaard	Technische Universiteit Delft
Prof. dr. J.M. Thijssen	Technische Universiteit Delft, reservelid

Keywords: quantum, quantum networks, quantum repeaters, quantum communication, entanglement distillation

Printed by: Gildeprint - www.gildeprint.nl

Front & Back: Visualisation of the norm function on $\mathbb{Z}[i]/p\mathbb{Z} \simeq \mathbb{F}_{p^2}$ for $p = 25087$ (with a Möbius transformation applied to the final image on the printed version). Created by K. Goodenough (with help from Joost de Jong).

ISBN 978-94-6419-405-0

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

So, the idea of Gordian Knot *(...) is that the moments of brilliance, convergence, beauty, and discovery that we attribute to those brilliant minds from history, appear to me as knots in that pervasive, ubiquitous fabric of the universe — as if to say, no matter how disparate they seem, be they advances in music, math, physics, etc., that they are all tied into that fabric and that they share a profound and common source.*

Sean Malone

CONTENTS

SUMMARY

The information-theoretic point of view on physics - and quantum mechanics in particular - has led to exciting insights on what distinguishes our world from a classical one. In particular, it was realised that quantum mechanics offers us communication capabilities above and beyond what can be realised in a non-quantum world. One important aspect of quantum communication is *quantum entanglement*. Entanglement is a strictly quantum-mechanical phenomenon, allowing for parties to share correlations that are impossible with only classical means. The art of devising quantum communication protocols is thus in exploiting these correlations. Initially the investigation of such quantum communication protocols was limited to idealised theoretical work, but in recent years there has been a flurry of proof-of-principle experiments and commercial availability of quantum communication devices.

While experimental advances have solidified the feasibility of practical quantum communication, distributing entanglement - or even only non-entangled states - remains a challenging experimental endeavour. This is because the distribution of quantum information is markedly different from its classical counterpart. The impossibility of cloning quantum information is exactly what provides security for certain quantum cryptographic protocols, but also prevents us from naively amplifying the signal to enlarge the distance over which qubits can be transmitted. Quantum bits (or *qubits*) can still be transmitted over large distances using the concept of *quantum repeaters*. However, their implementation also increases the decoherence experienced by qubits, a problem unique to quantum information. Thus, implementing quantum repeaters introduces a trade-off between the rate at which entanglement is established and the noise experienced by the states.

As is clear, there are still significant experimental and theoretical hurdles to overcome before a fully *quantum mechanical internet* can be realised. In such a quantum internet, it is envisioned that qubits can be distributed between any two points in the network, potentially on a global scale.

In this thesis, we investigate some of the pressing questions that need to be answered before a quantum internet can be realised. These questions mostly deal with how one should distribute entanglement. For example, what is the best way to distribute entanglement over a single repeater node? How about when there are multiple repeater nodes, not necessarily equally spaced? Is it possible to (more efficiently) counteract the effects of the noise incurred during the distribution?

The first two chapters are focused on an important building block for long-scale quantum communication - *a single quantum repeater*. We investigate the performance of such a quantum repeater and find parameters where it provides an advantage over direct transmission. Next, by linking quantum repeaters together into a *quantum repeater chain*, the distance over which quantum information can be effectively transmitted increases. However, the number of possible protocols that can be performed grows enor-

mously. We provide an algorithm for efficiently performing a heuristic optimisation over quantum repeater protocols, allowing us to better understand the form good quantum repeater protocols take and the required experimental parameters. Finally, we explore *entanglement distillation protocols*, which have as goal to combat the effects of noise in any realistic experimental setup. Entanglement distillation protocols condense multiple weakly entangled states into a (usually) smaller number of more strongly entangled states. We find all distillation protocols belonging to an experimentally relevant class of protocols for up to five arbitrary Bell-diagonal states. Furthermore, we use further symmetry reductions to find all such protocols for the case of up to eight copies of a Werner state. The found protocols surpass the fidelities and rates that were previously known to be achievable. Furthermore, we provide circuits with low depth and a small number of two-qubit gates that achieve the highest fidelity.

SAMENVATTING

Een informatie-theoretisch perspectief op de kwantummechanica heeft ons geleerd dat de wereld om ons heen zich op verrassende manieren onderscheidt van een Newtoniaans wereldbeeld. Men realiseerde zich dat kwantummechanica ons cryptografische mogelijkheden geeft, vele malen sterker dan wat gerealiseerd kan worden in een klassieke wereld. Een belangrijk aspect van kwantum communicatie is *kwantum verstrengeling*. Verstrengeling is een puur kwantum-mechanisch fenomeen, waardoor gesepareerde partijen correlaties kunnen delen die niet mogelijk zijn met enkel klassieke middelen. De kunst van het bedenken van kwantum communicatie toepassingen ligt in het uitbuiten van deze correlaties. Aanvankelijk werden dit soort toepassingen enkel theoretisch bestudeerd, maar in de laatste jaren zijn er steeds grotere stappen gemaakt met belangrijke *proof-of-principle* experimenten en zelfs commercieel verkrijgbare quantum communicatie apparaten.

Bovenstaande experimentele vorderingen hebben het duidelijk gemaakt dat kwantumcommunicatie niet enkel een abstract concept is maar ook in de praktijk toepasbaar zal zijn. Desondanks blijft het verspreiden van verspreiding en kwantum informatie een uitdaging in de praktijk. Dit komt omdat het verspreiden van kwantum informatie verschilt in vele belangrijke aspecten sterk van klassieke informatie. Het feit dat het onmogelijk is om kwantum informatie te kopiëren is vaak precies bepaalde cryptografische protocollen veilig maakt, maar zorgt er ook voor dat het onmogelijk is om naïef een signaal van qubits te versterken om zo de afstand waar kwantum informatie over verstuurd kan worden te vergroten. Desondanks is het nog steeds mogelijk om kwantum bits te versturen over grote afstanden door gebruik te maken van zogenaamde *kwantum herhalers*. Elke realistische implementatie van een kwantum herhaler zal ook extra *ruis* introduceren, wat de kwaliteit van de verstrengeling zal doen verminderen. Kwantum herhalers vergroten dus de snelheid van de generatie van verstrengeling, maar zorgen ook ervoor dat de verstrengeling minder bruikbaar wordt.

Het is evident dat er nog belangrijke experimentele en theoretische vraagstukken te overkomen voordat een *kwantum internet* gerealiseerd kan worden. In een mogelijke visie van het kwantum internet is het mogelijk om kwantum bits (beter bekend als *qubits*) te versturen tussen elke twee punten in het netwerk, mogelijk zelfs aan de andere kant van de wereld.

In dit proefschrift onderzoeken we enkele van de prangende vragen die beantwoord moeten worden voordat een kwantum internet gerealiseerd kan worden. De nadruk ligt hierbij vooral op de vraag hoe verstrengeling het best gedistribueerd kan worden. Wat is de beste manier voor het geval dat er maar een enkele kwantum herhaler is? Wat als er meerdere kwantum herhalers zijn, die niet noodzakelijk op gelijke afstand gepositioneerd zijn? Is het mogelijk om (efficiënter) de effecten van de ruis die is opgelopen gedurende de distributie tegen te gaan?

We beginnen bij de eerste twee hoofdstukken met een analyse van een enkele kwantum herhaler. We vinden experimentele parameters waar een dergelijke kwantum herhaler beter presteert dan directe transmissie. Vervolgens bestuderen we de situatie waar meerdere kwantum herhalers gekoppeld zijn aan elkaar tot een *kwantum herhaler ketting*. Het aantal mogelijke protocollen die uitvoerbaar zijn op een kwantum herhaler ketting groeit snel met het aantal kwantum herhalers in de ketting. Wij presenteren een heuristische optimalisatie, en gebruiken deze om de haalbare prestaties te analyseren van verschillende systemen en experimentele parameters. Tenslotte onderzoeken we *verstrengeling distillatie protocollen*. Dit zijn protocollen die als doel hebben om de effecten van ruis tegen te gaan. Dit gebeurt door meerdere zwakke verstrengelde paren te distilleren tot een kleiner aantal verstrengelde paren, maar die wel sterker verstrengeld zijn. We vinden alle mogelijke distillatie protocollen in een klasse van experimenteel relevante protocollen, die werken op vijf en minder arbitraire Bell-diagonale toestanden. Vervolgens gebruiken we de symmetrie in het geval van een n-voudige tensor kopie van een zogenaamde Werner state om alle protocollen te vinden voor acht en minder kopieën. Met de gevonden protocollen is het mogelijk om een hogere fideliteit en distillatie snelheid te behalen dan met voorgaande protocollen. Tenslotte geven we ook de circuits die de hoogste fideliteit behalen. Deze circuits hebben een lage diepte en een klein aantal twee-qubit poorten.

1

INTRODUCTION

The topic of this thesis is quantum communication, a relatively young subject. Initial research showing that quantum mechanics offers - in principle - relevant advantages over practical tasks, did not yet sway the pessimist that such experiments could actually be performed. Nowadays, experimental groups worldwide are pushing ever further what is possible, far exceeding what was imaginable when the first quantum information tasks were conceived. However, fault-tolerant quantum computation and a quantum internet - the two holy grails of quantum information processing - are yet to be realised [173].

A quantum internet is envisioned to ultimately allow any two parties in the world to generate entanglement between themselves, which they can then use for further quantum processing. Constructing and designing such a quantum internet requires not only immense experimental effort, but also guidance from a theoretical perspective. This thesis is but one of many steps required towards a theoretical understanding necessary for building the first quantum internet.

What would be the benefits of such a quantum internet? There exist certain tasks that are either impossible to perform using only classical resources, or for which quantum resources allow for advantages, such as increased security/privacy or reduced resources. Such tasks include the synchronisation of clocks [56, 80, 88, 133], distributed (quantum) computation [153], anonymous transmission of information [18, 33] and the distribution of secret keys [10, 49], which can be used for secure communication. A quantum internet would allow for two parties to perform such tasks at a global scale.

One of the main bottlenecks for quantum communication is photon loss — as (entangled) photons are transmitted through either free-space or fibre, they get lost with increasing probability the longer the distance. The rate at which classical and quantum communication tasks can be performed using only direct transmission is thus limited by the distance. For classical communication this problem is solved using intermediate optical amplifiers. A similar solution for the quantum scenario is prohibited by the no-cloning theorem [124]. Fortunately, while quantum mechanics rules out the amplification of quantum signals, it allows for a process called *entanglement swapping*. This

phenomenon is exploited [1] by so-called quantum repeaters [20, 111, 146], which work by first establishing entanglement over shorter distances. Performing entanglement swaps between these states results in entanglement over larger distances. The individual creation of the shorter entangled links are all independent, and do not have to happen simultaneously. Thus, by including a large number of quantum repeaters that can perform entanglement swaps, the effects of losses can be reduced. This shows that high-rate quantum communication is possible in the idealised scenario of noiseless apparatuses.

In practice, including more quantum repeaters will increase the noise. **It is important to understand in which experimental circumstances quantum repeaters will prove to be beneficial.** In fact, the current experimental frontier has seen only recently (non-scalable) implementations of single repeater nodes that can be claimed to improve over direct transmission for those distances [94, 108, 170].

Naively, it would thus appear that quantum communication is thus still restricted to limited distances in the presence of realistic noise. Luckily, quantum mechanics allows for *entanglement distillation*. This is a process where multiple entangled states are (possibly probabilistically) converted into a smaller number of more strongly entangled states [9, 11, 83]. The number of possible ways of distributing entanglement with multiple quantum repeaters grows enormously fast, especially when taking into account entanglement distillation. **How should one then distribute entanglement using multiple (near-term) quantum repeaters, and what are the resultant fidelities that can be achieved?**

As mentioned in the above paragraph, entanglement distillation is one key way noise can be combated in realistic quantum networks. This has motivated the theoretical study of entanglement distillation protocols with a small number of copies [38, 141, 181]. **However, even for a small number of copies, the possible operations one can perform increases significantly, rendering the analysis of entanglement distillation protocols difficult.** Improved distillation protocols would allow for pushing the capabilities of near-term quantum networks even further.

In this thesis we explore how one can improve in the near-term on some of the fundamental building blocks of quantum networks. In particular, we aim to solve (in part) the problems indicated in bold above.

[1] Here and in the remainder of the thesis, the term *quantum repeaters* refers to *first generation quantum repeaters* [113], unless stated otherwise.

1.1. THESIS OVERVIEW

This thesis contains four main chapters. In chapters 2 and 3 we investigate proof-of-principle quantum repeater experiments. We analyse the performance of several setups and investigate parameter regimes to find when an implemented device could be claimed to provide a benefit over direct transmission. In chapter 4, we examine schemes for entanglement distribution over quantum repeater chains. We utilise a heuristic optimisation to find better schemes than previously known and investigate the relevant parameters for long-distance entanglement distribution. Finally, we explore further entanglement distillation protocols in chapter 5. We provide a systematic way to find and optimise over distillation protocols belonging to an experimentally relevant class, allowing us to improve on previously known distillation protocols. Below we conclude with a more detailed summary of the remainder of the thesis.

Chapter 2: In this chapter we provide a fine-grained analysis of a specific scheme of a so-called single quantum repeater node. We are primarily interested in when a quantum repeater outperforms direct transmission, in particular for the task of quantum key distribution. This is motivated by quantum key distribution being one of the most mature quantum technologies. Thus, the metric we use for the performance of the single sequental quantum repeater is the asymptotic secret-key generated per attempt. We compare this quantity with several information-theoretic bounds on the achievable secret-key rate over quantum channels modelling direct transmission. For the modelling of direct transmission, we consider different *benchmarks*, depending on whether the input energy is constrained and whether thermal noise or additional losses are included. We use two methods for increasing the secret-key rate with the repeater. First, we introduce the *cut-off*, which allows to make a trade-off between the secret-key fraction and the generation rate. The cut-off imposes only a maximum storage time on the state to be stored, and is thus experimentally easy to implement. Secondly, we use advantage distillation for the classical post-processing to increase the generated secret-key rate. We perform an analysis for which parameters it is possible to beat each of the abovementioned benchmarks. We use this to find the important parameters to improve on for claiming a proper quantum repeater.

Chapter 3: In chapter 2 the analysis presented was tailored to a specific quantum repeater scheme. In this chapter, we consider four quantum repeater schemes, and assess their performance for generating secret key when implemented on a nitrogen-vacancy centre in diamond setup. The *single-photon scheme* performs the best out of all four schemes with near-term parameters. In fact, the single-photon scheme surpasses the capacity (which is the most stringent of the benchmarks introduced in chapter 2) by a factor of 7 with near-term parameters. Surprisingly, this scheme does not require any storage, highlighting the experimental feasibility of such a scheme.

Chapter 4: In the previous two chapters we were concerned with a single quantum repeater node. Multiple of these repeaters can be linked up together to form a quantum repeater chain. Each of the nodes has the potential to perform elementary link generation, Bell state measurements and distillation. However, the number of possible schemes that can be performed on such a repeater chain grows super-exponentially with the number of involved nodes. This complicates the optimisation over schemes for repeater chains. We provide an algorithm that can efficiently perform a heuristic optimisation over quantum repeater schemes. We find that our algorithm finds significant improvements in the generation rate in comparison to an optimisation over a simpler class of repeater schemes based on BDCZ repeater schemes. Furthermore, we apply our algorithm to three different implementations for an investigation of the important parameters to improve.

Chapter 5: In this chapter we study entanglement distillation protocols. We consider a class of experimentally relevant protocols, which are built from so-called *bilocal Clifford gates* and require only a single round of communication. We provide a full classification of such protocols for up to five arbitrary Bell-diagonal states. Furthermore, we consider the case of distilling an n-fold tensor product of a Werner state. By exploiting the symmetry of such an input state, we classify and subsequently optimise over all such protocols for up to 8 Werner states. We provide explicit circuits with modest depth and number of two-qubit gates that achieve the highest fidelity for the Werner case.

Chapter 6: We conclude with a summary of the results found in this thesis, and possible directions for future research.

2

PARAMETER REGIMES FOR A SINGLE SEQUENTIAL QUANTUM REPEATER

Filip Rozpędek*, Kenneth Goodenough*, Jérémy Ribeiro, Norbert Kalb, Valentina Caprara Vivoli, Andreas Reiserer, Ronald Hanson, Stephanie Wehner and David Elkouss

Quantum key distribution allows for the generation of a secret key between distant parties connected by a quantum channel such as optical fibre or free space. Unfortunately, the rate of generation of a secret key by direct transmission is fundamentally limited by the distance. This limit can be overcome by the implementation of so-called quantum repeaters. Here, we assess the performance of a specific but very natural setup called a single sequential repeater for quantum key distribution. We offer a fine-grained assessment of the repeater by introducing a series of benchmarks. *The benchmarks, which should be surpassed to claim a working repeater, are based on finite-energy considerations, thermal noise and the losses in the setup. In order to boost the performance of the studied repeaters we introduce two methods. The first one corresponds to the concept of a* cut-off, *which reduces the effect of decoherence during storage of a quantum state by introducing a maximum storage time. Secondly, we supplement the standard classical post-processing with an* advantage distillation *procedure. Using these methods, we find realistic parameters for which it is possible to achieve rates greater than each of the benchmarks, guiding the way towards implementing quantum repeaters.*

---

*These authors contributed equally. K. Goodenough contributed with the analysis of the cut-off and the benchmarks, implementing the code and co-writing the manuscript.

This chapter has been adapted from the following publication: Quantum Sci. Technol. 3, 034002 (2018)

2.1. INTRODUCTION

In this first chapter we evaluate a realistic setup of a so-called single sequential quantum repeater, one of the (possible) building blocks for the quantum internet. The general concept of a (not necessarily single sequential) quantum repeater was proposed in [20]. The authors of this scheme showed that by dividing the entire communication distance into smaller segments, generating entanglement over those short links and performing entanglement swapping operation at each of the intermediate nodes in a nested way, one can establish long-distance entanglement. It was also shown that by including the procedure of entanglement distillation, one can furthermore overcome the problem of noise.

Unfortunately, this model does not go into detail of how the physical imperfections of realistic devices, such as decoherence of the quantum memories with time or possibly the probabilistic nature of entanglement swapping, affect the performance. These observations have led to the development of significantly more detailed and accurate, but at the same time significantly more complex, repeater schemes [4, 45, 78, 110, 112]. Many quantum repeater proposals require significant resources and are thus not within experimental reach. However, the recent experimental progress in the development of quantum memories [101, 135, 152] has brought the realisation of a quantum repeater closer than ever.

This motivates us to study a quantum repeater setup that is close to experimental realisation. The setup that we will investigate here was originally proposed in [99], where the authors were inspired by the memory-assisted measurement-device-independent QKD setup (MA-MDI QKD) [121]. Alice and Bob use a single sequential quantum repeater located between them, where both of them are connected to the quantum repeater by optical fibre. The repeater is composed of two quantum memories, both of which have the ability to become entangled with a photon, see FIG. 2.1. However, the repeater has a single photonic interface, which means that it can only address Alice and Bob in a sequential fashion. Examples where only one of the qubit memories has an interface to the photonic channel include modular ion traps [71] and nitrogen-vacancy centres in diamond [14, 55, 135]. The situation is similar for atoms or ions trapped in a single cavity [136]. In this case, both memories can have a photonic interface. However, typically only one of the interfaces can be active at a given moment.

The figure of merit that we have chosen to evaluate the repeater is the secret-key rate. That is, the ratio between the number of generated secret bits and the number of uses of the quantum channel connecting the two parties. The secret-key rate is a very natural quantifier of the performance of the studied scheme for the task of the secret key generation. It depends both on the success rate of the protocol as well as on the quality of the transmission. We compare the secret-key rate achievable with the repeater with a set of benchmarks that we introduce here. The most strict of these benchmarks is the capacity of the channel [175]. That is, the optimal secret-key rate achievable over optical fibre unassisted by a quantum repeater [131]. The other benchmarks correspond to the optimal rates achievable with additional restrictions. In consequence, these benchmarks form a set of stepping stones towards the first quantum repeater able to produce a secure key over large distances.

The idea of assessing quantum repeaters by comparing with the optimal unassisted rates [6, 32, 59, 130, 131, 155, 176, 176] has spurred a significant amount of research devoted to developing sophisticated repeater proposals. Analysis of practical systems that utilise only parametric down-conversion sources and optical measurement setups [85] has shown that such systems do not allow for overcoming the channel capacity, which hints at the importance of quantum memories in repeater architectures. Specific architectures that utilise entangled-photon pair sources together with multimode quantum memories have also been considered in this context [63, 89]. Their analysis suggests that the required efficiency of those entangled-photon pair sources and number of storage modes might be experimentally very challenging for implementation in the very near fu-

ture. Finally, the so called all-optical repeaters that do not require quantum memories but allow to overcome the channel capacity have been proposed [122]. However, they necessitate the ability to create large photonic cluster states which are beyond current experimental capabilities.

A detailed analysis of a realistic, single-node proof-of-principle repeater that includes all the specific system imperfections has been recently performed [99]. In particular, the analysis identified parameter regimes where it would be possible to surpass the optimal direct transmission rates with a repeater scheme that is close to experimental implementation. We build upon the analysis of [99] by introducing two methods that allow us to achieve higher rates. The first of these methods is the introduction of a maximum storage time for the memories in the quantum repeater. This restriction effectively reduces the effect of decoherence. We derive tight analytical bounds for the secret-key rate as a function of the maximum storage time. In this way we can perform efficient optimisation of the secret-key rate over the maximum storage time. The second of these methods is advantage distillation [61], a two-way classical post-processing technique that allows for distilling secret key at a higher rate than achievable with only one-way post-processing. We note here that our analysis here — similar to the one from [99] — is on the implementation of a system without any further noise incurred by any potential eavesdropper, i.e. a so-called *simulation scenario*. Clearly this does not affect the security of the implementation of the setup for quantum key distribution.

This chapter is structured as follows. In Section 2.2 we detail our key distribution protocol. The sources of errors, such as losses in the apparatus and noisy operations and storage, are discussed in Section 2.3. In Section 2.4, we calculate the secret-key rate that the single sequential quantum repeater would achieve. We define the benchmarks in Section 2.5, in Section 2.6 we discuss a specific implementation for NV-centres, and in Section 2.7 we numerically explore the parameter regimes for which the quantum repeater implementation overcomes each benchmark and determine how the secret-key rate of the proposed protocol scales as a function of the distance. We end in Section 2.8 with some concluding remarks.

2.2. A SINGLE SEQUENTIAL QUANTUM REPEATER PROTOCOL

A quantum key distribution protocol consists of two main parts. First, Alice and Bob exchange quantum signals over a quantum channel and measure them to obtain a raw key that is post-processed in a second, purely classical part into a secure key [148]. Here, we focus our interest on the entanglement-based version of the BB84 [10] and the six-state [21] protocols. In this section, we describe the first part of both key distribution protocols.

The physical setup consists of two spatially separated parties Alice and Bob connected to an intermediate repeater via optical fibre channels. We note that such a repeater does not need to be positioned exactly half-way between Alice and Bob. The repeater is composed of two qubit quantum memories which we denote by QM_1 and QM_2. The repeater is then able to generate memory-photon entanglement, where the photonic degree of freedom in which the qubits are encoded is assumed to be time-bin. Alice and Bob each have an optical detector setup that performs a BB84 or a six-state measurement. For technical reasons (see Section 2.3), we consider slightly different setups for BB84 and six-state. More concretely, for BB84 we consider an active setup that switches randomly between the two measurement bases, while in the six-state protocol we consider a passive setup that chooses between the three measurement bases by a passive optical construction [57].

Let us now describe a first version of the protocol without a maximum storage time. First, the quantum repeater attempts to generate an entangled qubit-qubit state between a photon and the first quantum memory QM_1, after which the photon is sent through a fibre to Alice. Such a *trial*

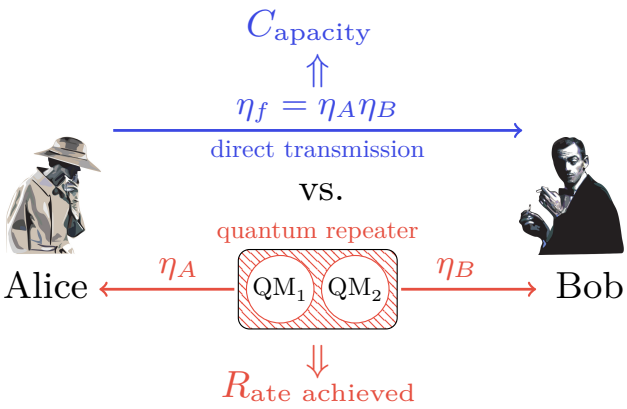

Figure 2.1: The quantum repeater will send photons entangled with the QM_1 to Alice through the optical fibre of transmissivity η_A. After receiving one photon she will perform a BB84 or six-state measurement. After Alice has measured a photon and communicated her success to the quantum repeater, the quantum repeater tries to send a photon entangled with the QM_2 to Bob through the optical fibre of transmissivity η_B. If Bob does not receive a photon within some pre-defined amount of trials (i.e. the cut-off), Alice and Bob will abort the round. This is done to prevent the state in the QM_1 from decohering excessively. If Bob does succeed, the quantum repeater performs a Bell state measurement on the two quantum memories.

is attempted repeatedly until a photon arrives at Alice's side, after which Alice performs either a BB84 or a six-state measurement. Second, the quantum repeater attempts to do the same on Bob's side with the second quantum memory QM_2 while the state in QM_1 is kept stored.

We denote the number of trials performed until a photon arrives at Alice's and Bob's sides n_A and n_B respectively. After Bob has received and measured a photon, a Bell state measurement is performed on the two states in QM_1 and QM_2. We denote by p_{bsm} the probability that the measurement succeeds. The classical outcome of the Bell state measurement is communicated to Bob. This concludes a single *round* of the protocol. We note that in this protocol every round ends with a successful generation of one bit of raw key. Such a protocol is closely related to the memory-assisted measurement-device-independent QKD setup (MA-MDI QKD) [121]. We discuss this connection in Appendix 2.9.3.

One of the main problems in a quantum repeater implementation is that a quantum state will decohere when it is stored in a quantum memory. This means that if it takes Bob a large amount of trials to receive a photon, the state in the quantum memory QM_1 will have significantly decohered, preventing the generation of secret key. This motivates the introduction of a *cut-off*. A cut-off is a limit on the amount of trials that Bob can attempt to receive a photon. We denote this maximum number by $n^\star$.
The protocol that we consider here modifies the protocol above as follows: if in a given round Bob reaches the cut-off without success, the round is interrupted and a new round starts from the beginning with the quantum repeater again attempting to send a photon to Alice. In this scheme a large number of rounds might be required until a single bit of raw key is successfully generated. See Algorithm 1 for a description of the modified protocol with the cut-off.

Algorithm 1: Generation of a bit of raw key with a single sequential quantum repeater

$k \leftarrow 0$;
Loop
  $n_A \leftarrow 0,\ n_B \leftarrow 0$;
  do
    $k \leftarrow k+1$ ▷ Increment the number of rounds;
    $n_A \leftarrow n_A + 1$ ▷ Increment the number of Alice's channel uses;
    Generate entangled photon-QM_1 pair ;
    Send entangled photon through fibre towards Alice
  while *Alice has not received photon*;
  Alice performs a BB84 or a six-state measurement, stores result;
  do
    $n_B \leftarrow n_B + 1$ ▷ Increment the number of Bob's channel uses;
    Generate entangled photon-QM_2 pair ;
    Send entangled photon through fibre towards Bob;
    if *Bob received photon* **then**
      Bob performs a BB84 or a six-state measurement, stores result;
      Perform the Bell state measurement on the memories;
      Communicate result;
      Store $\max(n_A, n_B)$ ▷ Store channel uses;
    end if
  while *Bob has not received photon and* $n_B < n^\star$;
EndLoop

2.3. SOURCES OF ERRORS

In this section, we model the different elements in the setup to identify the sources of losses and noise. The losses in the system are not only due to the transmissivity of the fibre; depending on the implementation a significant amount of photons is lost before they enter the fibre or due to the non-unit detector efficiency. The causes of noise are the experimental imperfections of the operations, measurements and quantum memories.

LOSSES

We model the process of generating and sending an entangled photon through a fibre as follows (see FIG. 2.2). First, the photon has to be generated at some photon source and be captured in the fibre. This process happens with probability p_{em}. Depending on the experimental implementation, only a fraction p_{ps} of the photons entering the fibre can be used for secret key generation. This can occur for any number of reasons, for instance photons might be filtered according to frequency or a certain time-window [55, 136]. The filtering can happen either before or after the transmission through the fibre. The fibre losses are modelled as an exponential decay of the transmissivity η_f with the distance L, i.e. $\eta_f = \exp\left(-\frac{L}{L_0}\right)$ for some fibre attenuation length L_0. We denote by η_A the fibre losses on Alice's side and by η_B the fibre losses on Bob's side. Finally, the arriving photons will be captured by the detectors with an efficiency p_{det}. This probability of detecting a photon will be increased by the presence of dark counts (which will also inevitably add noise to the system), see the discussion of the dark counts at the bottom of this section and in Appendix 2.9.1. We define the quantity $p_{\text{app}} = p_{\text{em}} p_{\text{det}}$ describing the total efficiency of our

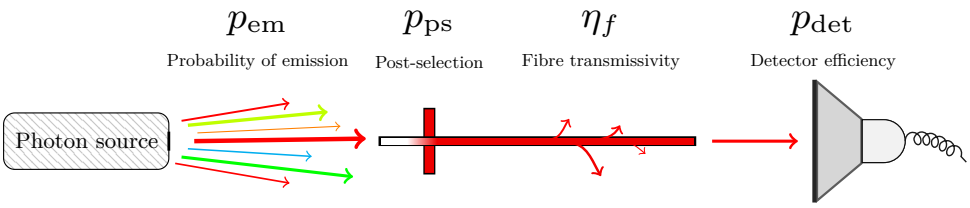

Figure 2.2: General model of all photon losses occurring in the repeater setup. p_{em} is the probability of generating and capturing a photon into the fibre. For experimental reasons a fraction $(1-p_{\text{ps}})$ of photons are additionally filtered out. The fibre has a transmissivity η_f. After exiting the fibre, the photons produce a click in the detector with probability p_{det}. The total efficiency of the apparatus is described by one parameter, $p_{\text{app}} = p_{\text{em}} p_{\text{det}}$.

apparatus.

NOISE

We model all noise processes either by the action of a dephasing channel

$$\mathscr{D}^{\lambda_1}_{\text{dephase}}(\rho) = \lambda_1 \rho + (1-\lambda_1) Z\rho Z \tag{2.1}$$

or that of a depolarising channel

$$\mathscr{D}^{\lambda_2}_{\text{depol}}(\rho) = \lambda_2 \rho + (1-\lambda_2) \frac{\mathbb{I}}{2} \tag{2.2}$$

where the parameters λ_1 and λ_2 quantify the noise, Z is the qubit gate $\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ and $\mathbb{I}/2$ is the maximally mixed state. The noise processes occur due to imperfect operations, decoherence of the state while stored in QM_1 and dark counts in the detectors.

The noise from imperfect quantum operations is captured by two parameters: F_{prep} and F_{gm}. F_{prep} is a dephasing parameter which corresponds to the preparation fidelity of the memory-photon entangled state [162]. F_{gm} is a depolarising parameter that describes the noise introduced by the imperfect gates and measurements performed on the two quantum memories during the protocol [35, 83]. Hence, the noise can be modelled by a dephasing and a depolarising channel with $\lambda_1 = F_{\text{prep}}$ and $\lambda_2 = F_{\text{gm}}$.

Besides the dephasing and depolarising noise with $\lambda_1 = F_{\text{prep}}$ and $\lambda_2 = F_{\text{gm}}$. , there is also decoherence over time. The decoherence is modelled by a decay of the fidelity in the number of trials n. This decoherence is caused by two distinct effects. Firstly, there is the decoherence due to the time that the quantum repeater has to wait between sending photons. This time is the time it takes to confirm whether the photon got lost plus the time it takes to generate a photon entangled with the memory. We model this effect through an exponential decay of fidelity with time [115], which is expected whenever excess dephasing is suppressed (e.g. by dynamical decoupling [36]). However, we note that this is not the only possible model of decay, in several experiments a Gaussian decay has been observed [71, 146, 152, 160]. Secondly, attempting to generate an entangled photon-memory pair at QM_2 might also decohere the state stored in the QM_1. For example, this effect is the most prominent decoherence mechanism in nitrogen-vacancy implementations [135], where an exponential decay of fidelity with the number of trials was observed. This is also how we model that effect here.

The quantum state ρ that is subjected to those effects undergoes an evolution given by the dephasing and depolarising channels with $\lambda_1 = (1+e^{-an})/2$ and $\lambda_2 = e^{-bn}$. The two parameters a and b are given by

$$a = a_0 + a_1\left(\frac{2n_{\text{ri}}L_B}{c} + t_{\text{prep}}\right), \tag{2.3}$$

$$b = b_0 + b_1\left(\frac{2n_{\text{ri}}L_B}{c} + t_{\text{prep}}\right), \tag{2.4}$$

where n_{ri} is the refractive index of the fibre, c is the speed of light in vacuum, L_B the distance from the quantum repeater to Bob and t_{prep} is the time it takes to prepare for the emission of an entangled photon. Here a_0 and b_0 quantify the noise due to a single attempt at generating an entangled state and a_1 and b_1 quantify the noise during storage per second. Finally, the dark counts in the detectors introduce depolarising noise. This model is justified for the two quantum key distribution protocols that we consider, see [8, 57]. We let $\alpha_{A/B}$ denote the corresponding depolarising parameter on Alice's/Bob's side. The details of this model are presented in Appendix 2.9.1.

2.4. SECRET-KEY RATE OF THE SETUP

The performance of a setup is assessed in this chapter by its ability to generate secret key between two parties, Alice and Bob. We note here that the ability of a quantum repeater to generate secret key can be measured in two different ways - in its throughput and its secret-key rate. The throughput is equal to the amount of secret key generated per unit of time, while the secret-key rate equals the amount of secret key generated per *channel use*. Here and in the following chapter, we will focus on the secret-key rate only. This is due to the fact that it allows us to make concrete information-theoretical statements about our ability to generate secret key. Moreover, we note that the secret-key rate is also more universal in the sense that it can be easily converted into the throughput by multiplying it with the repetition rate of our scheme (number of attempts we can perform in a unit of time). It must be also noted here that demonstrating repeater schemes that achieve higher throughput than the currently available QKD systems based on direct transmission will be a great challenge. This is because the sources of photonic states used within those QKD systems operate at the GHz repetition rates, while the performance of the repeater schemes will be limited by many additional factors such as transmission latency and time of local operations at the memory nodes. These issues are not captured by the secret-key rate directly. Nevertheless, as mentioned before, the universality of the secret-key rate allows for the interconversion between the two quantities. We further discuss the differences between the throughput and secret-key rate in Section 3.6.5 of the next chapter.

The secret-key rate R is defined as the amount of secret-key bits generated by a protocol divided by the number of channel uses and the number of optical modes. In the particular case of our sequential quantum repeater, the secret-key rate is given by

$$R = \frac{Y}{2}r\,. \tag{2.5}$$

The yield Y of the protocol is defined as the rate of raw bits per channel use. The secret-key fraction r is defined as the average amount of secret key that can be extracted from a single raw bit. The (conservative) factor of a half is due to the fact that the encoding uses two optical modes — in principle it would be possible to use those two modes to asymptotically generate $-2\log_2(1-\eta)$ key bits per two modes.

Since we consider two possible quantum key distribution protocols we take

$$r = \max\{r_{\text{BB84}}, r_{\text{six-state}}\}\,. \tag{2.6}$$

where r_{BB84} and $r_{\text{six-state}}$ are the secret-key fractions of the BB84 and six-state protocols, respectively (see Eq. (2.12) and Appendix 2.9.4).

YIELD

The yield can be calculated as p_{bsm} (i.e. the success probability of the Bell state measurement) divided by the (average) number of channel uses needed for the successful detection of a photon by both Alice and Bob in the same round. With a single sequential quantum repeater it is not obvious how to count the number of channel uses. As in [99], we count the *maximum* of the two channel uses on Alice's and Bob's sides respectively in this chapter,

$$Y = \frac{p_{\text{bsm}}}{\mathbb{E}[N]} = \frac{p_{\text{bsm}}}{\mathbb{E}[\max(N_A, N_B)]} . \tag{2.7}$$

where N, N_A and N_B are the random variables that model the number of channel uses, the number of channel uses at Alice's side and the number of channel uses at Bob's side, respectively. We note here that for a sequential repeater, the sum of the channel uses better captures the time spent occupying the channel. In this chapter, we use the maximum to match the approach from [99], but we consider the more conservative sum of N_A and N_B in the next chapter. We note here that the qualitative results from this and the following chapter do not strongly depend on which definition of the number of channel uses is used.

Without the cut-off, it is possible to obtain an analytical formula for the average number of channel uses [99, 121],

$$\mathbb{E}[\max(N_A, N_B)] = \frac{1}{p_A} + \frac{1}{p_B} - \frac{1}{p_A + p_B - p_A p_B} , \tag{2.8}$$

where p_A and p_B depend on the quantum key distribution protocol and are given by the following equations (see Appendix 2.9.1),

$$p_{A/B,\text{BB84}} = 1 - (1 - p_{\text{app}} p_{\text{ps}} \eta_{A/B})(1 - p_d)^2 , \tag{2.9}$$

$$p_{A/B,\text{six-state}} = 1 - (1 - p_{\text{app}} p_{\text{ps}} \eta_{A/B})(1 - p_d)^6 . \tag{2.10}$$

Here p_d is the probability of measuring a dark count.

Every time that Bob reaches $n^\star$ trials, Alice and Bob restart the round and start over again. The cut-off thus increases the average number of channel uses. We have developed an analytic approximation of $\mathbb{E}[N]$ which is essentially tight (see Appendix 2.9.5 for the derivation and error bounds)

$$\mathbb{E}[\max(N_A, N_B)] \approx \begin{cases} \frac{1}{p_A\left(1-(1-p_B)^{n^\star}\right)} & \frac{1}{p_A} > n^\star \\ \frac{1}{p_A} + \frac{1}{p_B} - \frac{1}{p_A + p_B - p_A p_B} & \frac{1}{p_A} \leq n^\star . \end{cases} \tag{2.11}$$

SECRET-KEY FRACTION

Here we consider the secret-key fraction of the BB84 and six-state protocols. As we discussed previously, we consider the BB84 protocol with an active measuring scheme and the six-state protocol with a passive one. Moreover, we consider a fully asymmetric version of BB84 and a fully symmetric version of six-state. Fully symmetric means that all bases are used with equal probability while fully asymmetric means that the ratio at which one of the bases is used is arbitrarily close to one. Finally, we consider a one-way key distillation scheme for BB84 [148] while for the six-state protocol we consider the advantage distillation scheme in [171]. Advantage distillation [61] is a classical post-processing technique that allows to increase the secret-key fraction at all levels of noise.

The reasons for not analysing the BB84 protocol with advantage distillation and the fully asymmetric six-state with advantage distillation are technical. In the case of BB84, computing the rate with advantage distillation requires the optimisation over a free parameter. The combination of the optimisation over the cut-off together with the extra free parameter was computationally too intensive to consider here.

For the six-state protocol there is, to our knowledge, no security proof that can deal with the asymmetric six-state protocol with photonic qubits without introducing extra noise [5, 57]. However, these protocol choices do not have a strong impact on our analysis. Advantage distillation does not significantly increase the amount of distillable key for low error rates. Hence, asymmetric BB84 without advantage distillation is only slightly suboptimal. For higher error rates, where advantage distillation plays a role, the symmetric six-state protocol with advantage distillation is a factor of three away from the asymmetric version.

The expression for the secret-key fraction of both protocols depends on the error rates in the X, Y and Z bases, which we denote by e_X, e_Y and e_Z. In the case of the BB84 protocol, [95, 148] it is given by

$$r_{\text{BB84}} = 1 - h(e_Z) - h(e_X)\,, \tag{2.12}$$

where $h(p) = -p\log_2 p - (1-p)\log_2(1-p)$ is the binary entropy function. The expression for $r_{\text{six-state}}$ is more complex; we leave its discussion to Appendix 2.9.4.

We can directly evaluate the error rates in each basis as a function of the general parameters of Section 2.3. For the single sequential quantum repeater these average errors are

$$e_X = e_Y = e_{XY} = \frac{1}{2} - \frac{1}{2} F_{\text{gm}} \alpha_A \alpha_B \left(2F_{\text{prep}} - 1\right)^2 \left\langle e^{-(a+b)n} \right\rangle\,, \tag{2.13}$$

$$e_Z = \frac{1}{2} - \frac{1}{2} F_{\text{gm}} \alpha_A \alpha_B \langle e^{-bn} \rangle\,. \tag{2.14}$$

where $\langle e^{-cn} \rangle$ is the average of the exponential e^{-cn} over a geometric distribution over the first $n^\star$ trials. The detailed derivation of the error expressions is presented in Appendix 2.9.2.

2.5. BENCHMARKS FOR ASSESSING QUANTUM REPEATERS

We introduce a set of benchmarks to assess the performance of a quantum repeater implementation.

The first benchmark that we consider is the rate that would be achieved with the same parameters for the system losses and dark counts and for the same protocol but without a quantum repeater. Overcoming this benchmark gives the first indication that the repeater setup is useful; it means that the repeater setup outperforms the setup without repeater. We call this benchmark the direct transmission benchmark.

The remaining benchmarks represent the optimal secret-key rate that Alice and Bob could achieve if they were to communicate over the same quantum channel without a repeater under some constraints.

The optimal secret-key rate without a repeater highly depends on the channel model. The first modelling decision is the placement of the boundary between Alice's and Bob's laboratories and the quantum channel. This is because it is not *a priori* clear where the channel begins and ends. However, this decision has a strong impact on the optimal achievable rate; if the channel includes most of Alice's and Bob's laboratories, then the channel is more lossy and noisy and the benchmark is easier to overcome. If, on the other hand, the channel is just the optical fibre cable the benchmark becomes more difficult to overcome.

We consider three cases in terms of the individual lossy components of our setup (see FIG. 2.1, FIG. 2.2 and their captions):

Case 1: Fibre only, in this case the transmissivity is: $\eta = \eta_f = \eta_A \eta_B$.

Case 2: Fibre and different filters, then the channel transmissivity becomes: $\eta = \eta_f p_{ps}$.

Case 3: Fibre, filters and Alice's and Bob's apparatus, then the transmissivity becomes: $\eta = \eta_f p_{ps} p_{app}$.

Note that although in the experimental implementation of the repeater the terms p_{ps} and p_{app} appear twice in the expression of the transmissivity, they appear only once in the benchmarks which include them. The reason is that in a scenario without a repeater the emission inefficiency and the filters only affect the transmissivity once.

The second design parameter for these benchmarks is the type of channel. Transmission of photons through fibres is modelled as a pure-loss channel [172], where only a fraction η of the input photons reach the end of the channel. The first type of channel that we consider is the pure-loss channel without any additional restriction. The optimal achievable rate over one mode of the pure-loss channel is given by the secret-key capacity [131]

$$-\log_2(1-\eta)\ . \tag{2.15}$$

Note that for high losses the scaling of this capacity with distance is proportional to $\eta_f = \exp\left(-\frac{L}{L_0}\right)$. At the same time with an ideal (noiseless) single quantum repeater placed half-way between Alice and Bob, the expected secret-key rate would scale proportionally to $\sqrt{\eta_f} = \exp\left(-\frac{L}{2L_0}\right)$ [99].

The second type of channel that we consider is the pure-loss channel when the transmitter has a limitation in the energy that can be introduced into the channel. There has been some recent work studying the optimal rate per mode of the finite-energy pure-loss channel [59, 156, 176]. However, the optimal rate remains unknown. The bound that we consider here [156] is given by

$$g\left((1+\eta)P/2\right) - g\left((1-\eta)P/2\right)\ , \tag{2.16}$$

where $g(x) := (x+1)\log_2(x+1) - x\log_2 x$ and P is the mean photon number. In our repeater setup, the finite energy restriction arises from the fact that, on average, only a fraction of a photon enters the fibre in each trial. More precisely, the average photon number satisfies $P = p_{em}$ in cases 1 and 2 above and $P = 1$ in case 3. Unfortunately, since Eq. (2.16) is an upper bound, it is only strictly smaller than the capacity of the pure-loss channel for small mean photon number. Expanding the bounds from equations Eq. (2.15) and Eq. (2.16) around $\eta = 0$ shows that the cross-over between the two bounds occurs when $p_{em}\log_2\left(\frac{p_{em}+2}{p_{em}}\right) = \frac{1}{\ln 2}$. In other words, for high losses the finite-energy bound is tighter when $p_{em} \lesssim 0.796$. This implies that the finite-energy bound does not yield an interesting benchmark in case 3.

The third type of channel that we consider is the thermal-loss channel. An upper bound on the capacity of the thermal-loss channel is

$$-\log_2[(1-\eta)\eta^{\overline{n}}] - g(\overline{n}), \tag{2.17}$$

if $\overline{n} < \frac{\eta}{1-\eta}$ and zero otherwise [131]. Here, $\overline{n}$ is the average number of thermal photons per channel use [172]. This is an interesting channel because the effect of dark counts can be seen as caused by the thermal photons. Hence this type of channel becomes relevant for case 3, where detectors, and therefore also the dark counts, are regarded as part of the channel. The details of the dark count model are presented in Appendix 2.9.1. There we also show how to easily convert the experimentally relevant dark count rate of the detector and the duration of the detection window t_{int} into $\overline{n}$ and p_d, the probability of getting a dark count within the given time window.

The combinations of a channel boundary together with a channel type give us a set of benchmarks. Not all combinations yield interesting benchmarks. In Table 2.1, we summarise the benchmarks that we consider.

Note that beating the benchmark in Eq. 2.15 would be the most convincing, especially since the other benchmarks make (stronger) assumptions on the channel that would be used for direct transmission. That is, for benchmark 3c one would have to assume one compares with direct transmission through a channel corresponding to certain values of the losses and dark counts. Our consideration of the other benchmarks also puts into a practical perspective the purely theoretical work from [59, 131, 156, 176].

	Infinite	Finite	Thermal	Direct transmission
Case 1: η_f	1a	1b	–	–
Case 2: $\eta_f p_{ps}$	2a	2b	–	–
Case 3: $\eta_f p_{ps} p_{app}$	–	–	3c	3d

Table 2.1: Labels of the benchmarks that we use to assess the performance of a quantum repeater. These labels are frequently referred to in the numerical results. Each row corresponds to a different channel boundary, which translates into an effective channel transmissivity. Each column corresponds to a different type of channel: pure loss, pure loss with energy constraint and thermal channel, and the final column corresponds to the direct transmission benchmark.

2.6. IMPLEMENTATION USING NITROGEN-VACANCY CENTRES

Our model is fully general and can be applied to a wide range of physical platforms. To illustrate its performance we will now consider one of such potential near-term realisations of a single sequential quantum repeater. For this particular example we choose to base our system on Nitrogen-Vacancy (NV) centres in diamond. NVs are a prime candidate for this task due to their optical interface featuring high-fidelity single-shot readout [139] and their recently demonstrated capabilities to distribute spin-photon entanglement while faithfully storing quantum states [83].

In the following we expand on the required experimental techniques (see Fig. 2.3). The NV centre itself can be readily used as a generator of spin-photon entanglement at cryogenic temperatures. The NV is encapsulated in an optical cavity of low-mode volume [138] to strongly enhance the emission into the zero phonon line (ZPL) via the Purcell effect. As no particular low-loss cavity design has been implemented with NVs yet, we rely purely on the aforementioned ZPL enhancement. More specific cavity configurations that allow for reflection based mechanisms rely on the realisation of a low-loss overcoupled cavity to be efficient [44] and might become available in the future.

Firstly, we generate spin-photon entanglement [126] and send the emitted photon off to Alice who reports successful detection events back to the repeater station. Note that electron spin decoherence during communication rounds is negligible since second-long coherence times have been demonstrated by employing XY8 dynamical decoupling sequences [1].

Upon success the optical interface of the NV is reused for communication with Bob. To this end, the NV spin state that is correlated with Alice's measurement outcome is stored on a ^{13}C nuclear spin in the vicinity of the electron spin, which itself is then reinitialised. We choose a configuration in which the always-on magnetic hyperfine coupling between both spins is weak (on the order of a few kHz). This configuration has been experimentally shown to result in a highly-addressable quantum memory which is resilient to optical excitation and reinitialisation of the NV spin [135]. Coherently swapping the NV state onto - and high-fidelity control over - such a weakly-coupled nuclear spin has been demonstrated recently [83, 158].

The protocol then proceeds as described in Section 2.2 by communicating with Bob. Note that repeated communication attempts will eventually decohere the memory state due to the necessity for frequent electron spin resets and the always-on hyperfine interaction between the two spins. This constitutes the main source of error in this system (parametrised by a_0 and b_0, see Sec. 2.3).

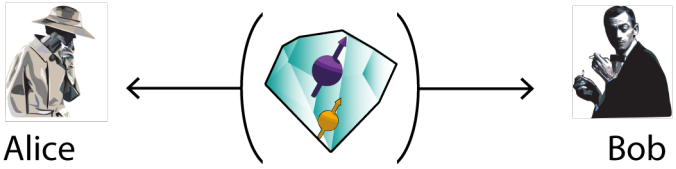

Figure 2.3: Single sequential quantum repeater based on an electron spin associated with an NV (purple) and ^{13}C nuclear spin (orange) in diamond. The previous quantum memories $QM_{1,2}$ are now represented by the electron and nuclear spin respectively. The optical interface of the NV is strongly Purcell-enhanced by an optical cavity with low-mode volume and allows for efficient photon transmission to Alice and Bob.

After a successful state transmission to Bob, we conduct a sequential two-step Bell state measurement and read-out the XX and ZZ parities of the combined nuclear-electron spin state, where X and Z denote the standard Pauli matrices. This can be achieved by means of the earlier mentioned universal control over the system or by introducing additional resource qubits such as the nitrogen nuclear spin associated with the NV [126].

2.7. NUMERICAL RESULTS

In this section, we perform a numerical analysis of our model applied to the physical system based on NV centres as described in Section 2.6. All numerical results have been obtained using a Mathematica notebook [177]. Unless specified otherwise, we use the following parameters that we call "expected parameters" These parameters represent best-case scenarios from the chosen references. These experimental capabilities do not fundamentally contradict or exclude each other and seem therefore achievable in a single experimental NV setup.

- a_0 (dephasing due to interaction) = $\frac{1}{2000}$ per attempt [135],
- a_1 (dephasing with time) = $\frac{1}{3}$ per second [107],
- b_0 (depolarisation due to interaction) = $\frac{1}{5000}$ per attempt [135],
- b_1 (depolarisation with time) = $\frac{1}{3}$ per second [107],
- t_{prep} (memory-photon entanglement preparation time) = 6 μs [65],
- F_{gm} (depolarising parameter for gates and measurements) = 0.9 [83],
- F_{prep} (dephasing parameter for the memory-photon state preparation) = 0.99 [65],
- p_{em} (probability of emission) = 0.49 [16, 65],
- p_{ps} (post-selection) = 0.46 [138],
- p_{det} (detector efficiency) = 0.8 [65],
- p_{bsm} (Bell state measurement success probability) = 1 [126],
- Dark count rate = 10 per second [65],
- t_{int} (detection window) = 30 ns [65],
- L_0 (attenuation length) = 0.542 km [65],
- n_{ri} (refractive index of the fibre) = 1.44 [125].

Before we present the results, we note that the emission frequency of the nitrogen-vacancy centres results in a relatively low L_0 which in turn does not allow to achieve large distances. In

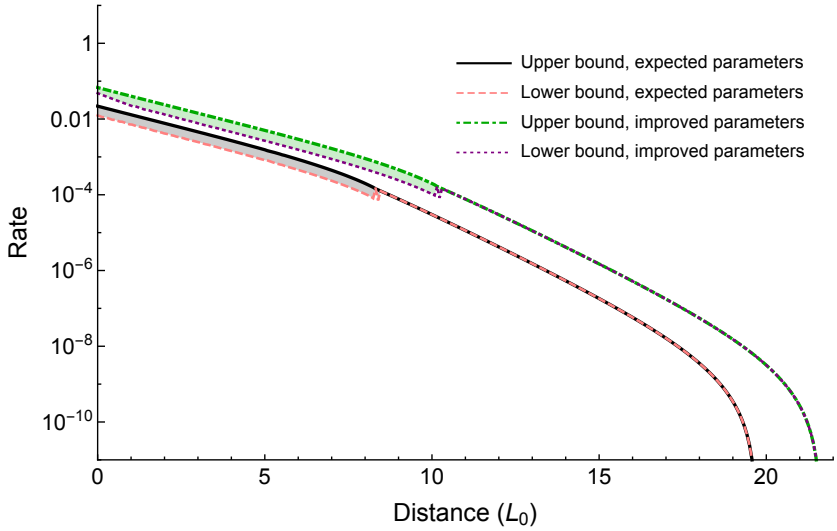

Figure 2.4: Upper- and lower bounds on the secret-key rate rate with a quantum repeater as a function of the distance in units of $L_0 = 0.542$ km. The repeater is positioned half-way between Alice and Bob. The curves correspond to the expected and improved parameters with optimised cut-off. The improved parameters correspond to setting $p_{\text{ps}} = p_{\text{em}} = 0.6$ and $F_{\text{gm}} = 0.97$. For high losses, the upper- and lower bounds become essentially tight. For this reason, the upper bound on the achieved rate forms a reliable estimate of the secret-key rate.

practical quantum key distribution networks, assuming that dedicated fibres are used for which one can choose which frequency mode one wants to transmit at, this problem might be overcome using the frequency conversion of the emitted photons into a telecom frequency, which will yield an increased L_0. Note that the benchmarks in Table 2.1 will scale accordingly.

There is a range of frequencies used in fibre-based communication and for each of those frequencies the attenuation length varies greatly depending on the type of the fibre used. To give some examples, the best fibres at 1560 nm have losses of 0.1419 dB/km ($L_0 \approx 30.6$ km) [159], while at 1310 nm standard single-mode fibres exhibit losses of 0.4 dB/km ($L_0 \approx 10.9$ km) [84].

Clearly our model is general and can be applied to a channel with any value of L_0. Here, throughout most of this section, we consider the transmission through the channel at the same wavelength as the emission line of the NV-centre setup, as such a channel for this specific physical system has been realised in an experiment [65] using fibre with losses of 8 dB/km ($L_0 = 0.542$ km as given in the list of parameters above). At the end we present an additional plot describing the scenario in which a telecom channel with the commonly used in the quantum repeater community attenuation length of $L_0 \approx 22$ km is available. In this case the frequency conversion of the emitted photons to telecom is applied.

Tightness of the error bounds for the secret-key rate. We have derived upper and lower bounds on the yield, and thus also on the secret-key rate, for the two studied protocols. In FIG. 2.4, we plot both the upper and the lower bound on the achieved rate with the current and improved parameters ($p_{\text{ps}} = p_{\text{em}} = 0.6$ and $F_{\text{gm}} = 0.97$) and optimised cut-off as a function of the distance in units of L_0. There are two regimes visible on the plot. This is a consequence of the fact that our bounds have a different analytical form in the two regimes (see Appendix 2.9.5). Since for practical purposes our bounds are essentially tight, from now on we will refer to the upper bound as the

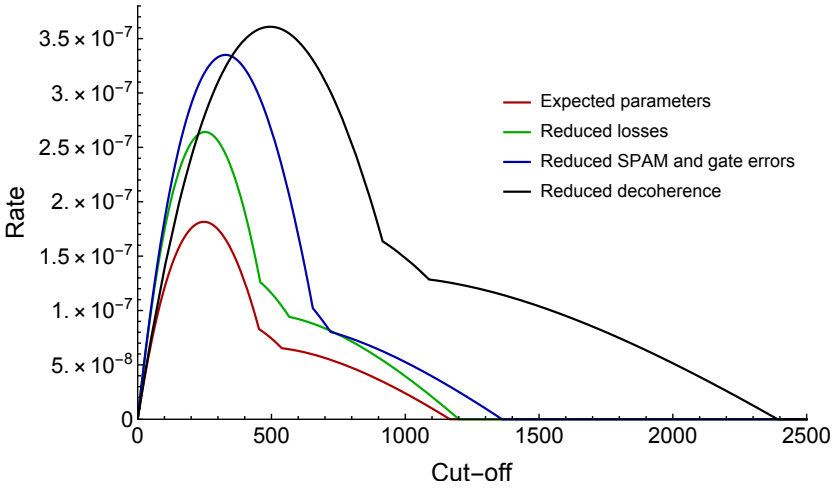

Figure 2.5: Secret-key rate as a function of the cut-off for the expected parameters with the repeater positioned half-way between Alice and Bob. The reduced losses are for $p'_{\text{app}} = (p_{\text{app}})^{0.9}$ and $p'_{\text{ps}} = (p_{\text{ps}})^{0.9}$, the reduced SPAM (state preparation and measurement) and gate errors are for $F'_{\text{gm}} = (F_{\text{gm}})^{0.7}$ and $F'_{\text{prep}} = (F_{\text{prep}})^{0.7}$ and the reduced decoherence is for $a' = a/2$ and $b' = b/2$. The optimal $n^\star$ shifts depending on the parameters. The kinks arise due to the fact that we optimise over two protocols: fully asymmetric BB84 and symmetric six-state protocol with advantage distillation which itself consists of two subprotocols. The optimal protocol depends on the bit error rates. The data have been plotted for the distance of $15L_0$, where $L_0 = 0.542$ km.

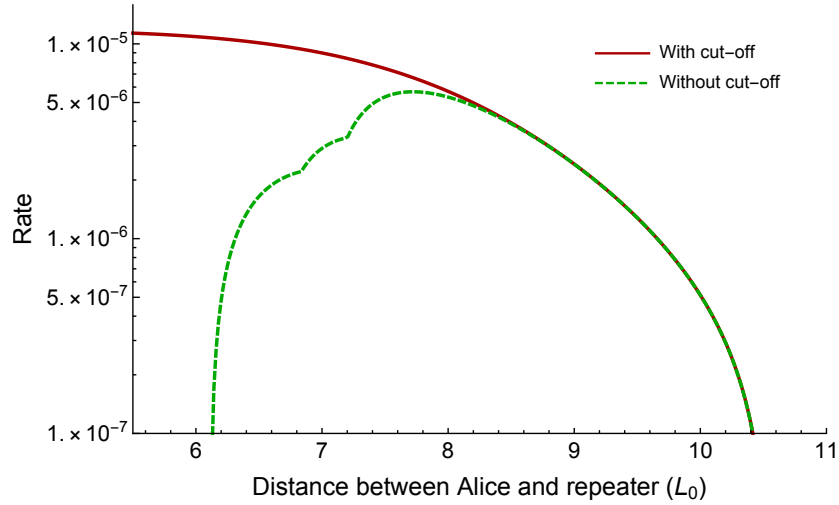

Figure 2.6: Secret-key rate with and without the cut-off as a function of the distance in units of $L_0 = 0.542$ km between Alice and quantum repeater. The total distance between Alice and Bob is fixed to $11L_0$. We see that with the cut-off optimisation, positioning the repeater half-way between Alice and Bob is optimal. This behaviour was also observed for other parameter regimes. This result contrasts with the optimal positioning for the no cut-off scenario, for which we see that shifting the repeater towards Bob is beneficial. We also note that the two rates overlap when the repeater is shifted towards Bob.

expected secret-key rate, and will omit the lower bound for the legibility of the plots.

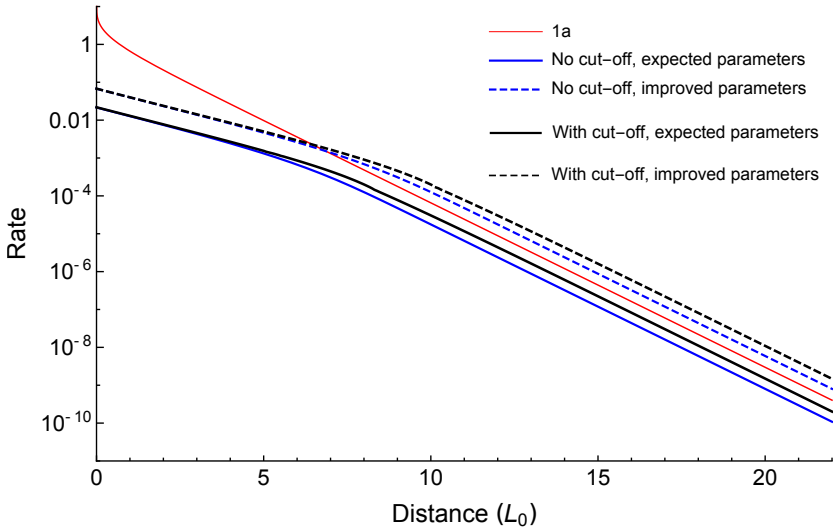

Figure 2.7: Secret-key rate as a function of the distance in units of $L_0 = 0.542$ km, assuming detectors without dark counts. The black lines correspond to the protocol with cut-off and the blue lines to the protocol without the cut-off but with optimised positioning of the repeater. We plot the data for both the expected and improved parameters. The improved parameters correspond to setting $p_{\text{ps}} = p_{\text{em}} = 0.6$ and $F_{\text{gm}} = 0.97$. Finally, the channel capacity (1a) is also included for comparison. It can be seen that both the cut-off and repositioning of the repeater allows to generate key for all distances.

The impact of the cut-off on the secret-key rate. In FIG. 2.5 we plot the secret-key rate versus the cut-off for different sets of parameters. The repeater is assumed to be positioned half-way between Alice and Bob. We observe a strong dependency of the secret-key rate on the cut-off. In particular, for large cut-off the secret-key rate drops to zero. This is due to the inclusion of rounds where the state has significantly decohered. This implies that the cut-off is essential for generating a key at large distances. Moreover, we observe that the optimal cut-off highly depends on the explored parameter regime.

Optimal positioning of the repeater. The asymmetry of the studied sequential protocol raises the question of whether it is best to position the repeater half-way between Alice and Bob. In fact, in the absence of a cut-off this is not the case [99]. For sufficiently large distances, shifting the repeater towards Bob can increase both the secret-key rate and the distance over which the secret-key rate is non-zero in the presence of dark counts. Specifically, the optimal positioning remains a fixed distance away from Bob independently of the actual total distance. Here, we find that with the cut-off and for the parameters considered this phenomenon disappears. We see in FIG. 2.6 that the optimal position with the cut-off optimisation appears to be exactly in the middle of Alice and Bob.

Nevertheless, we note that the bounds for the yield derived in Appendix 2.9.5 are valid under the condition $\eta_B \geq \eta_A$. This means that we can only study the effect of moving the repeater towards Bob. However, we do not expect any benefit in shifting the repeater towards Alice as this could only increase the noise due to decoherence. From now on for the scenarios with the cut-off optimisation, we always consider the repeater to be placed half-way between Alice and Bob. Interestingly, in FIG. 2.6 we also see that the rates for the two scenarios with and without the cut-off

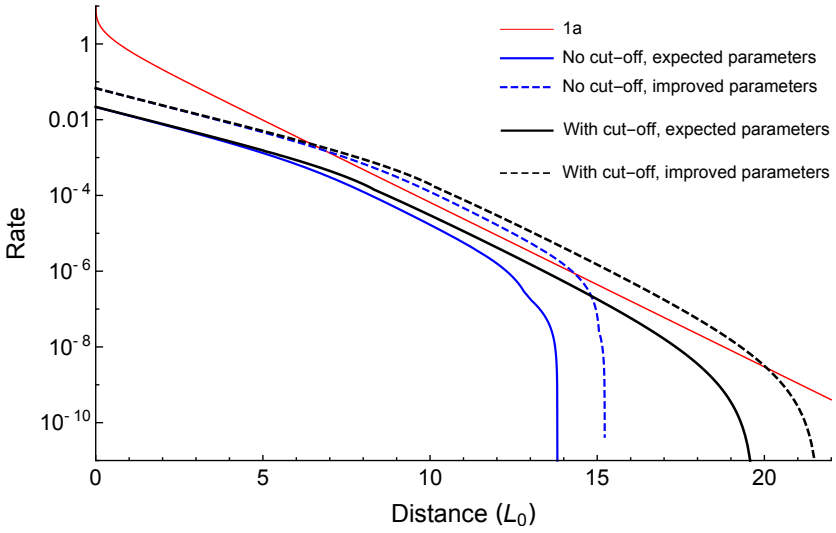

Figure 2.8: Secret-key rate as a function of the distance in units of $L_0 = 0.542$ km with dark counts. The black lines correspond to the protocol with cut-off and the blue lines to the protocol without the cut-off but with optimised positioning of the repeater. We plot the data for both the expected and improved parameters. The improved parameters correspond to setting $p_{\text{ps}} = p_{\text{em}} = 0.6$ and $F_{\text{gm}} = 0.97$. Finally, the channel capacity (1a) is also included for comparison. It can be seen that the protocol with the cut-off is more robust against dark counts than the protocol without the cut-off.

start to coincide after the quantum repeater is shifted within a certain distance of Bob. Intuitively this happens when the probability of Bob getting a photon is large enough so that the significance of the cut-off becomes marginal.

Cut-off versus no cut-off. Having established the optimal positioning of the repeater, we can now compare the two scenarios: optimised cut-off with middle positioning of the repeater and no cut-off with optimised positioning. We find that in the absence of dark counts the scaling with distance of both schemes is the same, with a small advantage of the cut-off scheme. However, the cut-off is more robust against dark counts. Hence, for imperfect detectors the cut-off allows distributing keys at larger distances. These results can be seen in FIG. 2.7 and FIG. 2.8, which show the secret-key rate as a function of distance for detectors without and with dark counts, together with the channel capacity of the optical fibre (i.e. benchmark 1a). We plot the data for the expected and improved parameters ($p_{\text{ps}} = p_{\text{em}} = 0.6$ and $F_{\text{gm}} = 0.97$).

In FIG. 2.7 where we assume no dark counts, we see that for small distances the rate scales approximately with the square root of the transmissivity for both scenarios. That is, they are proportional to the theoretical optimum [99] of $\sqrt{\eta_f} = e^{-L/2L_0}$. For sufficiently large distances time-dependent decoherence of the memory QM_1 becomes a problem. Both schemes overcome it at the expense of reducing the yield. As a result, the scaling becomes proportional to $\eta_f = e^{-L/L_0}$ for both schemes. In FIG. 2.8 however we see that the presence of dark counts affects the two schemes quite differently. While for both schemes the effect of dark counts becomes the dominant source of noise after a certain distance, this distance is shorter for the no cut-off scheme than for the scheme with the cut-off. In other words, we see that the cut-off is more robust towards dark counts than the repositioning method. This fact can be explained by noting that shifting the repeater towards Bob increases the losses on Alice's side and as a result makes the Alice-repeater link vulnerable to

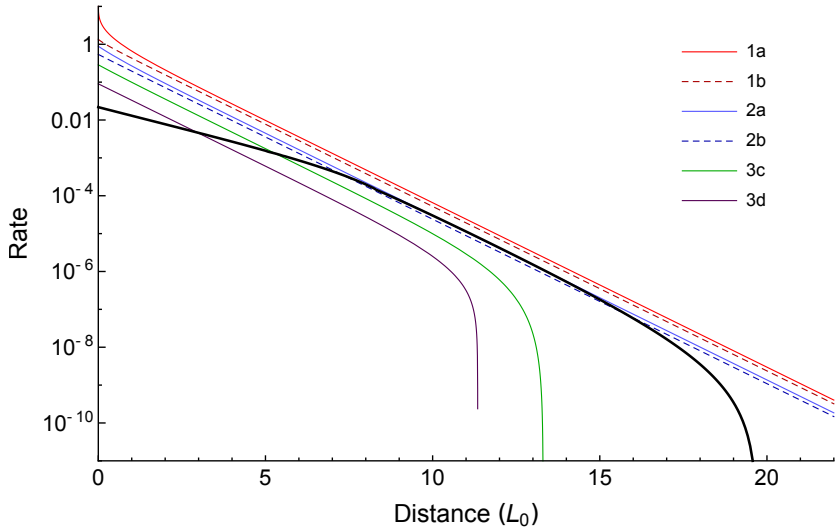

Figure 2.9: Secret-key rate with the quantum repeater implementation for the expected parameters with optimised cut-off as a function of the distance in units of $L_0 = 0.542$ km. The rate is compared to all the benchmarks defined in Table 2.1.

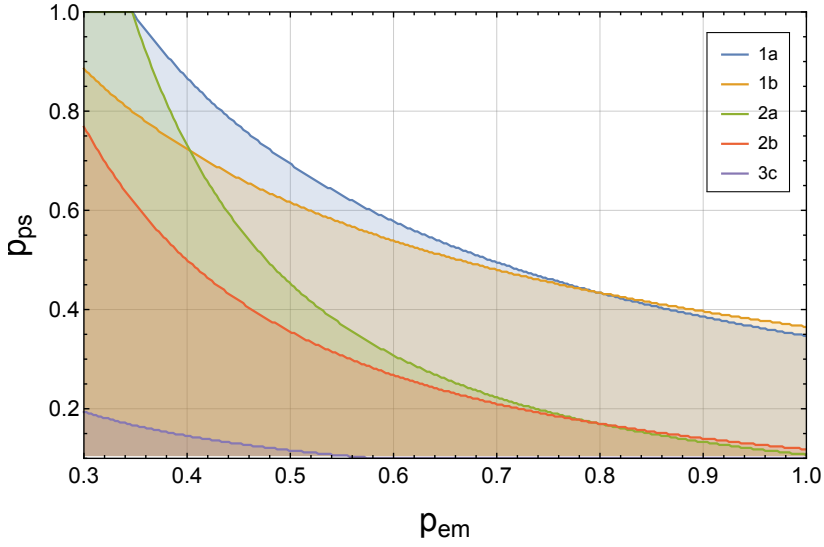

Figure 2.10: Contour plot of regions of p_{em} versus p_{ps} with the expected parameters where the benchmarks listed in Table 2.1 can be surpassed. The contour lines correspond to the parameters that achieve the corresponding benchmarks while the parameter regimes above the curves allow us to surpass them. The data is plotted for the distance of $9.6L_0$, where $L_0 = 0.542$ km.

dark counts. With the cut-off however, the repeater remains in the middle making both of the individual links Alice-repeater and repeater-Bob shorter than the Alice-repeater link in the no cut-off scheme. As a result the setup with the cut-off and with the improved parameters allows us to overcome the channel capacity (1a) more confidently and over larger range of distances, than without the cut-off.

Comparison with the proposed benchmarks. Let us now investigate the secret-key rate achievable with the expected parameters and how it compares with the proposed benchmarks. The comparison is depicted in FIG. 2.9. The benchmarks corresponding to direct transmission (3d), the thermal-loss channel (3c) and the pure-loss channel with energy constraint and inclusion of post-selection (2b) are outperformed. The achievable secret-key rate is also very close to the pure-loss channel benchmark with post-selection (2a). The other benchmarks are not overcome but are within experimental reach.

Parameter trade-off. Let us now give a general overview of how good the improved parameters need to be in order to overcome individual benchmarks. This information is presented on two contour plots. In FIG. 2.10, we study the parameter regions for which it is possible to beat the benchmarks in Table 2.1 as a function of p_{ps} and p_{em}. A similar plot as a function of F_{gm} and p_{em} can be seen in FIG. 2.11. We omit here the direct transmission benchmark which, as we have already seen, can be easily surpassed with the expected parameters. Moreover, we note that the capacity of the thermal channel in the benchmark (3c) goes to zero for very low p_{ps} and p_{em} for which it is still possible to generate key with the quantum repeater. Hence it is trivially easy to beat this benchmark for low p_{ps} and p_{em}. In that sense this benchmark is not so interesting in that regime. It is for this reason that this regime is not depicted on the contour plots. In both FIG. 2.10 and FIG. 2.11 we observe a crossing between the finite energy benchmarks (1b) and (2b) and their infinite energy counterparts (1a) and (2a) at $p_{\text{em}} \approx 0.796$, as discussed in Section 2.5.

Comparison with the proposed benchmarks for a commonly used telecom channel. Let us now again investigate the secret-key rate achievable with the expected parameters and how it compares with the proposed benchmarks, but this time assuming that we have an available channel at the commonly used telecom wavelength with attenuation length $L_0 = 22$ km. Hence in this case the frequency conversion of the emitted light into telecom would be applied. We consider such a conversion process with efficiency of 30% [180]. This parameter can be added to p_{em} so that we define $p'_{\text{em}} = 0.3\, p_{\text{em}}$. We note here that the assumed value of this parameter is a choice based on the specific experimental implementation. However, higher conversion efficiencies are in principle achievable. The comparison is depicted in FIG. 2.12. We see that for this choice of the direct channel, the benchmarks are more difficult to overcome. In particular only the benchmarks corresponding to direct transmission (3d) and the thermal-loss channel (3c) can be outperformed. The other benchmarks seem to be far from near-term experimental reach.

2.8. CONCLUSIONS

In this chapter, we have analysed numerically a realistic quantum repeater implementation for quantum key distribution. We have introduced two methods for improving the rates of the repeater with respect to previous proposals: advantage distillation and the cut-off. Advantage distillation is a classical post-processing method that increases the secret-key rate at all levels of noise. The cut-off on the other hand allows for a trade-off between the channel uses required and the secret-key fraction. Utilising the cut-off results in three benefits with respect to the previous scheme for the single sequential quantum repeater [99]. Firstly, the cut-off method achieves a higher rate for all distances. Secondly, the protocol is more robust against dark counts, in the sense that non-zero secret key can be generated over larger distances. Finally, the cut-off can be adjusted on the fly, unlike the repositioning of the repeater [99]. This is especially convenient in

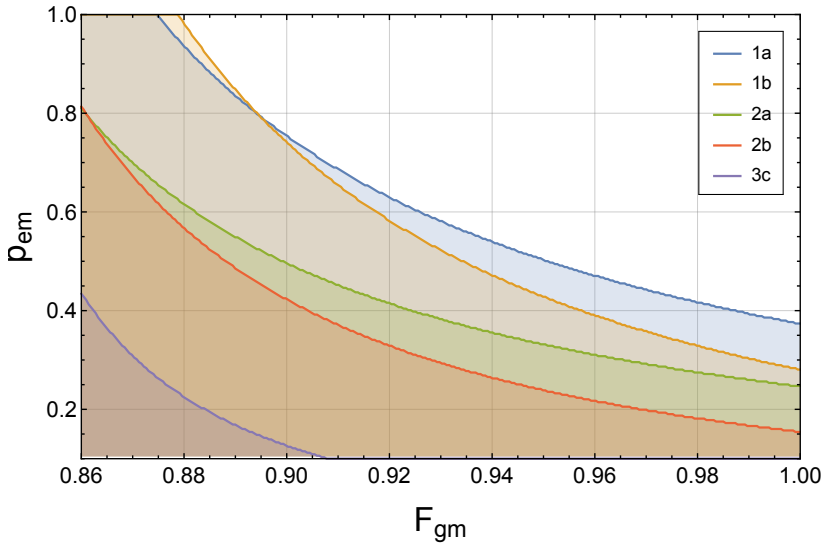

Figure 2.11: Contour plot of regions of F_{gm} versus p_{em} with the expected parameters where the benchmarks listed in Table 2.1 can be surpassed. The contour lines correspond to the parameters that achieve the corresponding benchmarks while the parameter regimes above the curves allow us to surpass them. The data is plotted for the distance of $9.6L_0$, where $L_0 = 0.542$ km.

the scenario where the experimental setup might be modified. With the previous scheme for example, improving the coherence times of the memories would lead to a new optimal position. The repositioning of the repeater node would be both costly and time-inefficient, while modifying the cut-off corresponds to a simple change in the programming of the devices.

We note here that one could also use the secret-key rate *per unit time* to assess the performance of a quantum repeater. The secret-key rate per unit time can be calculated by multiplying the secret-key fraction with the inverse of the (average) time it takes to generate a single raw bit between Alice and Bob. This time will depend on the travel time of the photons from the quantum repeater to Alice and Bob, the generation time of the entangled photon-memory pairs and the time it takes to perform the required operations such as the Bell state measurement. To compare the secret-key rate per unit time to the benchmarks, the benchmarks too must then be re-expressed in the secret-key rate per unit time. This can be achieved by multiplying the benchmarks with a fixed emission rate of a photon source [128]. Note that there is now an ambiguity in the benchmarks, as they depend on the fixed emission rate. Since the emission rate is limited by engineering constraints, the benchmarks are dependent on current technologies and cannot be claimed to be fundamental.

By optimising over the cut-off, we have found realistic parameter regions where it is possible to surpass several different benchmarks including the secret-key capacity. These benchmarks are relevant milestones towards claiming a quantum repeater, and thus form an important step in the creation of the first large-scale quantum networks.

To make our arguments concrete, we have chosen a specific parameter set induced by some recent experimental results. However, other platforms or technological advances might allow to improve upon our results and predict particularly simple setups for performing the first quantum repeater experiment. For example, our work could be extended by including other types of encod-

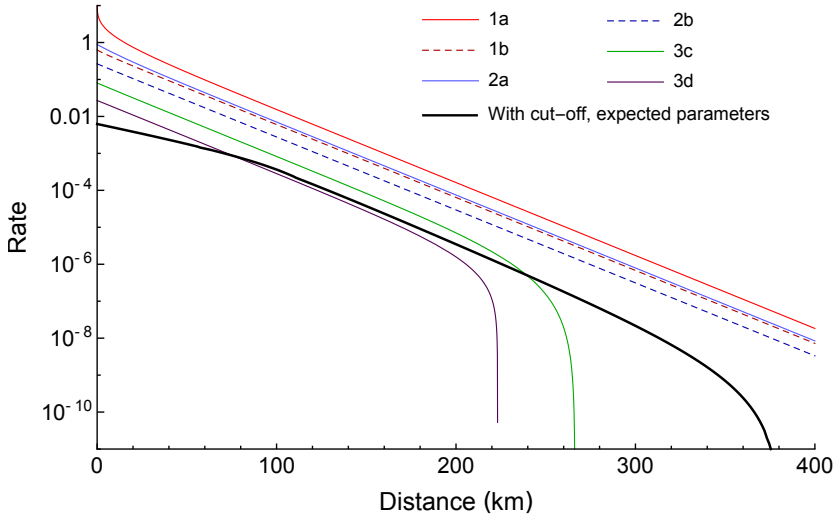

Figure 2.12: Secret-key rate for the telecom channel with $L_0 = 22$ km with the quantum repeater implementation for the expected parameters with optimised cut-off as a function of the distance in units of km. The rate is compared to all the benchmarks defined in Table 2.1.

ing, such as polarisation encoding, in which case additional depolarising noise in the fibre could become relevant. We leave the investigation of other parameter regimes open. In this respect our model has a very broad functionality, as it allows us to perform efficient optimisation of the secret-key rate over the cut-off for any set of parameters. We achieve this functionality by finding tight analytical bounds for the number of channel uses needed to generate one bit of raw key as a function of the cut-off. Our numerical package is freely available for further exploration [177].

2.9. APPENDIX

2.9.1. DARK COUNTS

In this section we detail the effect of dark counts in the detectors of Alice and Bob on our protocol. In particular, we briefly go over the concept of so-called *squashing models* [8, 57], after which we will be able to calculate the induced depolarising noise. We conclude with explaining how dark counts increase the yield.

Quantum states of light are naturally described by operators on an infinite-dimensional Hilbert space. However, a significant number of optical experiments have been performed where the infinite-dimensional states and operations are approximated by a lower dimensional description. An example of this is where the state of light is assumed to lie within a two-dimensional subspace spanned by the vacuum state and a single-photon excitation. Such an approximation is valid in the sense that the theoretical predictions of measurement statistics correspond accurately to those that are observed experimentally.

However, in cryptographic contexts one usually has to make unconditional statements about the information held by a third party. This third party might be malicious and all-powerful, and her measurement statistics are, by definition, unknown. This implies that there is not necessarily a bound on the information held by a malicious third party, despite the fact that the truncation of the Hilbert space is a good approximation for experimental statistics.

Since the theoretical analysis in an infinite-dimensional Hilbert space is difficult, one would prefer to be able to bound the information held by a third party, while at the same time applying a truncation to the finite-dimensional Hilbert space. This can be done if a so-called squashing model exists, which is a way of relating measurements performed on a high-dimensional state to a truncated space. As an approximation we consider here the squashing models for measurements of qubits encoded in the polarisation of photons. In this case squashing models exist for both the fully asymmetric BB84 protocol and the symmetric six-state protocol (with only passive measurements), implying that one can, without loss of generality, perform the fully asymmetric BB84 and symmetric (passive) six-state protocol with photons [8, 57]. The squashing model also dictates how multiple clicks in different detectors give rise to noise in the truncated space. In the next section, we discuss how to map the dark counts in the detectors to depolarising noise according to the corresponding squashing model.

The parameters typically used to quantify detectors are the dark counts per second and the detection window t_{int}, which is the duration of the integration period of the detectors. The number of thermal photons $\overline{n}$ relevant for the thermal benchmark is given by t_{int} times the dark counts per second. Assuming a Poisson distribution of the dark counts, it follows that the probability p_d of getting at least a single dark count click within the time window of awaiting the signal photon is given by $p_d = 1 - \exp(-\overline{n}) \approx \overline{n}$ for small $\overline{n}$.

The noise caused by the dark counts at Alice's or Bob's detector can then be modelled by a depolarising channel, where the depolarising parameter $\alpha_{A/B}$ depends on the implemented protocol,

$$\alpha_{A/B,\,\text{BB84}} = \frac{p_{\text{app}} p_{\text{ps}} \eta_{A/B}(1-p_d)}{1-(1-p_{\text{app}} p_{\text{ps}} \eta_{A/B})(1-p_d)^2}\,, \tag{2.18}$$

$$\alpha_{A/B,\,\text{six-state}} = \frac{p_{\text{app}} p_{\text{ps}} \eta_{A/B}(1-p_d)^5}{1-(1-p_{\text{app}} p_{\text{ps}} \eta_{A/B})(1-p_d)^6}\,. \tag{2.19}$$

That is, conditioned on a click in at least one of the detectors, Alice or Bob receive the desired state if they receive the signal photon and no other detector was triggered. Due to the squashing map all other events can be mapped onto a maximally mixed state [8, 57]. To explain the exponents, we note that the active BB84 protocol requires an optical measurement setup with two detectors, while for the six-state protocol such a measurement setup will consist of six detectors.

Furthermore, independent of the existence of a squashing map, the dark counts increase the total probability that Alice or Bob gets a click. This probability depends on whether the BB84 or six-state protocol is implemented, and is given by

$$p_{A/B,\,BB84} = 1-(1-p_{\text{app}} p_{\text{ps}} \eta_{A/B})(1-p_d)^2\,, \tag{2.20}$$

$$p_{A/B,\,\text{six-state}} = 1-(1-p_{\text{app}} p_{\text{ps}} \eta_{A/B})(1-p_d)^6\,. \tag{2.21}$$

2.9.2. QUANTUM BIT ERROR RATE

In this Appendix we derive the expressions for the average quantum bit error rate in the X, Y and Z basis as a function of the experimental parameters. It is given by

$$\langle e_X \rangle = \langle e_Y \rangle = \frac{1}{2} - \frac{1}{2} F_{\text{gm}} \alpha_A \alpha_B \left(2F_{\text{prep}} - 1\right)^2 \left\langle e^{-(a+b)n} \right\rangle\,, \tag{2.22}$$

$$\langle e_Z \rangle = \frac{1}{2} - \frac{1}{2} F_{\text{gm}} \alpha_A \alpha_B \langle e^{-b\cdot n} \rangle\,, \tag{2.23}$$

where the average is performed over the geometric distribution with only the first $n^\star$ trials. That is, the average of the exponential e^{-cn} is given by

$$\begin{aligned}\langle e^{-cn}\rangle &= \frac{\sum_{n=1}^{n^\star} p_B\left(1-p_B\right)^{n-1} e^{-cn}}{\sum_{n=1}^{n^\star} p_B\left(1-p_B\right)^{n-1}} \\ &= \frac{p_B e^{-c}}{1-\left(1-p\right)^{n^\star}} \frac{1-\left(1-p_B\right)^{n^\star} e^{-cn^\star}}{1-\left(1-p_B\right)e^{-c}} .\end{aligned} \tag{2.24}$$

To derive these quantum bit error rates, let us firstly define the two-qubit Bell states as

$$|\psi(\mathsf{x},\mathsf{z})\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0+\mathsf{x}\rangle + (-1)^{\mathsf{z}}|1\rangle|1+\mathsf{x}\,(\mathrm{mod}\,2)\rangle), \tag{2.25}$$

for $\mathsf{x},\mathsf{z}\in\{0,1\}$. The noise in the preparation can be modelled as dephasing noise [162]. The initially generated entangled state between the quantum memory and the state of the photon flying to Alice is then

$$\rho_{AR} = F_{\text{prep}}|\psi(1,0)\rangle\langle\psi(1,0)| + (1-F_{\text{prep}})|\psi(1,1)\rangle\langle\psi(1,1)| , \tag{2.26}$$

where F_{prep} is the preparation fidelity of this state. The state in the first quantum memory is now kept stored there. During this time, a second entangled photon-memory is attempted to be generated at the second quantum memory. During these attempts, the state stored in the first quantum memory decoheres through time-dependent dephasing and depolarising noise acting on it. This means that at the time when the second copy is generated, the first copy will have decohered. This second copy will be of the same form as the first one. The decohered first copy is of the form

$$\begin{aligned}\rho'_{AR} = F_{T_1}[&F_{\text{prep}}(F_{T_2}|\psi(1,0)\rangle\langle\psi(1,0)| + (1-F_{T_2})|\psi(1,1)\rangle\langle\psi(1,1)|) \\ &+ (1-F_{\text{prep}})\left(F_{T_2}|\psi(1,1)\rangle\langle\psi(1,1)| + (1-F_{T_2})|\psi(1,0)\rangle\langle\psi(1,0)|\right)] + (1-F_{T_1})\frac{\mathbb{I}}{4} ,\end{aligned} \tag{2.27}$$

where F_{T_1}, F_{T_2} are respectively the depolarising and dephasing parameters due to the decoherence processes on the stored state in the first memory. The fidelity decays exponentially with the number of attempts [135] and hence these parameters be written as

$$F_{T_1} = e^{-b\cdot n} , \tag{2.28}$$

$$F_{T_2} = \frac{1+e^{-a\cdot n}}{2} . \tag{2.29}$$

Here n is the number of attempts that have been performed on the second memory to successfully generate the repeater-Bob entanglement and the decay rates a and b are defined in the main text. Hence we can rewrite the state of ρ'_{AR} as

$$\rho'_{AR} = F_{T_1}(F_{\text{deph},AR}|\psi(1,0)\rangle\langle\psi(1,0)| \tag{2.30}$$

$$+(1-F_{\text{deph},AR})|\psi(1,1)\rangle\langle\psi(1,1)|) + (1-F_{T_1})\frac{\mathbb{I}}{4} . \tag{2.31}$$

where

$$F_{\text{deph},AR} = \frac{1+(2F_{\text{prep}}-1)e^{-an}}{2} . \tag{2.32}$$

The entanglement swapping is performed at the two memories at the repeater node. Since the situation is symmetric for all the four measurement outcomes, without loss of generality we can

consider the resulting state on AB as if the repeater measured $|\psi(1,0)\rangle$. If a different Bell state was measured, a Pauli rotation could be used to bring the state to this form. The state that we obtain is

$$\rho''_{AB} = F_{T_1}\Big(\Big[F_{\text{deph},AR}F_{\text{prep}} + (1-F_{\text{deph},AR})(1-F_{\text{prep}})\Big]|\psi(1,0)\rangle\langle\psi(1,0)| + \Big[F_{\text{deph},AR}(1-F_{\text{prep}}) + (1-F_{\text{deph},AR})F_{\text{prep}}\Big]|\psi(1,1)\rangle\langle\psi(1,1)|\Big) + \big(1-F_{T_1}\big)\frac{\mathbb{I}}{4}\,. \tag{2.33}$$

Finally we note that the operations such as Bell state measurements or any other required gates performed on the memories are also noisy. We will model them by the depolarising channel here [35]. The depolarising channel commutes with the dephasing channel. For the two copies of the Bell-diagonal state, it also commutes with the entanglement swapping, in the sense that applying it to one of our memory qubits is mathematically equivalent to applying the same channel to one of the photons flying to Alice or Bob. Hence independently of when exactly in the protocol those gates or measurements on the memories are applied, we can add the resulting depolarisation to the final state shared between Alice and Bob, so that we obtain

$$\rho''_{AB} = F_{\text{gm}}\alpha_A\alpha_B F_{T_1}\Big(\Big[F_{\text{deph},AR}F_{\text{prep}} + (1-F_{\text{deph},AR})(1-F_{\text{prep}})\Big]|\psi(1,0)\rangle\langle\psi(1,0)| + \Big[F_{\text{deph},AR}(1-F_{\text{prep}}) + (1-F_{\text{deph},AR})F_{\text{prep}}\Big]|\psi(1,1)\rangle\langle\psi(1,1)|\Big) + \big(1-F_{\text{gm}}\alpha_A\alpha_B F_{T_1}\big)\frac{\mathbb{I}}{4}\,. \tag{2.34}$$

Here by F_{gm} we denote the product of all the depolarising parameters corresponding to all noisy gates and measurements and $\alpha_{A/B}$ corresponds to the noise caused by the dark counts on Alice's/Bob's side. From the final state it follows that

$$\langle e_X\rangle = \langle e_Y\rangle = \frac{1}{2} - \frac{1}{2}F_{\text{gm}}\alpha_A\alpha_B\left(2F_{\text{prep}}-1\right)^2\left\langle e^{-(a+b)n}\right\rangle\,, \tag{2.35}$$

$$\langle e_Z\rangle = \frac{1}{2} - \frac{1}{2}F_{\text{gm}}\alpha_A\alpha_B\langle e^{-b\cdot n}\rangle\,. \tag{2.36}$$

where the average is over the geometric distribution with only the first $n^\star$ trials. This is due to the fact that, by construction, the state is never allowed to decohere more than $n^\star$ trials.

2.9.3. COMPARISON WITH MEMORY-ASSISTED MDI QKD SCHEMES

The setup of the proof-of-principle repeater analysed in this chapter bears close resemblance to the memory-assisted measurement-device-independent QKD (MA-MDI QKD) setups proposed in [121], which were analysed in more detail in the particular context of NV centres in [127]. However, in contrast to our focus on key per channel use, these schemes were mostly assessed on their performance of generating key per unit time. In this section, we will briefly discuss these schemes and their advantages and disadvantages in comparison to the scheme analysed in this chapter. In particular, we will focus both on their relevance in the context of secret-key generation per channel use, and on the complexity of their experimental implementation.

The three schemes that we compare with can be found in Figure 2.13. These schemes have the advantage of high expected rate per unit time, since heralding of the successful events now takes place at the repeater. Thus, after a failed attempt the repeater can immediately prepare for receiving another photon, without the need for waiting on any classical communication from Alice and Bob. Furthermore, these schemes are secure against detector side-channel attacks [96], since in each scheme there is no quantum information sent from the repeater to Alice or Bob.

However, these advantages, while relevant in practical QKD setups, might not necessarily translate directly in higher secret-key rate per channel use for proof-of-principle repeaters. Moreover, there are experimental challenges that make these MA-MDI QKD schemes more difficult to implement than the sequential quantum repeater that we consider. This is particularly important,

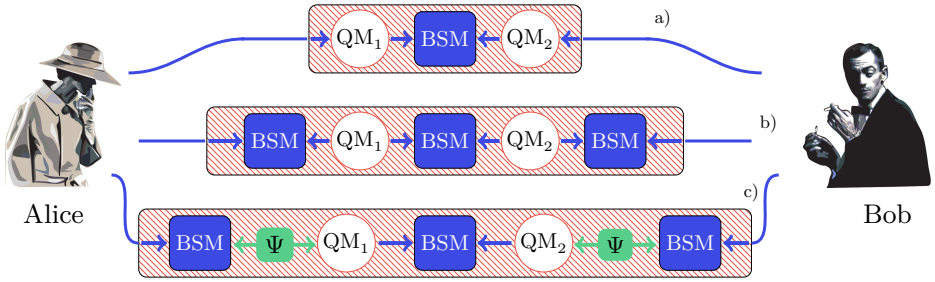

Figure 2.13: Three different setups for the memory-assisted measurement-device-independent quantum key distribution (MA-MDI QKD). Here, BSM stands for Bell state measurement. The first setup a) corresponds to the scheme of MA-MDI QKD with direct heralding. Specifically, the implementation of this setup requires that a photonic state can be transferred into a quantum memory QM_1 and QM_2 in a heralded fashion. That is, following the transfer attempt, one obtains the information whether the state of the photon emitted at Alice or Bob has been successfully transferred to the desired quantum memory. The second setup b) with indirect heralding is a modification of the first one. Here the requirement of the heralded state transfer has been dropped, at the cost of probabilistic Bell state measurements between two photonic qubits at the outer BSM stations. Finally, the setup in c) is a modification of b), which uses sources of entangled photons (Ψ). In this way, the attempt to transfer the quantum state of the photon into the memory is performed only after a successful Bell state measurement. This can increase the rate per unit time, since writing unto and resetting the memory is a time-consuming process.

since the goal of this chapter is to analyse a protocol that would be simple from the implementation perspective, and would have the capability to exceed the benchmarks in Section 2.5.

Let us now go over each of these schemes. Firstly, let us consider the first scheme a). This scheme seems to require a similar number of components as our proposed scheme, with the exception that the two detector setups have now been replaced with the sources of BB84 states. The main difficulty with implementing such a scheme lies in the requirement of heralded quantum state transfer from a single photon into the quantum memory. This is a great challenge from the experimental perspective and is not expected to be realised with high fidelity on a significant number of physical platforms in the near future. In systems that utilise cavities this task can be performed, provided that one can realise a low-loss overcoupled cavity with high cooperativity. While such a scenario has been demonstrated experimentally in trapped atoms by achieving the strong coupling regime [82], demonstrating high cooperativity is very challenging in general.

Due to the reasons explained above, scheme b) seems more realistic than scheme a) with the current state-of-the-art technology. However, a larger number of components is needed and the two additional optical Bell state measurements will reduce the rate by a factor of four. In particular, photonic states need to be emitted both from the quantum memories and the BB84 sources. These need to be synchronised such that the Bell state measurements can be performed on both of them. While there is nothing fundamentally challenging with this scheme, it requires larger number of components and is more complicated than the scheme analysed in this chapter. Similar conclusions apply to the more complex scheme proposed in c), which adds sources of entangled photons (denoted here by Ψ) into the scheme of b). A comparison of the achieved secret-key rate with the secret-key capacity, for a variant of scheme c), has been performed in [97].

2.9.4. SECRET-KEY FRACTION AND ADVANTAGE DISTILLATION

In this section the secret-key fraction formula for the six-state protocol with advantage distillation of [171] is briefly reviewed. We note here that while the analysis in Appendix 2.9.2 has the state $|\psi(1,0)\rangle$ as the target state, here we follow the analysis of [171] for which $|\psi(0,0)\rangle$ is the target state. This does not affect the overall analysis as the final state from Appendix 2.9.2 can be rotated locally such that $|\psi(0,0)\rangle$ could be made the target state. The secret key fraction can be expressed in terms of the Bell coefficients of the Bell diagonal state

$$\rho_{AB} = \sum_{x,z\in\{0,1\}} P_{XZ}(x,z)\,|\psi(x,z)\rangle\langle\psi(x,z)|\,. \tag{2.37}$$

Here P_{XZ} is a probability distribution and we will abbreviate $P_{XZ}(x,z)$ as p_{xz}. For the description of the advantage distillation protocol we refer the reader to [171]. It is shown there that the secret-key fraction can be written as

$$r_{\text{six-state}} \tag{2.38}$$

$$= \frac{1}{3}\max\left[1 - H(P_{XZ}) + \frac{P_{\tilde{X}}(1)}{2}h\left(\frac{p_{00}p_{10}+p_{01}p_{11}}{(p_{00}+p_{01})(p_{10}+p_{11})}\right), \frac{P_{\tilde{X}}(0)}{2}(1-H(P'_{XZ}))\right], \tag{2.39}$$

where

$$P_{\tilde{X}}(0) = (p_{00}+p_{01})^2 + (p_{10}+p_{11})^2\,, \tag{2.40}$$

$$P_{\tilde{X}}(1) = 2(p_{00}+p_{01})(p_{10}+p_{11})\,, \tag{2.41}$$

$$P'_{XZ}(0,0) = \frac{p_{00}^2+p_{01}^2}{(p_{00}+p_{01})^2+(p_{10}+p_{11})^2}, \tag{2.42}$$

$$P'_{XZ}(1,0) = \frac{2p_{00}p_{01}}{(p_{00}+p_{01})^2+(p_{10}+p_{11})^2}, \tag{2.43}$$

$$P'_{XZ}(0,1) = \frac{p_{10}^2+p_{11}^2}{(p_{00}+p_{01})^2+(p_{10}+p_{11})^2}, \tag{2.44}$$

$$P'_{XZ}(1,1) = \frac{2p_{10}p_{11}}{(p_{00}+p_{01})^2+(p_{10}+p_{11})^2}\,, \tag{2.45}$$

and $H(P_{XZ})$ is the Shannon entropy of the distribution P_{XZ}. The factor of a third arises from the fact that for a symmetric six-state protocol only a third of the measurements will be performed in the same basis by Alice and Bob.

In our model we only consider depolarising noise and dephasing noise in standard basis. Hence for the six-state protocol the error rates in X and Y basis will be the same. Therefore

$$p_{10}+p_{11} = e_Z\,, \tag{2.46}$$

$$p_{01}+p_{11} = e_{XY}\,, \tag{2.47}$$

$$p_{01}+p_{10} = e_{XY}\,, \tag{2.48}$$

$$p_{00}+p_{01}+p_{10}+p_{11} = 1\,. \tag{2.49}$$

Hence

$$p_{00} = 1 - \frac{e_Z}{2} - e_{XY}\,, \tag{2.50}$$

$$p_{01} = e_{XY} - \frac{e_Z}{2}\,, \tag{2.51}$$

$$p_{10} = p_{11} = \frac{e_Z}{2}\,. \tag{2.52}$$

And so

$$P_{\tilde{X}}(0) = 1 - 2e_Z + 2e_Z^2 \,, \tag{2.53}$$

$$P_{\tilde{X}}(1) = 2(1 - e_Z)e_Z \,. \tag{2.54}$$

2.9.5. YIELD

In this Appendix we derive the analytical approximation for the yield with the cut-off $n^\star$. The yield Y is given by

$$Y = \frac{p_{\text{bsm}}}{\mathbb{E}[N]} = \frac{p_{\text{bsm}}}{\mathbb{E}[\max(N_A, N_B)]} \,. \tag{2.55}$$

The approximation used for $\mathbb{E}[\max(N_A, N_B)]$ is

$$\mathbb{E}[\max(N_A, N_B)] \approx \begin{cases} \frac{1}{p_A\left(1-(1-p_B)^{n^\star}\right)} & \frac{1}{p_A} \geq n^\star \\ \frac{1}{p_A} + \frac{1}{p_B} - \frac{1}{p_A + p_B - p_A p_B} & \frac{1}{p_A} < n^\star \,, \end{cases} \tag{2.56}$$

where p_A and p_B are defined in Eq. (2.9) for BB84 and in Eq. (2.10) for the six-state protocol. In the rest of this Appendix, we will motivate this approximation by finding tight analytical lower and upper bounds on $\mathbb{E}[N]$.

We note that we consider separately two parameter regimes. One of them is the regime where on average the dominant number of channel uses per round is on Alice's side $\left(\frac{1}{p_A} > n^\star\right)$. This corresponds to the high-loss regime since the number of channel uses per round on Bob's side is upper bounded by the cut-off. The other regime is the low-loss regime $\left(\frac{1}{p_A} \leq n^\star\right)$. In this regime we will show that the cut-off does not play any significant role, so that in this regime the formula for the yield with no cut-off [99, 121] can be used. Moreover, for our derivation to be valid we require an additional constraint to be satisfied, namely $p_B \geq p_A$. This means that we cannot consider scenarios when the repeater is positioned closer to Alice than to Bob. Such a constraint is well-justified since the time-dependent decoherence in quantum memory QM_1 would only increase by shifting the repeater towards Alice.

HIGH-LOSS REGIME

The high-loss regime is the regime where the losses on Alice's side together with the cut-off on Bob's side ensure that the predominant number of channel uses is almost always on Alice's side, i.e. $\mathbb{E}[N] = \mathbb{E}[\max(N_A, N_B)] \approx \mathbb{E}[N_A]$. This regime is described by the condition $p_A n^\star < 1$. More specifically, as we will show in this section, if

$$\frac{1}{p_A} := \mu = \beta n^\star, \; \beta > 1 \,, \tag{2.57}$$

then

$$\mathbb{E}[N_A] \leq \mathbb{E}[N] \leq \left(g_{\text{err}}(p_A, p_B, n^\star) + 1\right)\mathbb{E}[N_A] \,, \tag{2.58}$$

where $\mathbb{E}[N_A] = \frac{1}{p_A(1-(1-p_B)^{n^\star})}$ (see Eq. (2.66)) and $g_{\text{err}}(p_A, p_B, n^\star) = \mathcal{O}\left(\frac{1}{\beta^2}\right)$ is a function defined in Eq. (2.86). This implies that for β large enough, $\mathbb{E}[N]$ can be accurately approximated by $\frac{1}{p_A(1-(1-p_B)^{n^\star})}$.

We start the proof of Eq. (2.58) by first noticing that $\mathbb{E}[N_A] \leq \mathbb{E}[N]$. It is, thus, only necessary to find an upper bound for $\mathbb{E}[N]$. Now, let $p(K = k) = (1 - p_r)^{k-1} p_r$ be the probability that Bob

succeeds in round k. Here $p_r = 1-(1-p_B)^{n^\star}$ is the probability that Bob succeeds in a given round. Then

$$\begin{aligned}\mathbb{E}[N] &= \mathbb{E}[\max(N_A, N_B)] \\ &= \sum_{k=1}^{\infty} p(K=k)\left(\sum_{n_A=k}^{\infty}\left(\sum_{n_B=(k-1)n^\star+1}^{kn^\star} p(N_A=n_A \wedge N_B=n_B|K=k)\max(n_A,n_B)\right)\right). \end{aligned} \tag{2.59}$$

One can split the sum over n_A in two, depending on whether n_A is greater than n_B or vice versa. We get

$$\begin{aligned}\mathbb{E}[N] = \sum_{k=1}^{\infty} p(k)\Bigg(&\sum_{n_B=(k-1)n^\star+1}^{kn^\star}\left(\sum_{n_A=k}^{n_B} p(n_A\wedge n_B|k)n_B\right) \\ &+ \sum_{n_B=(k-1)n^\star+1}^{kn^\star}\left(\sum_{n_A=n_B+1}^{\infty} p(n_A\wedge n_B|k)n_A\right)\Bigg),\end{aligned} \tag{2.60}$$

where $p(k) = p(K=k)$, and $p(n_A\wedge n_B|k) = p(N_A = n_A \wedge N_B = n_B|K=k)$. The first term of Eq. (2.60) can be upper bounded noticing that $n_B \leq kn^\star$, i.e.

$$\sum_{k=1}^{\infty} p(k)\left(\sum_{n_B=(k-1)n^\star+1}^{kn^\star}\left(\sum_{n_A=k}^{n_B} p(n_A\wedge n_B|k)n_B\right)\right) \leq \sum_{k=1}^{\infty} p(k)\,p\left(N_A \leq N_B|K=k\right)kn^\star. \tag{2.61}$$

The second term of Eq. (2.60) can be upper bounded in the following way

$$\begin{aligned}\sum_{k=1}^{\infty} p(k)\left(\sum_{n_B=(k-1)n^\star+1}^{kn^\star}\left(\sum_{n_A=n_B+1}^{\infty} p(n_A\wedge n_B|k)n_A\right)\right) &\leq \sum_{k=1}^{\infty} p(k)\left(\sum_{n_A=k}^{\infty} p(n_A|k)n_A\right) && (2.62)\\ &= \sum_{k=1}^{\infty} p(k)\sum_{n_A=1}^{\infty} p(n_A|k)n_A && (2.63)\\ &= \sum_{n_A=1}^{\infty} p(n_A)n_A = \mathbb{E}[N_A]\,. && (2.64)\end{aligned}$$

Inputting Eq. (2.61) and Eq. (2.64) back into Eq. (2.60), we obtain

$$\mathbb{E}[N] \leq \left(\frac{n^\star}{\mathbb{E}[N_A]}\sum_{k=1}^{\infty} p(k)\,p\left(N_A \leq N_B|k\right)k + 1\right)\mathbb{E}[N_A]\,. \tag{2.65}$$

Let N_A^i be the random variable describing the number of trials on Alice's side in round i. Since $p(N_A^i = n_A^i) = (1-p_A)^{n_A^i-1}p_A$, we clearly have that $\mathbb{E}[N_A^i] = \frac{1}{p_A} = \mu$. Then we note that

$$\begin{aligned}\mathbb{E}[N_A] &= \sum_{k=1}^{\infty} p(k)\sum_{i=1}^{k}\sum_{n_A^i=1}^{\infty} p(n_A^i)n_A^i = \sum_{k=1}^{\infty} p(k)\sum_{i=1}^{k}\mathbb{E}[N_A^i] \\ &= \mu\sum_{k=1}^{\infty} p(k)k = \mathbb{E}[K]\mu = \frac{1}{p_A p_r} = \frac{1}{p_A(1-(1-p_B)^{n^\star})}\,.\end{aligned} \tag{2.66}$$

Here, we first express $\mathbb{E}[N_A]$ by calculating the average number of trials in each of the k rounds. Then, we sum the k averages together, and finally, we average over the total number of rounds k.

Since all the rounds are independent, we replace each $\mathbb{E}[N_A^i]$ by μ as stated above. By inputting Eq. (2.66) into Eq. (2.65), we get

$$\mathbb{E}[N_A] \leq \mathbb{E}[N] \leq \left(\frac{1}{\mathbb{E}[K]\beta} \sum_{k=1}^{\infty} p(k)\, p(N_A \leq N_B | k)\, k + 1 \right) \mathbb{E}[N_A] \,. \tag{2.67}$$

We now upper bound the $p(N_A \leq N_B|k)$ term. Note that

$$p(N_A \leq N_B|k) = p\left(\sum_{i=1}^{k} N_A^i \leq \sum_{i=1}^{k} N_B^i \Big| k \right) . \tag{2.68}$$

We note that conditioned on $K = k$, we have that $\sum_{i=1}^{k} N_B^i = (k-1)n^\star + N_B^k$. It then follows that

$$p(N_A \leq N_B|k) = p\left(\sum_{i=1}^{k} N_A^i \leq (k-1)n^\star + N_B^k \Big| k \right) \leq p\left(\sum_{i=1}^{k} N_A^i \leq k n^\star \Big| k \right) . \tag{2.69}$$

Condition Eq. (2.57) and $-\sum_{i=1}^{k} N_A^i \geq -kn^\star$ is equivalent to $k\mu - \sum_{i=1}^{k} N_A^i \geq k(\beta-1)n^\star$. Hence,

$$p\left(\sum_{i=1}^{k} N_A^i \leq k n^\star \Big| k \right) = p\left(k\mu - \sum_{i=1}^{k} N_A^i \geq k(\beta-1)n^\star \Big| k \right) . \tag{2.70}$$

We can use the Chernoff bound to upper bound this probability. The Chernoff bound for a random variable X is

$$p(X \geq a) \leq \frac{\mathbb{E}[e^{tX}]}{e^{ta}}, \; t > 0 \,. \tag{2.71}$$

Let X be the sum of k random variables $X_1, X_2, \ldots, X_k$, where

$$X_i = \mu - N_A^i \,, \tag{2.72}$$

i.e. $X = \sum_{i=1}^{k} X_i = k\mu - \sum_{i=1}^{k} N_A^i$. From this we can now bound the desired probability. Using (2.71) and $a = k(\beta-1)n^\star$, we obtain the inequality

$$p\left(k\mu - \sum_{i=1}^{k} N_A^i \geq k(\beta-1)n^\star \Big| k \right) \leq \frac{\mathbb{E}\left[\exp\left(t\left(k\mu - \sum_{i=1}^{k} N_A^i \right) \right) \Big| k \right]}{e^{tk(\beta-1)n^\star}} \tag{2.73}$$

$$= \exp\left[tk\left(\mu - (\beta-1)n^\star \right) \right] \mathbb{E}\left[\Pi_{i=1}^{k} e^{-tN_A^i} | k \right] . \tag{2.74}$$

Let us now focus on $\mathbb{E}\left[\prod_{i=1}^{k} e^{-tN_A^i} | k \right]$,

$$\mathbb{E}\left[\prod_{i=1}^{k} e^{-tN_A^i} | k \right] = \prod_{i=1}^{k} \mathbb{E}\left[e^{-tN_A^i} | k \right] \tag{2.75}$$

$$= \prod_{i=1}^{k} \left(\sum_{n_A^i=1}^{\infty} p_A (1-p_A)^{n_A^i - 1} e^{-t n_A^i} \right) = \left(\frac{p_A e^{-t}}{1-(1-p_A)e^{-t}} \right)^k . \tag{2.76}$$

Here, after the first equality sign we have used the fact that the random variables N_A^i are independent for different i's. After the second equality we note that all of them have exactly the same geometric distribution over the k rounds. Specifically, it is now important to note that this holds provided that k is the value of K on which we have conditioned, i.e., the success on Bob's side

occurs exactly in the k'th round. Furthermore, the common ratio $(1-p_A)e^{-t}$ satisfies the convergence condition $\left|(1-p_A)e^{-t}\right| < 1$ for all $t > 0$. This yields

$$p\left(N_A \leq N_B | K = k\right) \leq \left(\exp\left[t\left(\frac{1}{p_A} - (\beta-1)n^\star\right)\right]\frac{p_A e^{-t}}{1-(1-p_A)e^{-t}}\right)^k . \tag{2.77}$$

Let's define the function $f(t)$ as

$$f(t) := \exp\left[t\left(\frac{1}{p_A} - (\beta-1)n^\star\right)\right]\frac{p_A e^{-t}}{1-(1-p_A)e^{-t}} . \tag{2.78}$$

This function should be minimised subject to $t > 0$ to obtain the tightest bound. A single stationary point is analytically found at

$$t_0 = \ln\left(\frac{(1-p_A)(p_A(\beta-1)n^\star - 1)}{p_A(\beta-1)n^\star + p_A - 1}\right) . \tag{2.79}$$

We now want to make sure that t_0 always satisfies the condition $t > 0$, necessary for applying the Chernoff bound. By condition Eq. (2.57), the denominator of the above expression inside the logarithm is $p_A(\beta-1)n^\star + p_A - 1 = 1 - p_A n^\star + p_A - 1 = p_A(1-n^\star) < 0$ as long as $n^\star > 1$. From this it follows that $t_0 > 0$ if and only if

$$(1-p_A)(p_A(\beta-1)n^\star - 1) < p_A(\beta-1)n^\star + p_A - 1 . \tag{2.80}$$

Clearly this condition is equivalent to $-p_A^2(\beta-1)n^\star < 0$ which is satisfied for $\beta > 1$. This means that $t_0 > 0$ is always satisfied. Now note that $f(t=0) = 1$. Moreover, one can also easily verify that $f'(t=0) = n^\star(1-\beta) < 0$ for $\beta > 1$, and that $\lim_{t\to\infty} f(t) \to \infty$ as long as $n^\star > 1$. These properties of $f(t)$, together with the continuity of $f(t)$, prove that $t = t_0$ corresponds to the global minimum of this function in the regime $t > 0$ and that $f(t_0) < 1$. Hence, we can now calculate $f(t_0)$ which gives

$$f(t_0) = \left(\frac{(p_A(\beta-1)n^\star - 1)(1-p_A)}{p_A(\beta-1)n^\star + p_A - 1}\right)^{\frac{1}{p_A} - (\beta-1)n^\star - 1} (1 - p_A(\beta-1)n^\star) . \tag{2.81}$$

This formula can be simplified by substituting the condition Eq. (2.57) to eliminate β

$$f(t_0) = p_A n^\star \left(\frac{n^\star(1-p_A)}{n^\star - 1}\right)^{n^\star - 1} . \tag{2.82}$$

$\mathbb{E}[N]$ can now be upper bounded by an expression that depends on $f(t_0)$, that is

$$\mathbb{E}[N] \leq \left(\frac{1}{\mathbb{E}[K]\beta}\sum_{k=1}^{\infty} p(K=k) f(t_0)^k k + 1\right)\mathbb{E}[N_A] . \tag{2.83}$$

We can now average over the number of rounds k,

$$\sum_{k=1}^{\infty}\frac{p_r}{(1-p_r)}\left[(1-p_r)f(t_0)\right]^k k = \frac{p_r f(t_0)}{\left[1-(1-p_r)f(t_0)\right]^2} . \tag{2.84}$$

Moreover, $\mathbb{E}[K] = \frac{1}{p_r}$ and again removing β through condition Eq. (2.57) yields

$$\begin{aligned}\mathbb{E}[N] &\leq \left(\frac{p_r^2 p_A n^\star f(t_0)}{\left[1-(1-p_r)f(t_0)\right]^2} + 1\right)\mathbb{E}[N_A] \\ &= \left(\frac{(1-(1-p_B)^{n^\star})^2 p_A n^\star f(t_0)}{\left[1-(1-p_B)^{n^\star} f(t_0)\right]^2} + 1\right)\mathbb{E}[N_A] .\end{aligned} \tag{2.85}$$

Now by taking the number of channel uses to be $\mathbb{E}[N_A]$, we can define the relative error $g_{\text{err}}(p_A, p_B, n^\star)$,

$$g_{\text{err}}(p_A, p_B, n^\star) := \frac{(1-(1-p_B)^{n^\star})^2 p_A n^\star f(t_0)}{\left[1-(1-p_B)^{n^\star} f(t_0)\right]^2}\,, \tag{2.86}$$

with $f(t_0)$ given in Eq. (2.82), so that

$$\mathbb{E}[N_A] \leq \mathbb{E}[N] \leq \left(g_{\text{err}}(p_A, p_B, n^\star)+1\right)\mathbb{E}[N_A]\,, \tag{2.87}$$

where the conditions required to satisfy the above formula are $n^\star > 1$ and $p_A n^\star < 1$. Finally, we can now show how $g_{\text{err}}(p_A, p_B, n^\star)$ scales with β. Note that

$$f(t_0) \leq p_A n^\star \left(1+\frac{1}{n^\star-1}\right)^{n^\star-1} \leq p_A n^\star e\,. \tag{2.88}$$

This together with $f(t_0) < 1$ gives

$$g_{\text{err}}(p_A, p_B, n^\star) < \frac{p_r^2 (p_A n^\star)^2 e}{p_r^2} = \frac{e}{\beta^2}\,. \tag{2.89}$$

Therefore $g_{\text{err}}(p_A, p_B, n^\star) = \mathcal{O}\left(\frac{1}{\beta^2}\right)$, implying that the bounds in the high-loss regime are good enough to tightly bound the achieved yield.

LOW-LOSS REGIME

Now we consider the complementary low-loss regime characterised by the condition $p_A n^\star \geq 1$. Firstly, since in our protocol there is never any benefit in placing the repeater closer to Alice than to Bob, we also have that $p_B \geq p_A$. This implies that $\frac{1}{p_B} \leq \frac{1}{p_A} = \mathbb{E}[N_A^i] \leq n^\star$. This is the regime where the cut-off is large in comparison with the average number of channel uses required to detect a single photon on Bob's side. That is,

$$\frac{\beta'}{p_B} = n^\star,\ n^\star \geq \beta' \geq 1\,. \tag{2.90}$$

As we will show in this section, in this region we can approximate $\mathbb{E}[N] = \mathbb{E}[\max(N_A, N_B)]$ by N_{NC}, where

$$N_{NC} = \frac{1}{p_A} + \frac{1}{p_B} - \frac{1}{p_A + p_B - p_A p_B}\,, \tag{2.91}$$

is the average number of channel uses in the no cut-off (NC) scenario [99, 121]. Intuitively, this is because Alice and Bob almost never have to restart due to Bob reaching the cut-off. More specifically, we show that

$$N_{NC} \leq \mathbb{E}[N] \leq \left(\tilde{g}_{\text{err}}(p_A, p_B, n^\star)+1\right) N_{NC}\,, \tag{2.92}$$

where $\tilde{g}_{\text{err}}(p_A, p_B, n^\star)$ is defined in Eq. (2.105). Since $\tilde{g}_{\text{err}}(p_A, p_B, n^\star) = \mathcal{O}\left(\beta' e^{-\beta'}\right)$, for sufficiently large β' the expectation value $\mathbb{E}[N]$ can be accurately approximated by N_{NC}.

Here we detail a proof of Eq. (2.92). We note that the presence of the cut-off increases the number of needed channel uses with respect to the no cut-off scenario, i.e. $N_{NC} \leq \mathbb{E}[N]$. For the

upper bound we can write now

$$\mathbb{E}[N] = \mathbb{E}[\max(N_A, N_B)] \tag{2.93}$$

$$= \sum_{k=1}^{\infty} p(K=k) \left(\sum_{n_A=k}^{\infty} \left(\sum_{n_B=(k-1)n^\star+1}^{kn^\star} p(n_A \wedge n_B | K=k) \max(n_A, n_B) \right) \right) \tag{2.94}$$

$$= p(K=1) \sum_{n_B=1}^{n^\star} \sum_{n_A=1}^{\infty} p(n_A|K=1) p(n_B|K=1) \max(n_A, n_B)$$

$$+ \sum_{k=2}^{\infty} p(K=k) \left(\sum_{n_B=(k-1)n^\star+1}^{kn^\star} \left(\sum_{n_A=k}^{\infty} p(n_A \wedge n_B|k) \max(n_A, n_B) \right) \right) . \tag{2.95}$$

In Eq. (2.95) we split the sum over k into two terms, one with $k=1$ and the other with $k>1$. Since the first term has fixed $k=1$, the variables N_A and N_B are independent here (there is only one round in which Bob for sure succeeds, so the value of n_B does not affect the value of n_A). Moreover, the geometric distribution of N_B is normalised over the interval $[1,\ldots,n^\star]$,

$$\mathbb{E}[N] \leq p(K=1) N_{NC} \tag{2.96}$$

$$+ \sum_{k=2}^{\infty} p(K=k) \left(\sum_{n_B=(k-1)n^\star+1}^{kn^\star} \left(\sum_{n_A=k}^{\infty} p(n_A \wedge n_B|k) \max(n_A, kn^\star) \right) \right) . \tag{2.97}$$

We have upper bounded the first term of Eq. (2.95) by upper bounding the sum $\sum_{n_B=1}^{n^\star}$ with $\sum_{n_B=1}^{\infty}$. In this case the expression after $p(K=1)$ in the first term becomes N_{NC}. In the second term we upper bound n_B by $kn^\star$. Since the second term does not depend on n_B anymore we upper bound it by removing the constraints on N_B completely from the probabilities $p(n_A \wedge n_B|K=k)$, i.e.

$$\mathbb{E}[N] \leq p(K=1) N_{NC} + \sum_{k=2}^{\infty} p(K=k) \sum_{n_A=k}^{\infty} p(n_A|K=k) \max(n_A, kn^\star) \tag{2.98}$$

$$= p(K=1) N_{NC} + \sum_{k=2}^{\infty} p(K=k) \left(\sum_{n_A=k}^{kn^\star} p(n_A|K=k) kn^\star + \sum_{n_A=kn^\star+1}^{\infty} p(n_A|K=k) n_A \right) , \tag{2.99}$$

where in the last line of Eq. (2.99) we split the second term into two terms corresponding to the regime where $kn^\star$ is larger than n_A and vice versa. Since $kn^\star$ does not depend on n_A, we upper bound this term by removing the constraints on n_A,

$$\mathbb{E}[N] \leq p(K=1) N_{NC} + \sum_{k=2}^{\infty} p(K=k) kn^\star + \sum_{k=2}^{\infty} p(K=k) \sum_{n_A=k}^{\infty} p(n_A|K=k) n_A . \tag{2.100}$$

Eq. (2.100) can be greatly simplified. We can perform the sum over n_A in the third term obtaining $k\mu$. The sums over k can then be easily evaluated so that the right hand side of Eq. (2.100) can be rewritten as

$$p(K=1) N_{NC} + \sum_{k=2}^{\infty} p(K=k) kn^\star + \sum_{k=2}^{\infty} p(K=k) k\mu \tag{2.101}$$

$$= p(K=1) N_{NC} + (n^\star + \mu)(\mathbb{E}(K) - p(K=1))$$

$$= \left(p_r + \frac{n^\star + \mu}{N_{NC}} \left(\frac{1}{p_r} - p_r \right) \right) N_{NC} \tag{2.102}$$

$$= \left(p_r + \left(\frac{n^\star + \mu}{N_{NC}} \right) \left(\frac{1 - p_r^2}{p_r} \right) \right) N_{NC} . \tag{2.103}$$

Hence we have that

$$N_{NC} \leq \mathbb{E}[N] \leq \left(\tilde{g}_{\text{err}}(p_A, p_B, n^\star) + 1\right) N_{NC} \,, \tag{2.104}$$

where $\tilde{g}_{\text{err}}(p_A, p_B, n^\star)$ is defined as

$$\tilde{g}_{\text{err}}(p_A, p_B, n^\star) := (1-p_B)^{n^\star} \left[\left(\frac{n^\star + \mu}{N_{NC}} \right) \left(\frac{2-(1-p_B)^{n^\star}}{1-(1-p_B)^{n^\star}} \right) - 1 \right] . \tag{2.105}$$

We now show that $\tilde{g}_{\text{err}}(p_A, p_B, n^\star)$ is small compared to the other quantities in Eq. (2.104). Observe that

$$(1-p_B)^{n^\star} = \left(1 - \frac{\beta'}{n^\star}\right)^{n^\star} \leq e^{-\beta'} \,. \tag{2.106}$$

From Eq. (2.105) it follows that

$$\tilde{g}_{\text{err}}(p_A, p_B, n^\star) \leq e^{-\beta'} \left[\frac{n^\star + \frac{1}{p_A}}{N_{NC}} \left(\frac{2}{1-e^{-\beta'}} \right) - 1 \right] . \tag{2.107}$$

To upper bound the relative error, we start by upper bounding the first term inside the brackets, namely

$$\begin{aligned} \frac{n^\star + \frac{1}{p_A}}{N_{NC}} &= \frac{n^\star + \frac{1}{p_A}}{\frac{1}{p_A} + \frac{1}{p_B} - \frac{1}{p_A + p_B - p_A p_B}} \\ &\leq \frac{n^\star + \frac{1}{p_A}}{\frac{1}{p_A} + \frac{1}{p_B} - \frac{1}{p_A + p_B - p_A}} = p_A n^\star + 1 \,. \end{aligned} \tag{2.108}$$

$\tilde{g}_{\text{err}}(p_A, p_B, n^\star)$, then, is upper bounded by

$$\tilde{g}_{\text{err}}(p_A, p_B, n^\star) \leq e^{-\beta'} \left[(p_A n^\star + 1) \left(\frac{2}{1-e^{-\beta'}} \right) - 1 \right] \tag{2.109}$$

$$= \frac{e^{-\beta'}}{1-e^{-\beta'}} (2 p_A n^\star + 1 + e^{-\beta'}) \tag{2.110}$$

$$\leq \frac{e^{-\beta'}}{1-e^{-\beta'}} (2\beta' + 1 + e^{-\beta'}) \tag{2.111}$$

$$= e^{-\beta'} \left(\frac{2\beta'}{1-e^{-\beta'}} + \coth\left(\frac{\beta'}{2}\right) \right) \tag{2.112}$$

$$< e^{-\beta'} \left(\frac{2\beta'}{1-e^{-1}} + \coth\left(\frac{1}{2}\right) \right) \tag{2.113}$$

$$< e^{-\beta'} \coth\left(\frac{1}{2}\right) (2\beta' + 1) \tag{2.114}$$

$$< 3 \coth\left(\frac{1}{2}\right) \beta' e^{-\beta'} \,. \tag{2.115}$$

Therefore $\tilde{g}_{\text{err}}(p_A, p_B, n^\star) = \mathcal{O}\left(\beta' e^{-\beta'}\right)$.

3

NEAR-TERM QUANTUM-REPEATER EXPERIMENTS WITH NV CENTRES

Filip Rozpędek*, Raja Yehia*, Kenneth Goodenough*, Maximilian Ruf, Peter Humphreys, Ronald Hanson, Stephanie Wehner and David Elkouss

Quantum channels enable the implementation of communication tasks inaccessible to their classical counterparts. The most famous example is the distribution of secret keys. However, in the absence of quantum repeaters the rate at which these tasks can be performed is dictated by the losses in the quantum channel. In practice, channel losses have limited the reach of quantum protocols to short distances. Quantum repeaters have the potential to significantly increase the rates and reach beyond the limits of direct transmission. However, no experimental implementation has overcome the direct transmission threshold. Here, we propose three quantum repeater schemes and assess their ability to generate secret key when implemented on a setup using NV centres in diamond with near-term experimental parameters. We find that one of these schemes - the so-called single-photon scheme, requiring no quantum storage - has the ability to surpass the capacity - the highest secret-key rate achievable with direct transmission - by a factor of seven for a distance of approximately 9.2 kilometres with near-term parameters, establishing it as a prime candidate for the first experimental realisation of a quantum repeater.

---

*These authors contributed equally. K. Goodenough contributed with implementing the code, analysing the duration of the proof-of-principle experiment, and co-writing the manuscript.
This chapter has been adapted from the following publication: Phys. Rev. A **99**, 052330 (2019)

3.1. INTRODUCTION

This chapter is a continuation of the investigation of proof-of-principle repeater setups in the previous chapter. Unlike the previous chapter, we will now investigate four schemes at the same time, as opposed to only one. As before, we will analyse the ability of each of these repeater schemes to generate secret-key. We focus once again on nitrogen-vacancy centres in diamond (NV), due to their properties lending themselves well to long-distance quantum communication tasks [1, 14, 16, 35, 55, 72, 83, 135, 158, 165].

Besides an explicit focus on four different NV-based setups, there are two other differences with the previous chapter we will now highlight beforehand. First, we will consider a smaller number of benchmarks, and include a new benchmark for one scheme in particular. Second, we now take a more conservative approach to counting the number of channel uses for our calculation of the yield. That is, we now take the sum of the channel uses on Alice's and Bob's side, which can be calculated exactly

The four considered schemes are: the single sequential quantum repeater node (first proposed and studied in [99], then further analysed in the previous chapter and [140]), the single-photon scheme (proposed originally in the context of remote entanglement generation [23], also studied in the context of secret-key generation without quantum memories [98]), and two schemes which are a combination of the first two. See Fig. 3.1 for a schematic overview of the repeater proposals considered in this chapter.

We show that one of these schemes, the *single-photon scheme*, can surpass the secret-key capacity by a factor of seven for a distance of ≈ 9.2 km with near-term parameters. This shows the viability of this scheme for the first experimental implementation of a quantum repeater.

In Section 3.2 we discuss and detail the different repeater proposals that will be assessed in this chapter. In Section 3.3 we expand on how the different components of the repeater proposals would be implemented experimentally. Section 3.4 details how to calculate the secret-key rate achieved with the quantum repeater proposals from the modelled components. In Section 3.5 we discuss how to assess the performance of a quantum repeater. The comparison of the different repeater proposals is performed in Section 3.6, which allows us to conclude with our results in Section 3.7. The numerical results of this article were produced with a Python and a Mathematica script, which are available upon request.

3.2. QUANTUM REPEATER SCHEMES

In the following section we present the quantum repeater schemes that will be assessed in this chapter. All these schemes use NV centre based setups which involve memory nodes consisting of an electron spin qubit acting as an optical interface and possibly an additional carbon ^{13}C nuclear spin qubit acting as a long-lived quantum memory. Specifically, the optical interface of the electron spin allows for the generation of spin-photon entanglement, where the photonic qubits can then be transmitted over large distances. The carbon nuclear spin acts as a long-lived memory, but can be accessed only through the interaction with the electron spin. Here, we briefly go over all the new proposed schemes, motivate why they are interesting from an experimental perspective and discuss their advantages and disadvantages. The first of these schemes was already studied in the previous chapter, and we will thus not discuss the implementation further here. From hereon, we will denote this scheme as a Single Sequential Quantum Repeater (SiSQuaRe).

3.2.1. THE SINGLE-PHOTON SCHEME

Cabrillo et al. [23] devised a procedure that allows for the heralded generation of entanglement between a separated pair of matter qubits (their proposal discusses specific implementation with single atoms, but the scheme can also be applied to other platforms such as NV centres or quan-

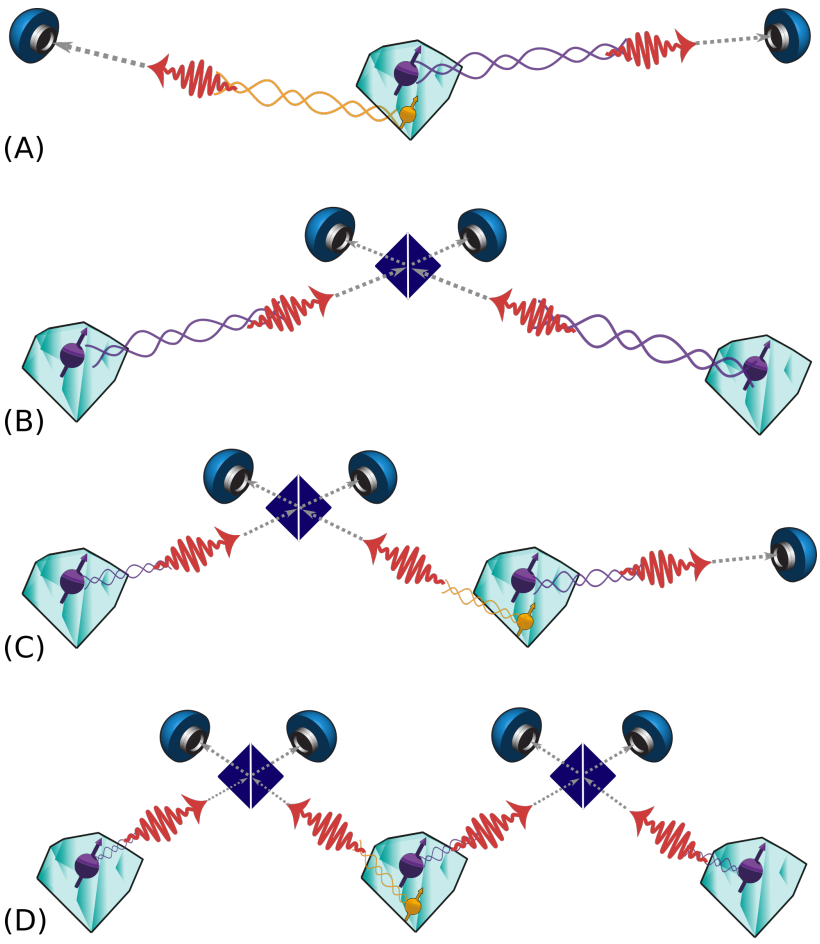

Figure 3.1: Schematic overview of the four quantum repeater schemes assessed in this chapter. From top to bottom: the Single Sequential Quantum Repeater (SiSQuaRe) scheme (A), the single-photon scheme (B), the Single-Photon with Additional Detection Setup (SPADS) scheme (C) and the Single-Photon Over Two Links (SPOTL) scheme (D). The purple particles represent NV electron spins capable of emitting photons (red wiggly arrows) while the yellow particles represent carbon ^{13}C nuclear spins. Dark blue squares depict the beam splitters used to erase the which-way information of the photons, followed by blue photon detectors. For more details on the different proposals, see Section 3.2.

tum dots) using linear optics. For the atomic ensemble platform this scheme also forms a building block of the DLCZ quantum repeater scheme [45]. Here we will refer to this scheme as a single-photon scheme as the entanglement generation is heralded by a detection of only a single photon. This requirement of successful transmission of only a single photon from one node makes it possible for this scheme to qualify as a quantum repeater (see below for more details).

The basic setup of the single-photon scheme consists of placing a beam splitter and two detectors between Alice and Bob, with both parties simultaneously sending a photonic quantum state towards the beam splitter. The transmitted quantum state is entangled with a quantum memory, and the state space of the photon is spanned by the two states corresponding to the presence and absence of a photon. Immediately after transmitting their photons through the fibre, both Alice and Bob measure their quantum memories in a BB84 or six-state basis (see the discussion of which quantum key distribution protocol is optimal for each scheme in Section 3.4.2 and in Section 3.6.1). Note that this is equivalent to preparing a specific state of the photonic qubit and therefore is closely linked to the measurement device independent quantum key distribution (MDI QKD) [96] as discussed in Appendix 3.8.9. However, preparing specific states that involve the superposition of the presence and absence of a photon on its own is generally experimentally challenging. The NV-implementation allows us to achieve this task precisely by preparing spin-photon entanglement and then measuring the spin qubit. Afterwards, by conditioning on the click of a single detector only, Alice and Bob can use the information of which detector clicked to generate a single raw bit of key, see Appendix 3.8.5 and [23] for more information.

The main motivation of this scheme is that, informally, we only need one photon to travel half the distance between the two parties to get an entangled state. This thus effectively reduces the effects of losses, and in the ideal scenario the secret-key rate would scale with the square root of the total transmissivity η, as opposed to linear scaling in η (which is the optimal scaling without a quantum repeater [129]).

However, one problem that one faces when implementing this scheme is that the fibre induces a phase shift on the transmitted photons. This shift can change over time, e.g. due to fluctuations in the temperature and vibrations of the fibre. The uncertainty of the phase shift induces dephasing noise on the state, reducing the quality of the state.

To overcome this problem, a two-photon scheme was proposed by Barrett and Kok [7], which does not place such high requirement on the optical stability of the setup. Specifically, in the Barrett and Kok scheme the problem of optical phase fluctuations is overcome by requiring two consecutive clicks and performing additional spin flip operations on both of the remote memories. The Barrett and Kok scheme has seen implementation in many experiments [13, 65, 66, 106]. However, the requirement of two consecutive clicks implies that a setup using only the Barrett and Kok scheme with two memory nodes will never be able to satisfy the demands of a quantum repeater. Specifically, the probability of getting two consecutive clicks will not be higher than the transmissivity of the fibre between the two parties and therefore will not surpass the secret-key capacity.

In the single-photon scheme, on the other hand, the dephasing caused by the unknown optical phase shift is overcome by using active *phase-stabilisation* of the fibre to reduce the fluctuations in the induced phase. This technique has been used in the experimental implementations of the single-photon scheme for remote entanglement generation using quantum dots [40, 154], NV centres [72] and atomic ensembles [31]. For experimental details relating to NV-implementation, we refer the reader to Section 3.3. This phase-stabilisation technique effectively reduces the optical phase uncertainty $\Delta\phi$, allowing us to significantly mitigate the resulting dephasing noise, see Appendix 3.8.1 for mathematical details.

In contrast to the Barrett and Kok scheme, the single-photon scheme cannot produce a perfect maximally entangled state, even in the case of perfect operations and perfect phase-stabilisation.

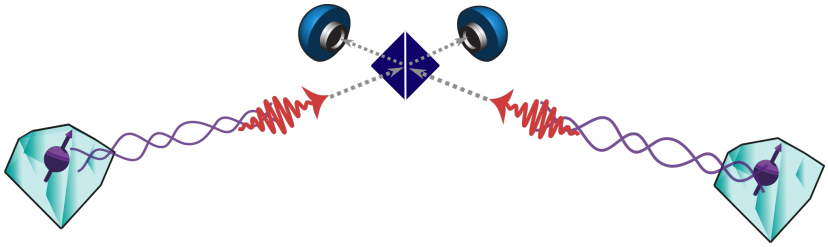

Figure 3.2: Schematic overview of the single-photon scheme. Alice and Bob simultaneously transmit a photonic state from their NV centres towards a balanced beam splitter in the centre. This photonic qubit, corresponding to the presence and absence of a photon, is initially entangled with the NV electron spin. If only one of the detectors (which can be seen at the top of the figure) registers a click, Alice and Bob can use the information of which detector clicked to generate a single raw bit of key.

This is because losses in the channel result in a significant probability of having both nodes emitting a photon which can also lead to a single click in one of the detectors, yet the memories will be projected onto a product state. As we discuss below, this noise can be traded versus the probability of success of the scheme by reducing the weight of the photon-presence term in the generated spin-photon entangled state. This is discussed in more detail below and the full analysis is presented in Appendix 3.8.5.

The single-photon scheme with phase-stabilisation is a promising candidate for a near-term quantum repeater with NV centres. We note here that recently other QKD schemes that use the MDI framework have been proposed. These *twin-field QKD schemes*, similarly to our proposal, use single-photon detection events to overcome the linear scaling of the secret-key rate with η [98, 102, 157]. In these so-called twin-field QKD proposals, in contrast to our single-photon scheme, no quantum memories are used, but instead Alice and Bob send phase-randomised optical pulses to the middle heralding station.

SETUP AND SCHEME

In the setup of the single-photon scheme Alice and Bob are separated by a fibre where in the centre there is a beam splitter with two detectors (see Fig. 3.2). They will both create entanglement between a photonic qubit and a stored spin and send the photonic qubit to the beam splitter.

Alice and Bob thus perform the following,

1. Alice and Bob both prepare a state $|\psi\rangle = \sin\theta\,|{\downarrow}\rangle|0\rangle + \cos\theta\,|{\uparrow}\rangle|1\rangle$ where $|{\downarrow}\rangle/|{\uparrow}\rangle$ refers to the dark/bright state of the electron-spin qubit, $|0\rangle/|1\rangle$ indicates the absence/presence of a photon, and θ is a tunable parameter.
2. Alice and Bob attempt to both separately send the photonic qubit to the beam splitter.
3. Alice and Bob both perform a BB84 or six-state measurement on their memories.
4. The previous steps are repeated until only one of the detectors between the parties clicks.
5. The information of which detector clicked gets sent to Alice and Bob for classical correction.
6. All the previous step are repeated until sufficient data have been generated.

The parameter θ can be chosen by preparing a non-uniform superposition of the dark and bright state of the electron spin $|\psi\rangle = \sin\theta\,|{\downarrow}\rangle + \cos\theta\,|{\uparrow}\rangle$ via coherent microwave pulses. This is done before applying the optical pulse to the electron which entangles it with the presence and absence of a photon. The parameter θ can then be tuned in such a way as to maximise the secret-key rate. In the next section, we will briefly expand on some of the issues arising when losses and imperfect detectors are present. We defer the full explanation and calculations until Appendix 3.8.5.

REALISTIC SETUP

In any realistic implementation of the single-photon scheme, a large number of attempts is needed before a photon detection event is observed. Furthermore, a single detector registering a click does not necessarily mean that the state of the memories is projected onto the maximally entangled state. This is due to multiple reasons, such as losing photons in the fibre or in some other loss process between the emission and detection, arrival of the emitted photons outside of the detection time-window and the fact that dark counts generate clicks at the detectors. Photon loss in the fibre effectively acts as amplitude-damping on the state of the photon when using the presence/absence state space [74, 131]. Dark counts are clicks in the detectors, caused by thermal excitations. These clicks introduce noise, since it is impossible to distinguish between clicks caused by thermal excitations and the photons traveling through the fibre if they arrive in the same time-window. All these sources of loss and noise acting on the photonic qubits are discussed in detail in Appendix 3.8.1. Finally we note that we assume here the application of non-number resolving detectors. This can lead to additional noise in the low loss regime, since the event in which two photons got emitted cannot be distinguished from the single-photon emission events even if no photons got lost. However, in any realistic loss regime this is not a problem, since the probability of two such photons arriving at the heralding station is quadratically suppressed with respect to events where only one photon arrives. In the realistic regime, almost all the noise coming from the impossibility of distinguishing two-photon from single-photon emission events is the result of photon loss. Namely, if a two-photon emission event occurs and the detector registers a click, then with dominant probability it is due to only a single photon arriving, while the other one being lost. Hence the use of photon-number resolving detectors would not give any visible benefit with respect to the use of the non-number resolving ones. For a detailed calculation of the effects of losses and dark counts for the single-photon scheme, see Appendix 3.8.5.

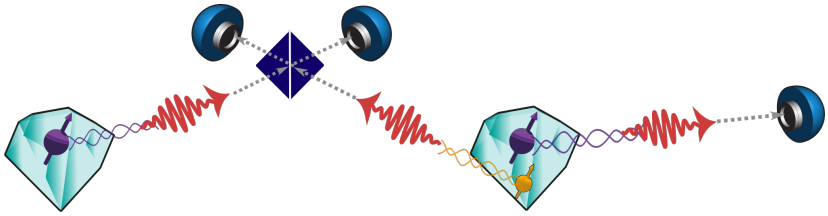

Figure 3.3: Schematic overview of the SPADS scheme. First, the two NV centres run the single-photon scheme, such that Alice measures her electron spin directly after every attempt. After success, the middle node swaps its state to the carbon spin. Then the middle node generates electron-photon entangled pairs where the photonic qubit is encoded in the time-bin degree of freedom and sent to Bob. This is attempted until Bob successfully measures the photon or until the cut-off is reached. If the cut-off is reached, the scheme gets restarted, otherwise the middle node performs an entanglement swapping on its two memories and communicates the classical outcome to Alice and Bob, who can correct their measurement outcomes to obtain a bit of raw key.

3.2.2. SINGLE-PHOTON WITH ADDITIONAL DETECTION SETUP (SPADS) SCHEME

The third scheme that we consider here is the Single-Photon with Additional Detection Setup (SPADS) scheme, which is effectively a combination of the single-photon scheme and the SiSQuaRe scheme as shown in Fig. 3.3. If the middle node is positioned at two-thirds of the total distance away from Alice, the rate of this setup would scale, ideally, with the cube root of the transmissivity η.

This scheme runs as follows:

1. Alice and the repeater run the single-photon scheme until success, however, only Alice performs her spin measurement immediately after each spin-photon entanglement generation attempt. This measurement is either in a six-state or BB84 basis.
2. The repeater swaps the state of the electron spin onto the carbon spin.
3. The repeater runs the second part of the SiSQuaRe scheme with Bob. This means it generates spin-photon entanglement between an electron and the time-bin encoded photonic qubit. Afterwards, it sends the photonic qubit to Bob. This is repeated until Bob successfully measures his photon in a six-state or BB84 basis or until the cut-off n^* is reached in which case the scheme is restarted with step 1.
4. After Bob has received the photon and communicated this to the repeater, the repeater performs a Bell-state measurement on its two quantum memories and communicates the classical result to Bob.
5. All the previous steps are repeated until sufficient data have been generated.

The motivation for introducing this scheme is two-fold. Firstly, we note that by using this scheme we divide the total distance between Alice and Bob into three segments: two segments corresponding to the single-photon subscheme and the third segment over which the time-bin encoded photons are sent. This gives us one additional independent segment with respect to the single-photon or the SiSQuaRe scheme on their own. Hence, for distances where no cut-off is required, we expect the scaling of the secret-key rate with the transmissivity to be better than the ideal square root scaling of the previous two schemes. Furthermore, dividing the total distance into more segments should also allow us to reach larger distances before dark counts become significant. When considering the resources necessary to run this scheme, we note that the additional third node needs to be equipped only with a photon detection setup.

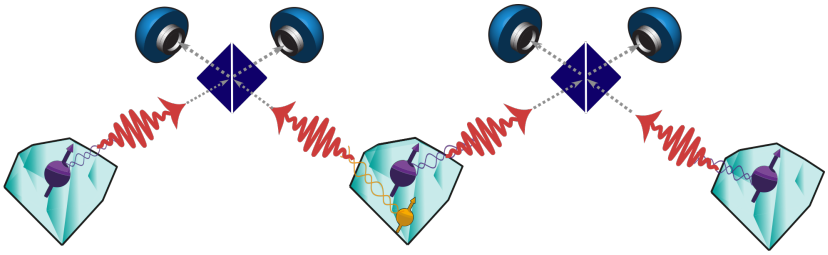

Figure 3.4: Schematic overview of the setup for the SPOTL scheme. This scheme is a combination of the SiSQuaRe and single-photon scheme. Instead of sending photons directly through the fibre as in the SiSQuaRe scheme, entanglement is established between the middle node and Alice/Bob using the single-photon scheme.

Secondly, we note that the SPADS scheme can also be naturally compared to the scenario in which an NV centre is used as a single photon source for direct transmission between Alice and Bob. Both the setup for the SPADS scheme and such direct transmission involve Alice using an NV for emission and Bob having only a detector setup. Hence, the SPADS scheme corresponds to inserting a new NV-node (the repeater) between Alice and Bob without changing their local experimental setups at all. This motivates us to compare the achievable secret-key rate of the SPADS scheme and direct transmission. We perform this comparison on a separate plot in Section 3.6.

3.2.3. SINGLE-PHOTON OVER TWO LINKS (SPOTL) SCHEME

The final scheme that we study here is the Single-Photon Over Two Links (SPOTL) scheme, and it is another combination of the single-photon and SiSQuaRe schemes. A node is placed between Alice and Bob which tries to sequentially generate entanglement with their quantum memories by using the single-photon scheme (see Fig. 3.4). The motivation for this scheme is that, while using relatively simple components and without imposing stricter requirement on the memories than in the previous schemes, its secret-key rate would ideally scale with the fourth root of the transmissivity η.

SETUP AND SCHEME

The setup that we study is the following:

1. Alice and the repeater run the single-photon scheme until success with the tunable parameter $\theta = \theta_A$. However, only Alice performs her spin measurement immediately after each spin-photon entanglement generation attempt. This measurement is in a BB84 or six-state basis.
2. The repeater swaps the state of the electron spin onto the carbon spin.
3. Bob and the repeater run the single-photon scheme until success or until the cut-off n^* is reached in which case the scheme is restarted with step 1. The tunable parameter is set here to $\theta = \theta_B$. Again, only Bob performs his spin measurement immediately after each spin-photon entanglement generation attempt and this measurement is in a six-state basis.
4. The quantum repeater performs a Bell-state measurement and communicates the result to Bob.
5. All the previous steps are repeated until sufficient data have been generated.

We note that for larger distances the optimal cut-off becomes smaller. Then, since we lose the independence of the attempts on both sides, the scaling of the secret-key rate with distance is

expected to drop to $\sqrt{\eta}$, which is the same as for the single-photon scheme. However, the total distance between Alice and Bob is now split into four segments. Alice and Bob thus send photons over only one fourth of the total distance. Thus, this scheme should be able to generate key over much larger distances than the previous ones, as the dark counts will start becoming significant for larger distances only.

3.3. NV-IMPLEMENTATION

Having proposed the different quantum repeater schemes, we now move on to describe their experimental implementation based on nitrogen-vacancy centers in diamond [43]. Since most of the components found in the schemes are the same as in the SiSQuaRe scheme discussed in the previous chapter, we will only discuss those components and operations that are not used in the SiSQuaRe scheme.

By applying selective optical pulses and coherent microwave rotations, we first generate spin-photon entanglement at an NV center node [13]. To generate entanglement between two distant NV electron spins, these emitted photons are then overlapped on a central beam splitter to remove their which-path information. Subsequent detection of a single photon heralds the generation of a spin-spin entangled state [13]. For all schemes based on single-photon entanglement generation, we need to employ active phase-stabilisation techniques to compensate for phase shifts of the transmitted photons, which will reduce the entangled state fidelity, as introduced in Section 3.2.1. These fluctuations arise from both mechanical vibrations and temperature induced changes in optical path length, as well as phase fluctuations of the lasers used during spin-photon entanglement generation. This problem can be mitigated by using light reflected off the diamond surface to probe the phase of an effectively formed interferometer between the two NV nodes and the central beam splitter, and by feeding the acquired error signal back to a fibre stretcher that changes the relative optical path length [72].

3.4. CALCULATION OF THE SECRET-KEY RATE

With the modeling of each of the components of the different setups in hand, the performance of each setup can be estimated. The performance of a setup is again assessed in this chapter by its ability to generate secret key between two parties Alice and Bob.

The secret-key rate R is equal to

$$R = \frac{Y \cdot r}{N_{\text{modes}}}, \tag{3.1}$$

where Y and r are the yield and secret-key fraction, respectively. The yield Y is defined as the average number of raw bits generated per channel use and the secret-key fraction r is defined as the amount of secret key that can be extracted from a single raw bit (in the limit of asymptotically many rounds). Here N_{modes} is the number of optical modes needed to run the scheme. Time-bin encoding requires two modes while the single-photon scheme uses only one mode. Hence $N_{\text{modes}} = 2$ for all the schemes that use time-bin encoding in at least one of the arms of the setup (such as in the previous chapter). For the schemes that use only the single-photon subschemes as their building blocks we have that $N_{\text{modes}} = 1$.

In the remainder of this section, we will briefly detail how to calculate the yield and secret-key fraction, from which we can estimate the secret-key rate of each scheme. As in the previous chapter, we consider here a *simulation scenario*, i.e. the scenario where there is no eavesdropper. As before, this does not impact the security of any of the statements we make with regards to the considered implementations.

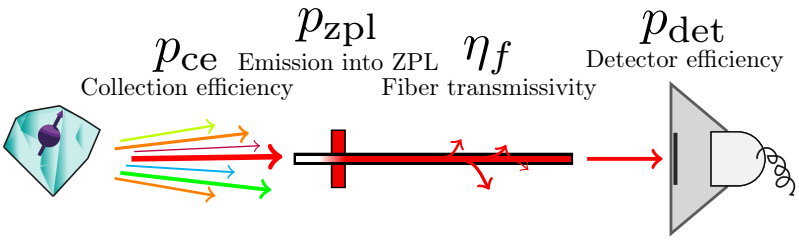

Figure 3.5: The model of photon loss proccesses occurring in our repeater setups. The parameter p_{ce} is the photon collection efficiency, which includes the probability that the photon is successfully coupled into the fibre. Only photons emitted at the zero phonon line (ZPL) can be used for quantum information processing. All non-ZPL photons are filtered out, such that a fraction p_{zpl} of the photons remains. The photons are then transmitted through a fibre with transmissivity η_f. Such successful transmissions are registered by the detector with probability p_{det}. Additionally, a significant fraction of photons can arrive in the detector outside of the detection time-window t_{w}. Such photons will effectively also get discarded. Here we describe the total efficiency of our apparatus by a single parameter, $p_{\text{app}} = p_{\text{ce}} p_{\text{zpl}} p_{\text{det}}$.

3.4.1. YIELD

The yield depends not only on the used scheme, but also on the losses in the system. We model the general emission and transmission of photons through fibres similar to the previous chapter. However, unlike the previous chapter, we focus now on NV centre setups exclusively and add a variable detection time-window for the photon detectors. This motivates us to use the NV centre specific terminology for this chapter.

Fig. 3.5 contains our model of photon losses. with probability p_{ce} spin-photon entanglement is generated and the photon is coupled into a fibre. The photons that successfully got coupled into the fibre might not be useful for quantum information processing since they are not coherent. Thus, we filter out those photons that are not emitted at the zero-phonon line, reducing the number of photons by a further factor of p_{zpl}. Then, over the length of the fibre, a photon gets lost with probability $1-\eta_f = 1-e^{-\frac{L}{L_0}}$, where L_0 is the attenuation length and η_f is the transmissivity. After exiting the fibre the photon gets registered as a click by the detector with probability p_{det}. Finally, the photon gets accepted as a successful click if the click happens within the time-window t_{w} of the detector (see Appendix 3.8.1 for more details).

The yield can then be calculated as the reciprocal of the expected number of channel uses needed to get one single raw bit,

$$Y = \frac{1}{\mathbb{E}[N]}\,, \tag{3.2}$$

with N being the random variable that models the number of channel uses needed for generating a single raw bit.

YIELD OF THE SINGLE-PHOTON SCHEME

The yield of the single-photon scheme is relatively easy to calculate, since the single condition heralding the success of the scheme is a single click in one of the detectors in the heralding station. Therefore the yield Y is simply the probability that an individual attempt will result in a single click in one of the detectors. This probability will depend on the losses in the system, dark counts and the angle θ. A full calculation of the yield is given in Appendix 3.8.5.

YIELD OF THE SISQUARE, SPADS AND SPOTL SCHEMES

The SiSQuaRe, SPADS and SPOTL schemes require two conditions for the heralding of the successful generation of a raw bit, namely the scheme needs to succeed both on Alice's and Bob's side independently. In this case we are going to take a more conservative perspective than in the previous chapter, and assume the total number of channel uses to be the sum of the required channel uses on Alice's and Bob's side of the memory repeater node,

$$\mathbb{E}[N] = \mathbb{E}[N_A + N_B] \,. \tag{3.3}$$

This should be contrasted with the previous chapter, where we considered the maximum of the channel uses on Alice's and Bob's side. Considering the sum of N_A and N_B also better corresponds to the time used for the experiment. Moreover, every time Bob reaches n^* attempts, both parties start the scheme over again. The cut-off increases the average number of channel uses, thus decreasing the yield. Denoting by p_A and p_B the probability that a single attempt of the subscheme on Alice's and Bob's side respectively succeeds, we find (see Appendix 3.8.3 for the derivation),

$$\mathbb{E}[N_A + N_B] = \frac{1}{p_A\left(1-\left(1-p_B\right)^{n^*}\right)} + \frac{1}{p_B} \,. \tag{3.4}$$

3.4.2. SECRET-KEY FRACTION

We use the same two protocols for generating secret-key as in the previous chapter, namely BB84 with standard one-way error correction and six-state with advantage distillation [171]. As discussed in the previous chapter, it is not possible to run an asymmetric six-state protocol when time-bin encoded photons are used.

We now state explicitly which QKD protocols will be considered for each scheme, which in turn depends on the type of measurements that Alice and Bob perform in that scheme. There are two physical implementations of measurements that Alice and Bob perform, depending on the scheme under consideration. That is, they either measure a quantum state of a spin or of a time-bin encoded photons. Since the fully asymmetric six-state protocol with advantage distillation has higher efficiency than both symmetric and asymmetric BB84 protocol with one-way error correction, we will use this six-state protocol for both the single-photon and SPOTL scheme. The SiSQuaRe and SPADS schemes involve direct measurement on time-bin encoded photons. Hence, for these schemes we consider the maximum of the amount of key that can be obtained using the fully asymmetric BB84 protocol and the symmetric six-state protocol with advantage distillation (which can tolerate more noise, but has three times lower efficiency than the fully asymmetric BB84 protocol).

To estimate the QBER, we model all the noisy and lossy processes that take place during the protocol run. From this, we calculate the qubit error rates and yield, from which we can retrieve the secret-key fraction. We invite the interested reader to read about the details of these calculations in Appendices 3.8.5 and 3.8.6. The derivation of the QBER and the yield for the SiSQuaRe scheme can be found in the previous chapter. Moreover, in this chapter we introduce certain refinements to the model which we discuss in Appendix 3.8.4. With the QBER in hand, we can calculate the resulting secret-key fraction for the considered protocols as presented in the previous chapter and Appendix 3.8.7.

We note here that we consider only the secret-key rate in the asymptotic limit, and that we thus do not have to deal with non-asymptotic statistics.

3.5. ASSESSING QUANTUM REPEATER SCHEMES

In this section we will detail four benchmarks that will be used to assess the performance of quantum repeaters, similar as was done in the previous chapter.

The considered benchmarks are defined with respect to the efficiencies of processes involving photon loss when emitting photons at NV centres, transmitting them through an optical fibre and detecting them at the end of the fibre as described in Section 3.4.1 and as shown in Fig. 3.5.

Having this picture in mind, we can now proceed to present the considered benchmarks. The first three of these benchmarks are inspired by fundamental limits on the maximum achievable secret-key rate if Alice and Bob are connected by quantum channels which model quantum key distribution over optical fibre without the use of a (possible) quantum repeater.

The first of these benchmarks is the *capacity of the pure-loss channel* introduced in the previous chapter. The capacity of the pure-loss channel is the maximum achievable secret-key rate over a channel modeling a fibre of transmissivity η_f, and is given by [131]

As noted in the previous chapter, surpassing the capacity is experimentally challenging. This motivates the introduction of other, easier to surpass, benchmarks. These benchmarks are still based on (upper bounds on) the secret-key capacity of quantum channels which model realistic implementations of quantum communications over fibres. The usage of upper bounds on the secret-key capacity (instead of the secret-key capacity itself) for certain channels is due to the fact that the secret-key capacity of those channels is still unknown.

The second benchmark is, as before, built on the idea of including the losses of the apparatus into the transmissivity of the fibre. The resultant channel with all those losses included we call here *the extended channel*. The benchmark is thus equal to

$$-\log_2\left(1-\eta_f p_{\text{app}}\right). \tag{3.5}$$

Here p_{app} describes all the intrinsic losses of the devices used. That is, the collection efficiency p_{ce} at the emitting diamond, the probability that the emitted photon is within the zero-phonon-line p_{zpl} (which is necessary for generating quantum correlations) and photon detection efficiency p_{det}, so that $p_{\text{app}} = p_{\text{ce}}\,p_{\text{zpl}}\,p_{\text{det}}$. This should be contrasted with the previous chapter, where the post-selection success probability p_{ps} was not included in the intrinsic losses of the device, while the corresponding probability of emitting into the zero-phonon-line p_{zpl} has been included.

The third benchmark is, as in the previous chapter, the *thermal channel bound*, which takes into account the effects of dark counts. As in the second benchmark, this benchmark is different from the benchmark from the previous chapter in the sense that p_{zpl} (p_{ps}) is included in the losses. We note here that the time-window of the detector t_{w} is not fixed in our model, but is optimised over for every distance in order to achieve the highest possible secret-key rate. Hence in this benchmark we fix $t_{\text{w}} = 5$ ns which is the shortest duration of the time-window that we consider in our secret-key rate optimisation.

Finally, the secret-key rate achieved with *direct transmission using the same devices* can be seen as **the fourth benchmark**. Specifically, here we mean the secret-key rate achieved when Alice uses her electron spin to generate spin-photon entanglement and sends the time-bin encoded photon to Bob. She then measures her electron spin while Bob measures the arriving photon. However, to take a conservative view, we will only use this direct transmission benchmark for the SPADS scheme. This is motivated by the fact that for both the SPADS scheme and the direct transmission scheme the experimental setups on Alice's and Bob's side are the same, ensuring that the two rates can be compared fairly. We note that similarly as in the modelled secret-key rates achievable with our proposed repeater schemes, also for this direct transmission benchmark we optimise over the time-window t_{w} for each distance.

The secret-key capacity is the main benchmark that we consider. Surpassing it establishes the considered scheme as a quantum repeater. The remaining benchmarks further guide the way

Parameter	Notation	Value
Depolarising parameter for electron measurement	F_m	0.95 [72]
Depolarising parameter for two qubit gates	F_g	0.98 [83]
Characteristic time of the NV emission	τ	6.48 ns [53, 138]
Detection window offset	t_w^{offset}	1.28 ns [65]
Optical phase uncertainty of the spin-spin state	$\Delta\phi$	14.3° [72]

Table 3.1: Additional parameters used for the nitrogen-vacancy centre setups considered in this chapter.

towards implementation of a quantum repeater. We define all the considered benchmarks for the channel with the same fibre attenuation length L_0 as the channel used for the corresponding achievable secret-key rate.

3.6. NUMERICAL RESULTS

We now have a full model of the rate of the presented quantum repeater protocols as a function of the underlying experimental parameters. In this section we will firstly state all the parameters required by our model and then present the results and conclusions drawn from the numerical implementation of this model. In particular, in Section 3.6.1 we will first provide a deeper insight into the benefits of using the six-state protocol and advantage distillation in specific schemes. In Section 3.6.2 we determine the optimal positioning of the repeater nodes for our schemes and investigate the dependence of the secret-key rate achievable with those schemes on the photon emission angle θ and the cutoff n^* for the appropriate schemes. In Section 3.6.3 we then use the insights acquired in the previous section to compare the achievable secret-key rates for all the proposed repeater schemes with the secret-key capacity and other proposed benchmarks. In particular, we show that the single-photon scheme significantly outperforms the secret-key capacity and hence can be used to demonstrate a quantum repeater. Finally, in Section 3.6.4 we determine the duration of the experiment that would allow us to demonstrate such a quantum repeater with the single-photon scheme.

We will use the same parameters as in the previous chapter, but with some additional changes. These modifications are due to the model presented here being more fine-grained (see 3.8.4) and that there are certain parameters relevant only for the single-photon scheme. Furthermore, certain parameters from the previous chapter have been renamed to fit their NV description, while retaining the same value. That is, p_{em} and p_{ps} have been renamed to p_{ce} and p_{zpl}, respectively. Furthermore, the detector time-window t_w is now not fixed, but optimised over to maximise the secret-key rate. To reiterate, the parameters that we use are either parameters that have been achieved in an experiment, or correspond to expected parameters when the NV centre is embedded in an optical Fabry-Perot microcavity. The parameters we will use are listed in Table 3.1.

The parameters that have not been discussed in the main text are discussed in the appendix.

3.6.1. COMPARING BB84 AND SIX-STATE ADVANTAGE DISTILLATION PROTOCOLS

We first investigate here when the BB84 or six-state advantage distillation protocol performs better. Advantage distillation is a specific way of post-processing the gathered data, allowing for a higher secret-key rate. It was shown in the previous chapter that in the SiSQuaRe scheme there is a trade-off - for the low noise regime (small distances) the fully asymmetric BB84 protocol is preferable, while in the high noise regime (large distances) the problem of noise can be overcome by using a six-state protocol supplemented with advantage distillation. This technique allows us to increase

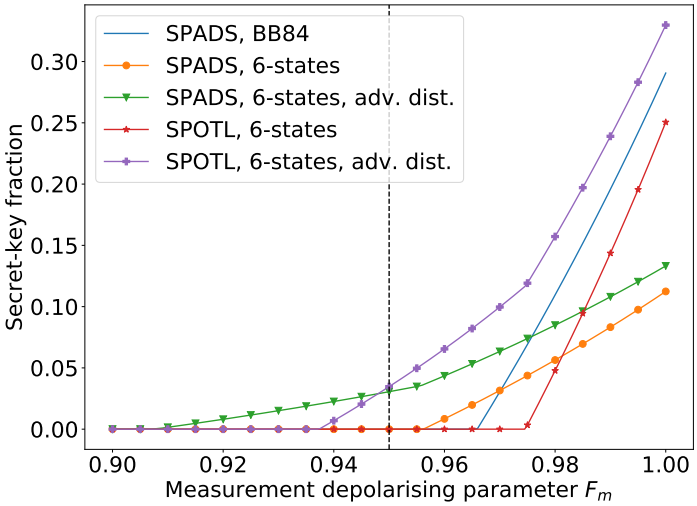

Figure 3.6: Secret-key fraction as a function of the depolarising parameter due to noisy measurement F_m for the total distance of $12.5L_0$. We see that for the current experimental value of $F_m = 0.95$ (marked with a dashed black vertical line) both schemes can generate key only if the advantage distillation post-processing is used. As F_m increases the protocols that do not utilise advantage distillation also start generating key. We also see that the curves can be divided into two groups in terms of their slope in the regime where they generate non-zero amount of key. Those two groups correspond to the scenarios where a fully asymmetric (bigger slope) or a symmetric (smaller slope) protocol is used. For all the plotted protocols the cutoff n^* is set to one and $t_w = 5$ ns (the smallest detection time-window we use) to maximize the secret-key fraction. Moreover, for each value of F_m we optimize the secret-key fraction over the angle θ. For the SPOTL scheme we assume $\theta_A = \theta_B$. For the SPADS scheme we position the repeater node 2/3 away of the total distance from Alice and in the middle between Alice and Bob for the SPOTL scheme.

the secret-key fraction at the expense of reducing the yield by a factor of three, since a six-state protocol in which Alice and Bob perform measurements on photonic qubits does not allow for the (fully) asymmetric protocol within our model. Numerically, we find that for the SPADS and SPOTL scheme advantage distillation is *necessary* to generate non-zero secret-key at any distance. This is due to the fact that there is a significant amount of noise in these schemes. Thus, for the SPADS (SPOTL) scheme the (a)symmetric six-state protocol with advantage distillation is optimal.

To provide more insight into the performance of those different QKD schemes for different parameter regimes, we plot the achievable secret-key fraction for the SPADS and SPOTL schemes as a function of the depolarising parameter due to imperfect electron spin measurement F_m in Figure 3.6 (see Appendix 3.8.2 for the discussion of the corresponding noise model). Noise due to imperfect measurements is one of the significant noise sources in our setup, since the SPADS scheme involves three and the SPOTL scheme four single-qubit measurements on the memory qubits. The data have been plotted for a fixed distance of $12.5L_0$, where $L_0 = 0.542$ km is the attenuation length of the fibre. Moreover, since on this plot we aim at maximising only the secret-key fraction over the tunable parameters, we set the cutoff n^* to one and the detection time-window t_w to 5 ns (the smallest detection time-window we use) for both schemes. Furthermore, within the single-photon subscheme the heralding station is always placed exactly in the middle between the two memory nodes. We also consider the positioning of the memory repeater node to be two-thirds away from Alice for the SPADS scheme and in the middle for the SPOTL scheme as discussed in the next section. For the SPOTL scheme we also assume $\theta_A = \theta_B$ which we will justify in the next section.

We see that for the current experimental value of $F_m = 0.95$ both schemes can generate key only if the advantage distillation post-processing is used. As F_m increases, we observe that for the SPADS scheme firstly the six-state protocol without advantage distillation and then the BB84 protocol start generating key. For the SPOTL scheme the value of F_m at which the six-state protocol without advantage distillation starts generating key is much larger than the corresponding value of F_m for any of the studied protocols for the SPADS scheme. This is because the SPOTL scheme involves more noisy processes than the SPADS scheme. This also provides an approximate quantification of the benefit of using advantage distillation. Specifically, looking at the SPOTL scheme, it can be observed that while at the current experimental value of $F_m = 0.95$ advantage distillation allows for generating key, at a higher value of the depolarising parameter $F_m = 0.97$, still no key can be generated with standard one-way post-processing. Moreover, we see that utilising advantage distillation for the SPADS scheme allows for the generation of key, even with very noisy measurements when $F_m = 0.91$. We also observe two distinct scalings of the secret-key fraction with F_m in the regime where non-zero amount of key is generated. These two scalings depend on whether we use a symmetric or asymmetric protocol. Specifically, for the SPADS scheme the symmetric six-state protocol is used. Therefore the corresponding two curves have a slope that is approximately three times smaller than the other three curves corresponding to the protocols that run in the fully asymmetric mode.

3.6.2. OPTIMAL SETTINGS

We see that the above described repeater schemes include several tunable parameters. These parameters are the cut-off $n^\star$ for Bob's number of attempts until restart, the angle θ in the single-photon scheme and the positioning of the repeater. These parameters can be optimised to maximise the secret-key rate. Here we will approach this optimisation in a consistent way - we gradually restrict the parameter space by making specific observations based on numerical evidence.

The first claim that we will make is in relation to the *optimal positioning of the repeater.* In the previous chapter we have conjectured that for the SiSQuaRe scheme the middle positioning of the repeater is optimal. For the single-photon scheme we want the probability of transmitting

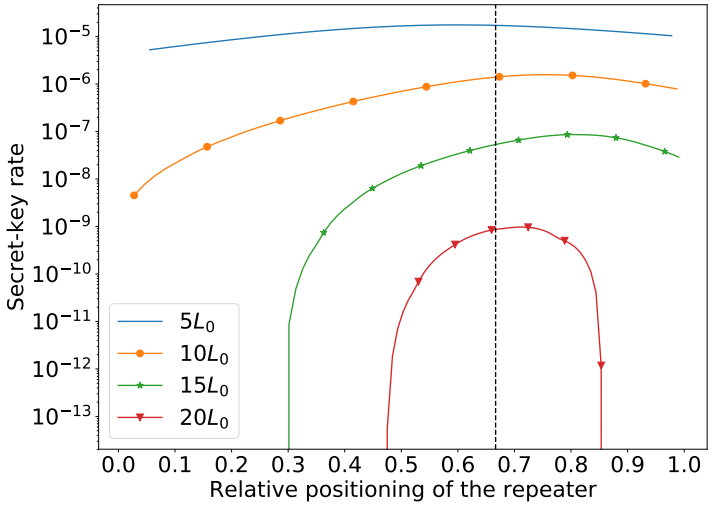

Figure 3.7: Secret-key rate as a function of the relative positioning of the repeater for few different total distances for the SPADS scheme. The total distances are expressed in terms of the fibre attenuation length $L_0 = 0.542$ km. We see that positioning the repeater two-thirds of the distance away from Alice (marked by the vertical black dashed line) is a good positioning for all the distances. For each total distance considered and each positioning the secret-key rate is optimised over the cutoff n^*, the angle θ and the time-window of the detector t_w.

the photons from each of the two nodes to the beam splitter heralding station to be equal. This effectively sets the target state between the electron spins to be the maximally entangled state. Hence, if we restrict ourselves to the case where the emission angles θ of both Alice and Bob are the same, then it is natural to position the heralding station symmetrically in the middle between them. Hence, the only non-obvious optimal positioning is for the SPADS and SPOTL scheme.

For the SPADS scheme, positioning the repeater at two-thirds of the relative distance away from Alice could intuitively be expected to be optimal. This is due to the fact that the single-photon scheme runs on two segments: Alice-beam splitter, beam splitter-repeater, while the one half of the SiSQuaRe scheme runs only over a single segment between repeater and Bob. By segment we mean here a distance over which we need to be able to independently transmit a photon. In Fig. 3.7 we show the secret-key rate as a function of the relative positioning of the repeater for a set of different total distances. We see there that despite the fact that positioning the repeater at two-thirds is not always optimal, it is a good enough positioning for all distances for our purposes. For each data point on the plot we independently optimise over the cut-off n^*, the angle θ of the single-photon subscheme and the duration of the detector time-window t_w.

The SPOTL scheme has the same symmetry as the SiSQuaRe scheme, in the sense that the part of the scheme performed on Alice's side is exactly the same as on Bob's side. This symmetry is only broken by the sequential nature of the scheme. Since we have already observed that the middle positioning is optimal for the SiSQuaRe scheme, we expect to see the same behavior for the SPOTL scheme. Indeed, we confirm this expectation numerically in Fig. 3.8. Here for each data point we independently optimise over the cut-off $n^\star$, the angle θ_A (θ_B) of the single-photon subscheme on Alice's (Bob's) side and the duration of the detection time-window.

To conclude, we will always place the heralding station within the single-photon (sub)protocol exactly in the middle between the two corresponding memory nodes. Moreover, we will also al-

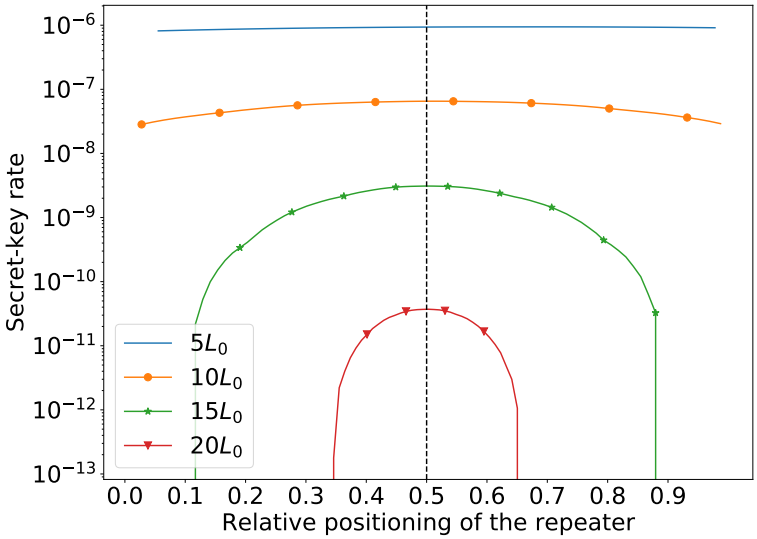

Figure 3.8: Secret-key rate as a function of the relative positioning of the repeater for few different total distances for the SPOTL scheme. The total distances are expressed in terms of the fibre attenuation length $L_0 = 0.542$ km. We see that positioning the repeater in the middle between Alice and Bob (marked by the vertical black dashed line) is a good positioning for all the distances. For each total distance considered and each positioning the secret-key rate is optimised over the cutoff n^*, the angles θ_A and θ_B and the time-window of the detector t_w.

ways place the memory repeater node in the middle for the SPOTL scheme and two-thirds of the distance away from Alice for the SPADS scheme.

Having established the optimal positioning of the repeater, we look into the relation between θ_A and θ_B for the SPOTL scheme. We observe that the relative error resulting from optimising the secret-key rate over a single angle $\theta_A = \theta_B$ rather than two independent ones is smaller than 1% for all distances. Hence from now on we will restrict ourselves to optimising only over one angle θ for the SPOTL scheme.

Having resolved the issues of the optimal positioning of the repeater for all schemes and reducing the number of angles to optimise over for the SPOTL scheme to one, we now investigate how our secret-key rate depends on the remaining parameters. These parameters are the angle θ, the cut-off $n^\star$ and the duration of the detection time-window t_w. The optimal time-window follows a simple behavior for all schemes: for short distances the probability of getting a dark count p_d is negligible compared to the probability of detecting the signal photon. Hence for those distances we can use a time-window of 30 ns to make sure that almost all the emitted photons which are not polluted by the photons from the optical excitation pulse arrive inside the detection time-window. We always need to sacrifice the photons arriving within the time t_w^{offset} after the optical pulse has been applied to filter out the photons from that pulse, see Appendix 3.8.1 for details. Then, for larger distances where p_d starts to become comparable with the probability of detecting the signal photon, the duration of the time-window is gradually reduced. This reduces the effect of dark counts at the expense of having more and more photons arriving outside of the time-window. See Appendix 3.8.1 for the modeling of the losses resulting from photons arriving outside of the time-window.

The dependence of the secret-key rate on the angle θ, the tunable parameter that Alice and

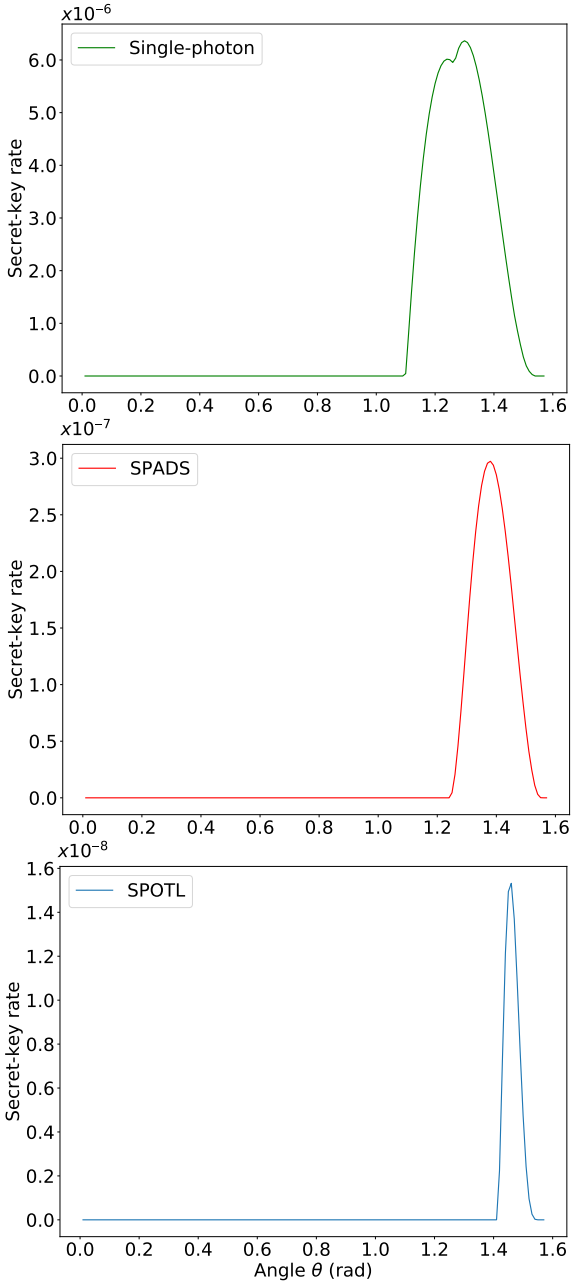

Figure 3.9: Secret-key rate as a function of the angle $\theta = \theta_A = \theta_B$ for the single-photon, SPADS and SPOTL schemes for the total distance of $12.5L_0$, where $L_0 = 0.542$ km. For each value of θ the secret-key rate is optimised over the cutoff n^* and time-window t_w. The schemes have a decreasing range of θ for which secret-key can be generated. This is due to the additional noisy processes in the SPADS scheme, and the overwhelming dominance of the dark state of the spin (no emission of the photon) in order to avoid any extra noise coming from the photon loss for SPOTL scheme. The visible kink for the single-photon scheme is a consequence of the fact that the six-state protocol with advantage distillation involves optimisation over of two subprotocols.

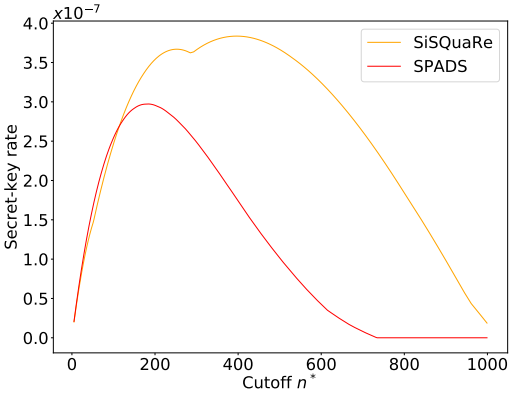

(a) Secret-key rates for the SiSQuaRe and SPADS scheme.

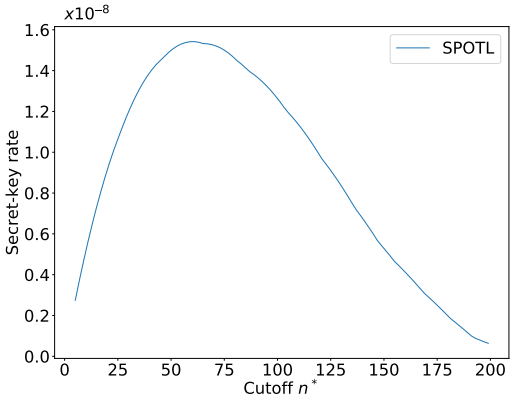

(b) Secret-key rates for the SPOTL scheme.

Figure 3.10: Secret-key rate as a function of the cut-off for the SiSQuaRe, SPADS and SPOTL scheme for the total distance of $12.5L_0$, where $L_0 = 0.542$ km. We observe a drop in the needed cut-off to maximise the secret-key rate, due to the schemes becoming progressively more noisy. For each value of the cutoff n^* we optimise the secret-key rate over the time-window t_w and for the SPADS scheme also over the θ angle. The kink for the SiSQuaRe scheme arises because of the optimisation over the fully asymmetric one-way BB84 protocol and symmetric six-state protocol with advantage distillation, which itself involves optimization over two subprotocols.

Bob choose in their starting state $|\psi\rangle = \sin\theta|\downarrow\rangle|0\rangle + \cos\theta|\uparrow\rangle|1\rangle$ in the single-photon scheme, is more complex. We observe that the optimal value of θ is closer to $\frac{\pi}{2}$ for schemes that involve more noisy processes. Informally, this means that Alice and Bob send 'less' photons towards the beam splitter, to overcome the noise coming from events in which both nodes emit a photon. At $\frac{\pi}{2}$ however, no photons are emitted and the rate drops down to zero. We illustrate this in Fig 3.9. We see that for the SPADS and SPOTL scheme, there is only a restricted regime of the angle θ for which one can generate non-zero amount of key. In particular, the SPOTL scheme requires a larger number of noisy operations, and therefore cannot tolerate much noise arising from the effect of photon loss in the single-photon subscheme. This means that there is only a small range of θ that allows for production of secret key. The single-photon scheme involves much less operations and can tolerate more noise, and so lower values of the parameter θ still allow for the generation of key.

We also investigate the dependence of the rate on the cut-off. Both the SPADS and SPOTL scheme require a lower cut-off than the SiSQuaRe scheme, see Fig. 3.10a and 3.10b. This is caused by the fact that each of them involves more noisy operations, and hence less noise tolerance is possible.

3.6.3. ACHIEVED SECRET-KEY RATES OF THE REPEATER PROPOSALS

Now we are ready to present the main results, the secret-key rate for all the considered schemes as a function of the total distance when optimised over θ, the cut-off $n^\star$ and the duration of the time-window t_w. We compare the rates to the benchmarks from Section 3.5.

In Fig. 3.11 we plot the rate of all four of the quantum repeater schemes as a function of the distance between Alice and Bob. We observe that already for realistic near-term parameters, the single-photon scheme can outperform the secret-key capacity of the pure-loss channel by a factor of seven for a distance of ≈ 9.2 km.

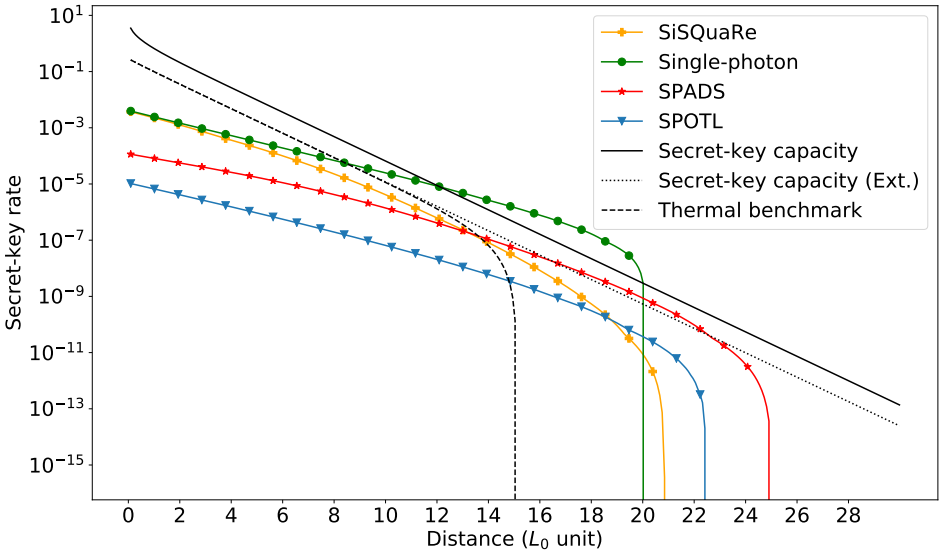

Figure 3.11: Rate of all studied quantum repeater schemes as a function of the distance between Alice and Bob, expressed in the units of $L_0 = 0.542$ km. We also plot the different benchmarks from Section 3.5. We see that the single-photon scheme outperforms the secret-key capacity. For the achievable rates the secret-key rate is optimised over the cutoff n^*, the angle θ and the time-window t_w independently for each distance.

We have also investigated what improvements would need to be done in order for the SPADS and SPOTL schemes to also overcome the secret-key capacity. An example scenario in which the SPADS scheme outperforms this repeaterless bound includes better phase stabilisation such that $\Delta\phi = 5°$ and reduction of the decoherence effects in the carbon spin during subsequent entanglement generation attempts such that $a_0 = 1/8000$ and $b_0 = 1/20000$. Further improvement of these effective coherence times to $a_0 = 1/20000$ and $b_0 = 1/50000$ allows the SPOTL scheme to also overcome the secret-key capacity. We note that maintaining coherence of the carbon-spin memory qubit for such large number of subsequent remote entanglement generation attempts is expected to be possible using the method of decoherence-protected subspaces [81, 135].

As mentioned before, the SPADS scheme can be naturally compared against the benchmark of the direct transmission using NV as a source. The results are depicted in Fig. 3.12. We see that the SPADS scheme easily overcomes the NV-based direct transmission and the thermal benchmark for larger distances for which these benchmarks drop to zero.

In Fig. 3.11 we observe that for the SPOTL scheme, the total distance over which key can be generated is significantly smaller than for the SPADS scheme. This is despite the fact that the full distance is divided into four segments. The rather weak performance of this scheme is due to the fact that it involves a larger number of noisy operations. As a result, the scheme can tolerate little noise from the single-photon subscheme, requiring the angle θ to be close to $\frac{\pi}{2}$ as can be seen in Fig. 3.9. As a result, the probability of photon emission becomes greatly diminished and so the distance after which dark counts start becoming significant is much smaller than for the SPADS scheme. To overcome this problem one would need to reduce the amount of noise in the system. One of the main sources of noise is the imperfect single-qubit measurement. Hence we illustrate the achievable rates for the scenario with the boosted measurement depolarising parameter $F_m = 0.98$ in Fig. 3.13. Additionally, in this plot we also consider the application of probabilis-

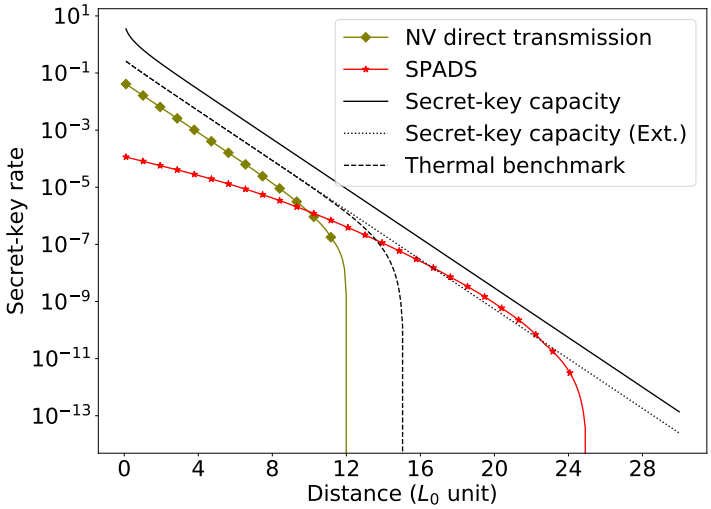

Figure 3.12: Comparison of the SPADS scheme with the rate achievable using the direct transmission, with NV being the photon source. The secret-key rates for those schemes are plotted as a function of the distance between Alice and Bob, expressed in the units of $L_0 = 0.542$ km. We also plot the different benchmarks. We see that the SPADS scheme easily overcomes the direct transmission and the thermal benchmark (see Section 3.5). For the secret-key rate achievable with the SPADS scheme we perform optimisation over the cutoff n^*, the angle θ and the time-window t_w independently for each distance. Similarly, we also optimise the secret-key rate achievable with direct transmission over the time-window t_w.

tic frequency conversion to the telecom wavelength at which $L_0 = 22$ km. Frequency conversion has already been achieved experimentally in the single-photon regime with success probability of 30% [180]. This is also the success probability that we consider here. The corresponding benchmarks have also been plotted for the new channel with $L_0 = 22$ km. We see in Fig. 3.13 that with the improved measurement and using frequency conversion, the SPOTL scheme allows now to generate secret key over more than 550 km. We also see that under those conditions the single-photon scheme can also overcome the secret-key capacity of the telecom channel.

3.6.4. RUNTIME OF THE EXPERIMENT

While the theoretical capability of an experimental setup to surpass the secret-key capacity is a necessary requirement to claim a working quantum repeater, it does not necessarily mean that this can be experimentally verified in practice. Indeed, if a quantum repeater proposal only surpasses the secret-key capacity by a narrow margin at a large distance, the running time of an experiment could be too long for practical purposes. In this section, we will discuss an experiment which can validate a quantum repeater setup and calculate the running time of such an experiment, where we demonstrate that the single-photon scheme could be validated to be a quantum repeater within twelve hours.

A straightforward way of validating a quantum repeater would consist of first generating secret-key, calculating the achieved (finite-size) secret-key rate and then comparing the rate with the secret-key capacity. However, this requires a large number of raw bits to be generated, partially due to the loose bounds on finite-size secret-key generation. What we propose here is an experiment where the QBER and yield are separately estimated to lie within a certain confidence interval. Then, if with the (worst-case) values of the yield and the QBER the corresponding asymptotic

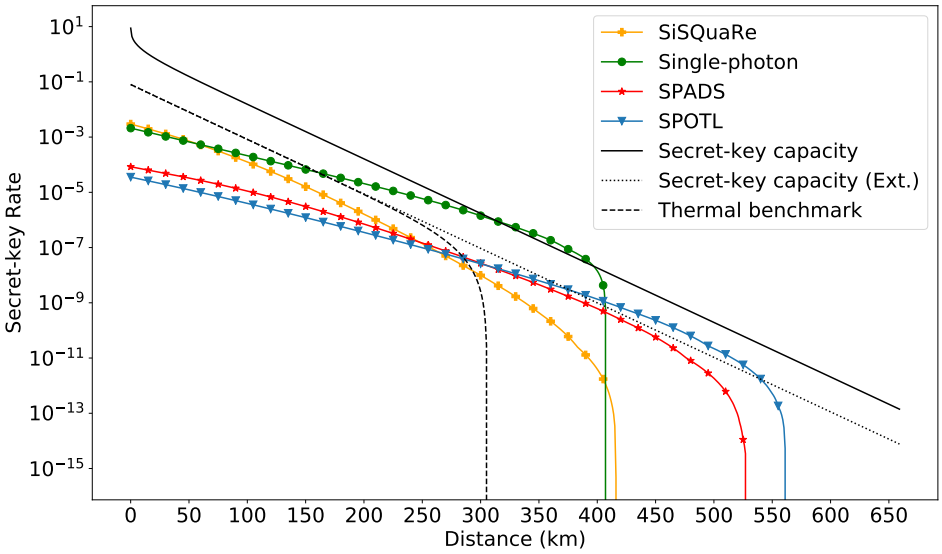

Figure 3.13: Secret-key rate as a function of distance in units of km for transmission at telecom channel with $L_0 = 22$ km, along with the benchmarks from Section 3.5. We consider an improved measurement depolarising parameter of $F_m = 0.98$. The frequency conversion efficiency is assumed to be 0.3. We observe that the SPOTL scheme allows for the generation of secret-key over a distance of more than 550 km. For the achievable rates the secret-key rate is optimised over the cutoff n^*, the angle θ and the time-window t_w independently for each distance.

secret-key rate still confidently beats the benchmarks, one could claim that, in the asymptotic regime, the setup would qualify as a quantum repeater.

As we show in Appendix 3.8.8, it is possible to run the single-photon scheme over a distance of $17L_0 \approx 9.2$ km for approximately twelve hours to find with high confidence ($\geq 1 - 1.5 \cdot 10^{-4}$) that the scheme beats the capacity (see Eq. (2.15)) at that distance by a factor of at least three.

3.6.5. DISCUSSION AND FUTURE OUTLOOK

It is worth noting that our figure of merit - the secret-key rate - is weakly impacted by the latency of transmission, which grows linearly with distance for the SiSQuaRe, SPADS and SPOTL schemes. Its only effect on the secret-key rate is the resulting decoherence time in the quantum memories while the memory nodes await the success/failure signals. This decoherence due to the waiting time is negligible in comparison to the noise due to interaction, arising from subsequent entanglement generation attempts. On the other hand, this latency would clearly be very visible in low throughput of these schemes. The single-photon scheme on the other hand has the advantage of the repetition rate being limited only by the local processing of the memory nodes which would result in a higher throughput. We observe this fact in the modest expected duration of the experiment, even in the high loss regime needed for overcoming the secret-key capacity. It is worth noting that while the single-photon scheme maintains constant latency for QKD, there exist schemes where such constant latency can be maintained also for remote entanglement generation, see e.g. [79]. It is hence clear that there are certain important properties of an efficient quantum repeater scheme that are not captured by the secret-key rate. However, achieving high throughputs for arbitrary distances would require almost all the components to be efficient in terms of rates and memories

to be of high quality in terms of operational and long-storage fidelities. It is clear that demonstrating all these features together in a single experiment is still a future goal. The advantage of the secret-key rate is that overcoming the secret-key capacity would form a crucial step towards an implementation of an efficient and practical, long-distance quantum repeater architecture whose validity would carry an information-theoretic significance and will therefore be totally independent of any hardware-based reference scenario.

In our model we have identified significant amount of noise arising in the system. As a result, we find that it is not always beneficial to just divide the fixed distance into more elementary links. Hence, it is a natural question whether this noise could be eliminated e.g. using entanglement distillation. In fact for the noise arising due to photon loss in the single-photon scheme not only does there exist an efficient distillation procedure [24, 118], but it has also already been demonstrated in the NV-platform [83]. Moreover, in the ideal case of noiseless operations and storage, a scheme based on generating two entangled states through the single-photon scheme and then distilling them as demonstrated in [83] should effectively also be able to overcome the secret-key capacity [165] and provide a significant boost by completely removing the noise due to photon loss. Furthermore, an implementation of such a distillation-based remote entanglement generation scheme would alleviate the requirement of the optical phase stabilisation of the system. Therefore this distillation based scheme could be a natural fifth candidate for a proof-of-principle repeater. Nevertheless, we believe that the fidelities of quantum operations and the effective coherence times of the memories used in this chapter might need to be improved before this distillation would prove useful.

Since the publication of this work, there has been an impressive improvement on performed experiments on two fronts. First, several twin-field QKD [98] experiments have been performed [25, 26, 52], beating the secret-key capacity for those distances. Twin-field QKD is closely related to the single-photon scheme, in the sense that both allow for surpassing the secret-key capacity (by achieving a $\sqrt{\eta}$ scaling of the secret-key rate), while not relying on any swapping operations. One drawback is that both the single-photon scheme and twin-field QKD do not scale, in the sense that achieving a better than square root scaling in the transmissivity is not possible using only these techniques. Still, the usefulness of twin-field QKD has been demonstrated by the implementation of QKD over a distance of over 500 kilometre in a non-laboratory environment [25]. Second, one of the main sources of errors in the memory-based repeater schemes in this and the previous chapter was the decoherence experienced during subsequent entanglement generation. Recent experiments have shown a significant reduction of the incurred dephasing by increasing the applied magnetic field [67, 132]. It would be of strong interest to see how the memory-based repeater proposals (which are thus scalable unlike twin-field QKD and the single-photon scheme) here would perform with such new improved parameters.

3.7. CONCLUSIONS

We analysed four experimentally relevant quantum repeater schemes on their ability to generate secret key. More specifically, the schemes were assessed by contrasting their achievable secret-key rate with the secret-key capacity of the channel corresponding to direct transmission. The secret-key rates have been estimated using near-term experimental parameters for the NV centre platform. The majority of these parameters have already been demonstrated across multiple experiments. A remaining challenging element of our proposed schemes is the implementation of optical cavities. These cavities would enable the enhancement of both the photon emission probability into the zero-phonon line and the photon collection efficiency to the desired level.

With these near-term experimental parameters, our assessment shows the viability of one of the schemes, the single-photon scheme, for the first experimental demonstration of a quantum re-

peater. In fact, the single-photon scheme achieves a secret-key rate more than seven times greater than the secret-key capacity. We also estimated the duration of an experiment to conclude that a rate larger than the secret-key capacity is achievable. The duration of the experiment would be approximately twelve hours.

Finally, we show that a scheme based on concatenating the single-photon scheme twice (i.e. the SPOTL scheme), has the capability to generate secret-key at large distances. However, this requires converting the frequency of the emitted photons to the telecom wavelength and modestly improving the fidelity at which measurements can be performed.

3.8. APPENDIX

3.8.1. LOSSES AND NOISE ON THE PHOTONIC QUBITS

In this Appendix we describe how the losses and noise affect our photonic qubits. In particular, we first recall how the two types of encoding result in the losses acting as different quantum channels on the states. Then, we study the effects of a finite detector time-window. More specifically, we firstly show that the arrival of a photon outside the time-window is equivalent to all the other loss processes and secondly we calculate the probability of registering a dark count within the time-window. We also show how to model the noise arising from those dark counts for the SiSQuaRe and SPADS schemes. Finally, we calculate the dephasing induced by the unknown phase shift for the single-photon scheme.

EFFECTS OF LOSSES FOR THE DIFFERENT ENCODINGS

The physical process of probabilistically losing photons corresponds to different quantum channels depending on the qubit encoding. In our repeater schemes we use two types of encoding: time-bin and presence-absence of a photon. For a time-bin encoded qubit in the ideal scenario of no loss we always expect to obtain a click in one of the detectors. Hence loss of a photon resulting in a no-click event raises an erasure flag which carries the failure information. Therefore it is clear that for this encoding the physical photon loss process corresponds to an erasure channel with the erasure probability given by one minus the corresponding transmissivity,

$$D(\rho) = \eta\rho + (1-\eta)|\perp\rangle\langle\perp| \,. \tag{3.6}$$

Here $|\perp\rangle$ is the loss flag, corresponding to the non-detection of a photon. Since we are only interested in the quantum state of the system for the successful events when a detection event has occurred, we effectively post-select on the non-erasure events. For presence-absence encoding the situation is different since now there is no flag available that could explicitly tell us whether a photon got lost or not. In fact for this encoding the photon loss results in an amplitude-damping channel applied to the photonic qubit. Here the damping parameter equals one minus the transmissivity of the channel [34].

EFFECTS OF THE DETECTOR TIME-WINDOW

The detector only registers clicks that fall within a certain time-window. It is *a priori* not clear what kind of noisy or lossy channel should be used to model the loss of information due to non-detection of photons arriving outside of the time-window. This is because in a typical loss process we have a probabilistic leakage of information to the environment. In the scenario considered here, the situation is slightly different as effectively no leakage occurs, but rather certain part of the incoming signal effectively gets discarded. Here we will show that despite this qualitative difference, within our model this process can effectively be modelled as any other loss process.

Now, let us provide a brief description of the physics of this process. Firstly, the detection time-window is chosen such that the probability of detecting a photon from the optical excitation pulse

used to entangle the electron spin with the photonic qubit is negligible [65]. For that reason the detection time-window is opened after a fixed offset t_w^{offset} with respect to the beginning of the decay of the optical excited state of the electron spin. We note that for the considered enhancement of the ZPL-emission using the optical cavity we predict the characteristic time of the NV emission τ to be approximately a half of the corresponding value of τ if no cavity is used [53, 65, 138]. Therefore here we consider the scenario where the duration of the optical excitation pulse is made twice shorter with respect to the one used in [65]. This will allow us to filter out the unwanted photons from the excitation pulse by setting t_w^{offset} to half of the offset used in [65].

Secondly, we note that the detection time-window cannot last too long, specifically, it needs to be chosen such that there is a good trade-off between detecting coherent and non-coherent (i.e. dark counts) photons. In this subsection we will discuss the effects of photons arriving outside of this time-window and the effects of registering dark counts within this time-window.

Losses from the detector time-window. The NV centre emits a photon through an exponential decay process with characteristic time τ. Therefore the probability of detecting a photon during a time-window starting at t_w^{offset} and lasting for t_w is

$$p_{\text{in}}(t_w) = \frac{1}{\tau}\int_{t_w^{\text{offset}}}^{t_w^{\text{offset}}+t_w} \mathrm{d}t \exp\left(-\frac{t}{\tau}\right) = \exp\left(-\frac{t_w^{\text{offset}}}{\tau}\right) - \exp\left(-\frac{t_w^{\text{offset}}+t_w}{\tau}\right). \tag{3.7}$$

Clearly the process of a photon arriving outside of the time-window is qualitatively different from the loss process where the photons get lost to the environment. In the remainder of this section we will now look at the difference between these two phenomena in more detail.

The emission process of the NV centre is a coherent process over time. Consider a generic scenario in which we divide the emission time into two intervals, denoted by "in" and "out", respectively. Coherent emission then means that the state of the photon emitted by the electron spin in state $|\uparrow\rangle$ will be

$$|\psi\rangle = \sqrt{p_{\text{in}}}\,|1\rangle_{\text{in}}|0\rangle_{\text{out}} + \sqrt{1-p_{\text{in}}}\,|0\rangle_{\text{in}}|1\rangle_{\text{out}}\,. \tag{3.8}$$

Now let us come back to our specific model, in which the "in" mode corresponds to the interval $\left[t_w^{\text{offset}},\ t_w^{\text{offset}}+t_w\right]$ and the "out" mode to all the times $t \geq 0$ lying outside of this interval ($t=0$ is the earliest possible emission time). Here, the emission into the "in" mode occurs with probability $p_{\text{in}}(t_w)$. Hence the spin-photon state resulting from the emission by the $\alpha|\downarrow\rangle + \beta|\uparrow\rangle$ spin state is

$$|\psi\rangle = \alpha|\downarrow\rangle|0\rangle_{\text{in}}|0\rangle_{\text{out}} + \beta|\uparrow\rangle\left(\sqrt{p_{\text{in}}(t_w)}\,|1\rangle_{\text{in}}|0\rangle_{\text{out}} + \sqrt{1-p_{\text{in}}(t_w)}\,|0\rangle_{\text{in}}|1\rangle_{\text{out}}\right). \tag{3.9}$$

If the presence-absence encoding is used, such a photonic qubit is then transmitted to the detector. Since only the spin and the "in" mode of the photon will be measured, we can now trace out the "out" mode

$$\rho = \left(|\alpha|^2 + |\beta|^2 p_{\text{in}}(t_w)\right)|\phi\rangle\langle\phi| + |\beta|^2(1-p_{\text{in}}(t_w))|\uparrow\rangle\langle\uparrow|\otimes|0\rangle\langle 0|_{\text{in}}\,, \tag{3.10}$$

where

$$|\phi\rangle = \frac{1}{\sqrt{|\alpha|^2 + |\beta|^2 p_{\text{in}}(t_w)}}\left(\alpha|\downarrow\rangle|0\rangle_{\text{in}} + \beta\sqrt{p_{\text{in}}(t_w)}\,|\uparrow\rangle|1\rangle_{\text{in}}\right). \tag{3.11}$$

Note that this state can be obtained by passing the photonic qubit of the state

$$|\psi\rangle = \alpha|\downarrow\rangle|0\rangle + \beta|\uparrow\rangle|1\rangle\,, \tag{3.12}$$

through the amplitude-damping channel with the damping parameter given by $1-p_{\text{in}}(t_w)$. Hence we can conclude that for the photon number encoding, the possibility of the photon arriving outside of the time-window of the detector can be modelled in the same way as any other photon loss process, namely an amplitude-damping channel applied to that photonic qubit.

In the case of time-bin encoding we effectively have four photonic qubits, since now we have an "in" and "out" mode for both the early (denoted by "e") and the late (denoted by "l") time-window. We assume here that the slots do not overlap. That is, a photon emitted in the "out" mode of the early time-window is always distinct from any photon in the late time-window. This can be achieved by making the time gap between the "in" modes of the early and late window long enough. In this case the emission process results in a state

$$|\psi\rangle = \alpha\,|\downarrow\rangle\left(\sqrt{p_{\text{in}}(t_{\text{w}})}\,|1\rangle_{e,\text{in}}\,|0\rangle_{e,\text{out}}\,|0\rangle_{l,\text{in}}\,|0\rangle_{l,\text{out}} + \sqrt{1-p_{\text{in}}(t_{\text{w}})}\,|0\rangle_{e,\text{in}}\,|1\rangle_{e,\text{out}}\,|0\rangle_{l,\text{in}}\,|0\rangle_{l,\text{out}}\right)$$
$$+\,\beta\,|\uparrow\rangle\left(\sqrt{p_{\text{in}}(t_{\text{w}})}\,|0\rangle_{e,\text{in}}\,|0\rangle_{e,\text{out}}\,|1\rangle_{l,\text{in}}\,|0\rangle_{l,\text{out}} + \sqrt{1-p_{\text{in}}(t_{\text{w}})}\,|0\rangle_{e,\text{in}}\,|0\rangle_{e,\text{out}}\,|0\rangle_{l,\text{in}}\,|1\rangle_{l,\text{out}}\right).$$

Again, tracing out the "out" modes results in a state

$$\rho = p_{\text{in}}(t_{\text{w}})|\phi\rangle\langle\phi| + (1-p_{\text{in}}(t_{\text{w}}))\left(|\alpha|^2|\downarrow\rangle\langle\downarrow| + |\beta|^2|\uparrow\rangle\langle\uparrow|\right)\otimes|00\rangle\langle 00|_{e,l}\,, \tag{3.13}$$

where

$$|\phi\rangle = \alpha\,|\downarrow\rangle\,|1\rangle_e\,|0\rangle_l + \beta\,|\uparrow\rangle\,|0\rangle_e\,|1\rangle_l = \alpha\,|\downarrow\rangle\,|e\rangle + \beta\,|\uparrow\rangle\,|l\rangle\,. \tag{3.14}$$

Here $|00\rangle_{e,l}$ corresponds to the loss flag from which we see that for the time-bin encoding the possible arrival of a photon outside of the time-window results in an erasure channel with the erasure probability given by $(1-p_{\text{in}}(t_{\text{w}}))$. Hence this process can be also modelled as any other loss process for this encoding.

We have just shown that for both photon presence/absence and time-bin encodings the process of the photon arriving outside of the time-window can be modelled by the source which prepares photons in a coherent superposition of the "in" and "out" modes and the detector tracing out (losing) the "out" modes. We have also shown that those two elements combined together result effectively in a loss process corresponding to the same channel as any other loss process for that encoding (amplitude-damping for photon presence/absence and erasure channel for time-bin encoding).

However, between the source and the detector there are other lossy or noisy components resulting in other quantum channels that need to be applied before the tracing out of the "out" mode at the detector. Now we show that for all loss and noise processes that occur in our model, the tracing out of the "out" mode can be mathematically commuted through all those additional noise/lossy processes. This means that the tracing out can be applied directly after the source, such that the above described reductions to amplitude-damping or erasure channel can be applied.

Consider the quantum channels acting on the photonic qubits of the form

$$\mathcal{N} = \sum_i p_i \mathcal{N}^i_{\text{in}} \otimes \mathcal{N}^i_{\text{out}}\,. \tag{3.15}$$

Effectively these are the channels that do not couple the "in" and "out" modes. Since in reality "in" and "out" modes correspond to different time modes, their coupling would require some kind of memory inside the channel. Hence we can think of the above defined channels as channels without memory. Now it is clear that for a quantum state ρ that among its registers includes both the "in" and the "out" mode, we have that

$$\text{tr}_{\text{out}}[\mathcal{N}(\rho)] = \text{tr}_{\text{out}}\left[\sum_i p_i \mathcal{N}^i_{\text{in}} \otimes \mathcal{N}^i_{\text{out}}(\rho)\right] = \sum_i p_i \mathcal{N}^i_{\text{in}}(\rho_{\text{in}})\,. \tag{3.16}$$

Now, firstly tracing out the "out" modes and then applying the channel $\mathcal{N}$ (only the "in" part can be applied now) also results in $\sum_i p_i \mathcal{N}^i_{\text{in}}(\rho_{\text{in}})$ at the output. Hence the tracing out of the "out"

modes commutes with all the channels that are of the form (3.15), which correspond to channels without memory. Clearly the noise/loss processes that occur before the detection, such as photon loss or dephasing due to uncertainty in the optical phase of the photon, belong to this class of channels. In particular this means that for photon presence/absence the amplitude-damping due to photon loss in the channel and due to photon arrival outside of the time-window can be both combined into one channel with the single damping parameter given by $1-\eta p_{\text{in}}(t_{\text{w}})$ (η denotes the transmissivity due to the loss process e.g. the transmissivity of the fibre). The same applies to time-bin encoding where we now have a single erasure channel with erasure probability $1-\eta p_{\text{in}}(t_{\text{w}})$.

To conclude, the arrival of the photon outside of the time-window can be modelled in the same way as any other loss process for both photon encodings used and therefore we can now redefine the detector efficiency $p'_{\text{det}} = p_{\text{det}} \cdot p_{\text{in}}(t_{\text{w}})$ and the total apparatus efficiency $p'_{\text{app}} = p_{\text{ce}} p_{\text{zpl}} p'_{\text{det}}$. We can then define $\eta_{\text{total}} = p'_{\text{app}} \eta_f$ as the total transmissivity - with probability η_{total} a photon will be successfully transmitted from the sender to the receiver.

Dark counts within the detector time-window.

Photon detectors are imperfect, and due to thermal excitations, they will register clicks that do not correspond to any incoming photons. These undesired clicks are called dark counts and can effectively be seen as a source of noise. The magnitude of this noise depends on the ratio between the probability of detecting the signal photon and measuring a dark count. Clearly, dark counts become a dominant source of noise when the probability of detecting the signal photon becomes comparable to the probability of a dark count click. The probability p_d of getting at least one dark count within the time-window t_{w} of awaiting the signal photon is given by $p_d = 1-\exp(-t_{\text{w}} \cdot \text{DCpS})$, where DCpS is the number of dark count per second of the detector, see the previous chapter.

In the SiSQuaRe scheme Alice and Bob perform measurements on time-bin encoded photons. The same applies to Bob in the SPADS scheme. Since at least two detectors are required to perform this measurement, the presence of dark counts means that the outcome may lie outside of the qubit space. Moreover, this measurement needs to be trusted. In consequence, a squashing map needs to be used to process the multi-click events in a secure way. Here as an approximation we consider the squashing map for the polarisation encoding [57] in the same way as described in the previous chapter. Hence this measurement can also be modelled as a perfect measurement preceded by a depolarising channel with parameter α which depends on whether the BB84 or six-state protocol is used. The parameter α is given by (see previous chapter):

$$\alpha_{A/B,\,\text{BB84}} = \frac{p'_{\text{app}} \eta_B (1-p_d)}{1-(1-p'_{\text{app}}\eta_{A/B})(1-p_d)^2}\,, \tag{3.17}$$

$$\alpha_{A/B,\,\text{six-state}} = \frac{p'_{\text{app}} \eta_{A/B} (1-p_d)^5}{1-(1-p'_{\text{app}}\eta_{A/B})(1-p_d)^6}\,. \tag{3.18}$$

Here $\eta_{A/B}$ denotes the transmissivity of the fibre between the memory repeater node and Alice's/Bob's detector setup. Finally we note that dark counts increase the probability of registering a successfull measurement event. For the optical measurement schemes utilising the squashing map the probability of registering a click in at least one detector is given by (see preceding chapter):

$$p_{A/B,\,BB84} = 1-(1-p'_{\text{app}}\eta_{A/B})(1-p_d)^2\,, \tag{3.19}$$

$$p_{A/B,\,\text{six-state}} = 1-(1-p'_{\text{app}}\eta_{A/B})(1-p_d)^6\,. \tag{3.20}$$

The effect of dark counts in the single-photon scheme, which carries over to the SPOTL scheme, is analyzed in Appendix 3.8.5.

NOISE DUE TO OPTICAL PHASE UNCERTAINTY

Another important noise process affecting photonic qubits is related to the fact that for the photon presence/absence encoding the spin-photon entangled state will also depend on the optical phase of the apparatus used. Specifically, it will depend on the phase of the lasers used to generate the spin photon entanglement as well as the optical phase acquired by the photons during the transmission of the photonic qubit. Knowledge about this phase is crucial for being able to generate entanglement through the single-photon scheme. In any realistic setup however, there would be a certain degree of the lack of knowledge about this phase acquired by the photons. Since in the end what matters is the knowledge about the relative phase between the two photons, we can model this source of noise as the lack of knowledge of the phase on only one of the incoming photonic qubits. This noise process can be effectively modelled as dephasing. In this section we will show that the phase uncertainty induces dephasing with a parameter λ equal to

$$\lambda = \frac{I_1\left(\frac{1}{(\Delta\phi)^2}\right)}{2I_0\left(\frac{1}{(\Delta\phi)^2}\right)} + \frac{1}{2}\,, \tag{3.21}$$

where $\Delta\phi$ is the uncertainty in the phase and $I_{0/1}$ is the Bessel function of order 0/1. Let us assume that for Alice, the local phase of the photonic qubit has a Gaussian-like distribution on a circle, i.e. it corresponds to a *wrapped distribution* with standard deviation $\Delta\phi$ as observed in [72]. This motivates us to model the distribution as a von Mises distribution [75]. The von Mises distribution reads

$$f(\phi) = \frac{e^{\kappa\cos(\phi-\mu)}}{2\pi I_0(\kappa)}\,. \tag{3.22}$$

Here μ is the measure of location, i.e. it corresponds to the centre of the distribution, κ is a measure of concentration and can be effectively seen as the inverse of the variance and I_0 is the modified Bessel function of the first kind of order 0. One can then show [75] that

$$\int_{-\pi}^{\pi} \mathrm{d}\phi f(\phi) e^{\pm i\phi} = \frac{I_1(\kappa)}{I_0(\kappa)} e^{\pm i\mu}\,. \tag{3.23}$$

Since we are only interested in the noise arising from the lack of knowledge about the phase rather than the actual value of this phase, without loss of generality we can assume $\mu = 0$. Moreover, the experimental parameter that we use here is effectively the standard deviation of the distribution $\Delta\phi$ and therefore we can write $\kappa = \frac{1}{(\Delta\phi)^2}$.

Hence, let us write the spin-photon entangled state that depends on the optical phase ϕ.

$$|\psi^{\pm}(\phi)\rangle = \sin(\theta)\,|{\downarrow}\,0\rangle \pm e^{i\phi}\cos(\theta)\,|{\uparrow}\,1\rangle\,. \tag{3.24}$$

Now, the lack of knowledge about this phase leads to a mixed state:

$$\begin{aligned}\int_{-\pi}^{\pi} f(\phi)|\psi^{\pm}(\phi)\rangle\langle\psi^{\pm}(\phi)|\mathrm{d}\phi = \sin^2(\theta)|{\downarrow}\,0\rangle\langle{\downarrow}\,0| + \cos^2(\theta)|{\uparrow}\,1\rangle\langle{\uparrow}\,1| \\ \pm \sin(\theta)\cos(\theta)\int_{-\pi}^{\pi} f(\phi)(e^{i\phi}|{\uparrow}\,1\rangle\langle{\downarrow}\,0| + e^{-i\phi}|{\downarrow}\,0\rangle\langle{\uparrow}\,1|)\mathrm{d}\phi\,.\end{aligned} \tag{3.25}$$

Let us now try to map this state onto a dephased state

$$\begin{aligned}\lambda|\psi^{\pm}(0)\rangle\langle\psi^{\pm}(0)| + (1-\lambda)|\psi^{\mp}(0)\rangle\langle\psi^{\mp}(0)| = \sin^2(\theta)|{\downarrow}\,0\rangle\langle{\downarrow}\,0| + \cos^2(\theta)|{\uparrow}\,1\rangle\langle{\uparrow}\,1| \\ \pm \sin(\theta)\cos(\theta)(2\lambda-1)(|{\uparrow}\,1\rangle\langle{\downarrow}\,0| + |{\downarrow}\,0\rangle\langle{\uparrow}\,1|)\,.\end{aligned} \tag{3.26}$$

Hence, we observe that

$$2\lambda - 1 = \frac{I_1\left(\frac{1}{(\Delta\phi)^2}\right)}{I_0\left(\frac{1}{(\Delta\phi)^2}\right)}. \tag{3.27}$$

$$\rightarrow \lambda = \frac{I_1\left(\frac{1}{(\Delta\phi)^2}\right)}{2I_0\left(\frac{1}{(\Delta\phi)^2}\right)} + \frac{1}{2}. \tag{3.28}$$

3.8.2. NOISY PROCESSES IN NV-BASED QUANTUM MEMORIES

In our setups we use ^{13}C nuclear spins in diamond as long-lived memory qubits next to a Nitrogen-Vacancy (NV) electron spin taking the role of a communication qubit. In this Appendix, we will focus on the modifications to the model presented in the previous chapter, since for most error processes the model is the same.

Storage of quantum states in the carbon spin memory during entanglement generation attempts are modelled as applying dephasing noise with parameter

$$\lambda_1 = F_{T_2} = \frac{1 + e^{-a\cdot n}}{2}, \tag{3.29}$$

$$\text{where } a = a_0 + a_1\left(L_s \cdot \frac{n_{ri}}{c} + t_{\text{prep}}\right), \tag{3.30}$$

$$\tag{3.31}$$

and depolarising with parameter,

$$\lambda_2 = F_{T_1} = e^{-b\cdot n}, \tag{3.32}$$

$$\text{where } b = b_0 + b_1\left(L_s \cdot \frac{n_{ri}}{c} + t_{\text{prep}}\right). \tag{3.33}$$

Here n_{ri} is the refractive index of the fibre, c is the speed of light in vacuum, t_{prep} is the time it takes to prepare for the emission of an entangled photon and L_s is the distance the signal needs to travel before the repeater receives the information about failure or success of the attempt. Let L_B denote the distance between the memory repeater node and Bob. Then for the SiSQuaRe and SPADS schemes $L_s = 2L_B$ as in each attempt first the quantum signal needs to travel to Bob who then sends back to the middle node the classical information about success or failure. For the SPOTL scheme $L_s = L_B$ as in this case both the quantum and the classical signals need to travel only half of the distance between the middle node and Bob since the signals are exchanged with the heralding station which is located half-way between the middle memory node and Bob. The parameters a_0 and b_0 quantify the noise due to a single attempt at generating an entangled spin-photon, induced by stochastic electron spin reset operations, quasi static noise and microwave control infidelities. The parameters a_1 and b_1 quantify the noise during storage per second.

Gates and measurements in the quantum memory are also imperfect. We model those imperfections via two depolarising channels. The first one acts on a single qubit with depolarising parameter $\lambda_2 = F_m$ corresponding to the measurement of the electron spin. The second one acts on two qubits with depolarising parameter $\lambda_2 = F_g$ corresponding to applying a two-qubit gate to both the electron spin and the ^{13}C spin. This means that every time a measurement is done on a e^- qubit of a quantum state ρ, it is actually done on $\mathscr{D}^{F_m}_{\text{depol}}(\rho)$. Also a swapping operation between the e^- spin and the nuclear spin (done experimentally via two two-qubit gates, see main text) leads to an error modelled by a depolarising channel of parameter $F_{\text{swap}} = F_g^2$. Following the same logic, a Bell state measurement will cause the state to undergo an evolution given by a

depolarising channel. Specifically, following the decomposition of the Bell measurement into elementary gates for the NV-implementation as described in Section 3.3, this evolution will consist of a depolarising channel with parameter F_g^2 acting on both of the measured qubits and the depolarising channel with parameter F_m^2 acting only on the electron spin qubit, see the previous chapter for more information.

3.8.3. EXPECTATION OF THE NUMBER OF CHANNEL USES WITH A CUT-OFF

In this Appendix we derive an analytical formula for the expectation value of the number of channel uses between Alice and Bob needed to generate one bit of raw key for the SiSQuaRe, SPADS and SPOTL schemes,

$$\mathbb{E}[N] = \frac{1}{p_A \cdot \left(1-(1-p_B)^{n^*}\right)} + \frac{1}{p_B} \,. \tag{3.34}$$

For these three schemes, we implement a cut-off which is used to prevent decoherence. Each time the number of channel uses between the repeater node and Bob reaches the cut-off n^*, the entire protocol restarts from the beginning. Here we take a conservative view and define the number of channel uses N between Alice and Bob as the sum $N_A + N_B$, where N_A (N_B) corresponds to the number of channel uses between Alice (Bob) and the middle node. From the linearity of the expectation value we have that

$$\mathbb{E}[N_A + N_B] = \mathbb{E}[N_A] + \mathbb{E}[N_B] \,. \tag{3.35}$$

We denote by p_A and p_B the probability of a successful attempt on Alice's and Bob's side respectively. Bob's number of channel uses follows a geometric distribution with parameter $p = p_B$, so that $\mathbb{E}[N_B] = \frac{1}{p_B}$. Without the cut-off, Alice's number of channel use would follow a geometric distribution with parameter $p = p_A$. However, the cut-off parameter adds additional channel uses on Alice side. Since the probability that Bob succeeds within n^* trials is $p_{\text{succ}} = 1-(1-p_B)^{n^*}$, we in fact have that Alice's number of channel uses follows a geometric distribution with parameter $p'_A = p_A \cdot p_{\text{succ}}$. Hence it is straightforward to see that

$$\mathbb{E}[N_A + N_B] = \frac{1}{p'_A} + \frac{1}{p_B} \tag{3.36}$$

$$= \frac{1}{p_A \cdot \left(1-(1-p_B)^{n^*}\right)} + \frac{1}{p_B} \,. \tag{3.37}$$

3.8.4. SISQUARE SCHEME ANALYSIS

The analysis of the SiSQuare scheme has been performed in the previous chapter. In this chapter we use the estimates of the yield and QBER as derived before with the following modifications:

- For the calculation of the yield we now adopt a conservative perspective and calculate the number of channel uses as $\mathbb{E}[N_A+N_B]$, as derived in Appendix 3.8.3, rather than $\mathbb{E}[\max(N_A, N_B)]$. Note that $\mathbb{E}[\max(N_A, N_B)] \leq \mathbb{E}[N_A + N_B] \leq 2\mathbb{E}[\max(N_A, N_B)]$.
- The total depolarising parameter for gates and measurements F_{gm} defined in the previous chapter is now decomposed into individual operations as described in Appendix 3.8.2. That is, in this chapter depolarisation due to imperfect operations on the memories is expressed in terms of depolarising parameter due to imperfect measurement, F_m, and imperfect two-qubit gate, F_g. Since in the analysis of the SiSQuaRe scheme we only deal with Bell diagonal states, the overall noise due to imperfect swap gate and the Bell measurement leads to $F_{\text{gm}} = F_g^4 F_m^2$.

- In the previous chapter we assumed the duration of the detection time-window to be fixed to 30 ns and assumed that all the emitted photons will fall into that time-window. Here, similarly as for other schemes, we perform a more refined analysis in which we include the trade-off between the duration of the time-window and the dark count probability as described in Appendix 3.8.1.
- We do not limit ourselves only to extracting key from the Z-basis for the advantage distillation scheme, as was done in the preceding chapter. By allowing extracting secret-key from the X- or Y-basis as well, a higher secret-key fraction can be achieved, see 3.8.7 for more information.

3.8.5. SINGLE-PHOTON SCHEME ANALYSIS

In this Appendix we provide a detailed analysis of the single-photon scheme between two remote NV-centre nodes. This section is structured as follows. First, we describe the creation of the spin-photon entangled state followed by the action of the lossy channel on the photonic part of this state, including the noise due to the uncertainty in the phase of the state induced by the fibre. Second, we apply the optical Bell measurement. Then we evaluate the effect of dark counts which introduce additional errors to the generated state. Finally we calculate the yield of this scheme and extract the QBER from the resulting state.

SPIN-PHOTON ENTANGLEMENT AND ACTION OF A LOSSY FIBRE ON THE PHOTONIC QUBIT

Firstly, both Alice and Bob generate spin-photon entangled states, parameterised by θ. As we will later see, this parameter allows for trading off the quality of the final entangled state of the two spins with the yield of the generation process. The ideal spin-photon state would then be described as

$$|\psi^+\rangle = \sin(\theta)|\downarrow\rangle|0\rangle + \cos(\theta)|\uparrow\rangle|1\rangle \ . \tag{3.38}$$

The preparation of the spin-photon entangled state is not ideal. That is, the spin-photon entangled state is not actually as described above, but rather of the form (see Appendix 3.8.2)

$$\rho = F_{\text{prep}}|\psi^+\rangle\langle\psi^+| + (1-F_{\text{prep}})(\mathbb{I}\otimes Z)|\psi^+\rangle\langle\psi^+|(\mathbb{I}\otimes Z) \tag{3.39}$$

$$= F_{\text{prep}}|\psi^+\rangle\langle\psi^+| + (1-F_{\text{prep}})|\psi^-\rangle\langle\psi^-| \ . \tag{3.40}$$

Here

$$|\psi^-\rangle = \sin(\theta)|\downarrow\rangle|0\rangle - \cos(\theta)|\uparrow\rangle|1\rangle \ . \tag{3.41}$$

For the next step we need to consider two additional noise processes that affect the photonic qubits before the optical Bell measurement is performed. The first one is the loss of the photonic qubit. This can happen at the emission, while filtering the photons that are not of the required ZPL frequency, in the lossy fibre, in the imperfect detectors, or due to the arrival outside of the time window in which detectors expect a click. All these losses can be combined into a single loss parameter

$$\eta = \eta_{\text{total}} = p_{\text{ce}}\, p_{\text{zpl}} \sqrt{\eta_f}\, p'_{\text{det}} \ , \tag{3.42}$$

with $\eta_f = \exp\left(-\frac{L}{L_0}\right)$, where L is the distance between the two remote NV-centre nodes in the scheme (see Fig. 3.5 and Appendix A). Hence, a photon is successfully transmitted through the fibre and detected in the middle heralding station with probability η. Now we note that the action of the pure-loss channel on the qubit encoded in the presence or absence of a photon corresponds to the action of the amplitude-damping channel with the damping parameter $1-\eta$ [34].

The second process that effectively happens at the same time as loss, is the dephasing noise arising from the optical instability of the apparatus as described in Appendix 3.8.1. We note that the amplitude-damping and dephasing channel commute, hence it does not matter in which order

we apply the two noise processes corresponding to the loss of the photonic qubit and unknown drifts of the phase of the photonic qubit in our model. Here we firstly apply the dephasing due to the lack of knowledge of the phase on Alice's photon and then amplitude-damping on both photons due to all the loss processes.

Following the model in Appendix 3.8.1, the lack of knowledge about the optical phase will effectively transform Alice's state to

$$\begin{aligned}\rho_A = &\left(F_{\text{prep}}\lambda + (1-F_{\text{prep}})(1-\lambda)\right)|\psi^+\rangle\langle\psi^+| \\ &+\left((1-F_{\text{prep}})\lambda + F_{\text{prep}}(1-\lambda)\right)|\psi^-\rangle\langle\psi^-|\,.\end{aligned}$$

where

$$\lambda = \frac{I_1\left(\frac{1}{(\Delta\phi)^2}\right)}{2I_0\left(\frac{1}{(\Delta\phi)^2}\right)} + \frac{1}{2}. \tag{3.43}$$

Now we can apply all the transmission losses modelled as the amplitude-damping channel. The action of this channel on the photonic part of the state ρ results in the state that we can describe as follows. Firstly, let us introduce two new states

$$|\psi_\eta^\pm\rangle = \frac{1}{\sqrt{\sin^2(\theta)+\eta\cos^2(\theta)}}\left(\sin(\theta)\,|\downarrow\rangle\,|0\rangle \pm \sqrt{\eta}\cos(\theta)\,|\uparrow\rangle\,|1\rangle\right). \tag{3.44}$$

Then, after the losses and before the Bell measurement, the state of Alice can be written as

$$\begin{aligned}\rho'_A = &\left(\sin^2(\theta)+\eta\cos^2(\theta)\right)\Big(\left(F_{\text{prep}}\lambda + (1-F_{\text{prep}})(1-\lambda)\right)|\psi_\eta^+\rangle\langle\psi_\eta^+| \\ &+\left((1-F_{\text{prep}})\lambda + F_{\text{prep}}(1-\lambda)\right)|\psi_\eta^-\rangle\langle\psi_\eta^-|\Big) + (1-\eta)\cos^2(\theta)|\uparrow\rangle\langle\uparrow|\,|0\rangle\langle 0|\,,\end{aligned}$$

and for Bob

$$\begin{aligned}\rho'_B = &\left(\sin^2(\theta)+\eta\cos^2(\theta)\right)\Big(F_{\text{prep}}|\psi_\eta^+\rangle\langle\psi_\eta^+| + (1-F_{\text{prep}})|\psi_\eta^-\rangle\langle\psi_\eta^-|\Big) \\ &+(1-\eta)\cos^2(\theta)|\uparrow\rangle\langle\uparrow|\,|0\rangle\langle 0|\,.\end{aligned} \tag{3.45}$$

STATES AFTER THE BELL MEASUREMENT

Now we need to perform a Bell measurement on the photonic qubits within the states ρ'_A and ρ'_B. Here we consider the scenario with non photon-number resolving detectors. Assuming for the moment the scenario without dark counts, we have at most two photons in the system. Hence we can consider three possible outcomes of our optical measurement: left detector clicked, right detector clicked, none of the detectors clicked. The measurement operators can be easily derived by noting that in our scenario without dark counts, each of the detectors can be triggered either by one or two photons and no cross-clicks between detectors are possible due to the photon-bunching effect. Then we can apply the reverse of the beam splitter mode transformations to the projectors on the events with one or two photons in each of the detectors to obtain these projectors in terms of the input modes. Finally we truncate the resulting projectors to the qubit space since in our scenario it is not possible for more than one photon to be present in each of the input modes of the beam splitter. In this way we obtain the following measurement operators

$$\begin{aligned}A_0 &= |\Psi^+\rangle\langle\Psi^+| + \frac{1}{\sqrt{2}}|11\rangle\langle 11|\,,\\ A_1 &= |\Psi^-\rangle\langle\Psi^-| + \frac{1}{\sqrt{2}}|11\rangle\langle 11|\,,\\ A_2 &= |00\rangle\langle 00|\,.\end{aligned} \tag{3.46}$$

These outcomes occur with the following probabilities,

$$p_0 = p_1 = \eta\cos^2(\theta)\left(1-\frac{\eta}{2}\cos^2(\theta)\right), \tag{3.47}$$

$$p_2 = (1-\eta\cos^2(\theta))^2 . \tag{3.48}$$

The post-measurement state of the two spins for the outcome A_0 is

$$\rho_0 = \frac{2\sin^2(\theta)}{2-\eta\cos^2(\theta)}\left(a|\Psi^+\rangle\langle\Psi^+| + b|\Psi^-\rangle\langle\Psi^-|\right) + \frac{\cos^2(\theta)(2-\eta)}{2-\eta\cos^2(\theta)}|\uparrow\uparrow\rangle\langle\uparrow\uparrow| . \tag{3.49}$$

Here

$$|\Psi^\pm\rangle = \frac{1}{\sqrt{2}}\left(|\downarrow\uparrow\rangle \pm |\uparrow\downarrow\rangle\right), \tag{3.50}$$

$$a = \lambda(F_{\text{prep}}^2 + (1-F_{\text{prep}})^2) + 2F_{\text{prep}}(1-F_{\text{prep}})(1-\lambda), \tag{3.51}$$

$$b = (1-\lambda)(F_{\text{prep}}^2 + (1-F_{\text{prep}})^2) + 2F_{\text{prep}}(1-F_{\text{prep}})\lambda . \tag{3.52}$$

For the outcome A_1 the post-measurement state of the spins is the same up to a local Z gate which Bob can apply following the trigger of the A_1 outcome. The post-measurement state of the spins for the outcome A_2, that is when none of the detector clicked, is

$$\rho_2 = \frac{1}{(1-\eta\cos^2(\theta))^2}\Big(\sin^4(\theta)|\downarrow\downarrow\rangle\langle\downarrow\downarrow| + (1-\eta)\cos^2(\theta)\sin^2(\theta)\left(|\downarrow\uparrow\rangle\langle\downarrow\uparrow| + |\uparrow\downarrow\rangle\langle\uparrow\downarrow|\right) + (1-\eta)^2\cos^4(\theta)|\uparrow\uparrow\rangle\langle\uparrow\uparrow|\Big) . \tag{3.53}$$

This is a separable state and so events corresponding to outcome A_2 (that is, no click in any of the detectors) will be discarded as failure. However, dark counts on our detectors can make us draw wrong conclusions about which of the three outcomes we actually obtained.

The effect of dark counts can be seen as follows

- We measured A_2 (no actual detection) but one of the detectors had a dark count. This event will happen with probability $2p_2p_d(1-p_d)$ and will make us accept the state ρ_2. Note that this is a classical state so application of the Z correction by Bob does not affect this state at all.
- We measured A_1 or A_2 but we also got a dark count in the other detector. This event will happen with probability $(p_0+p_1)\cdot p_d$. This will effectively lead us to rejection of the desired state ρ_0. Hence effectively ρ_0 will only be accepted if we measured A_1 or A_2 but the other detector did not have a dark count, which will happen with probability $(p_0+p_1)\cdot(1-p_d)$.

THE YIELD AND QBER

Taking dark counts into account, we see that the yield of the single-photon scheme, which is just the probability of registering a click in only one of the detectors, will be

$$Y = (p_0+p_1)(1-p_d) + 2p_2p_d(1-p_d) = 2(1-p_d)\left[\eta\cos^2(\theta)\left(1-\frac{\eta}{2}\cos^2(\theta)\right) + (1-\eta\cos^2(\theta))^2 p_d\right] .$$

The effective accepted state after a click in one of the detectors will then be

$$\rho_{\text{out}} = \frac{1}{Y}\left((p_0+p_1)(1-p_d)\rho_0 + 2p_2p_d(1-p_d)\rho_2\right) . \tag{3.54}$$

Note that both Alice and Bob perform a measurement on their electron spins immediately after each of the spin-photon entanglement generation events. This measurement causes an error

modelled as a depolarising channel of parameter F_m on each qubit, which means that after a successful run of the single-photon protocol, the effective state shared by Alice and Bob including the noise of their measurements will be given by

$$\rho_{AB} = F_m^2 \rho_{\text{out}} + (1-F_m)F_m \left[\frac{\mathbb{I}_{2,A}}{2} \otimes \text{tr}_A[\rho_{\text{out}}] + \text{tr}_B[\rho_{\text{out}}] \otimes \frac{\mathbb{I}_{2,B}}{2} \right] + (1-F_m)^2 \frac{\mathbb{I}_{4,AB}}{4} .$$

One can then extract the QBER for this state in all the three bases using the appropriate correlated/anti-correlated projectors such that:

$$e_z = \text{Tr}\big((|00\rangle\langle 00| + |11\rangle\langle 11|)\rho_{AB}\big) , \tag{3.55}$$

$$e_{xy} = \text{Tr}\big((|+-\rangle\langle +-| + |-+\rangle\langle -+|)\rho_{AB}\big) = \text{Tr}\big((|0_y 1_y\rangle\langle 0_y 1_y| + |1_y 0_y\rangle\langle 1_y 0_y|)\rho_{AB}\big) . \tag{3.56}$$

Here $|+\rangle$ and $|-\rangle$ denote the two eigenstates of X and $|0_y\rangle$ and $|1_y\rangle$ denote the two eigenstates of Y. We note that for our model of the single-photon scheme the QBER in X- and Y- bases are the same and therefore we denote both by a single symbol e_{xy}.

3.8.6. SPADS AND SPOTL SCHEMES ANALYSIS

In order to compute the quantum bit error rate (QBER) of the Single-Photon with Additional Detection Setup (SPADS) scheme and the Single-Photon Over Two Links (SPOTL) scheme, we derive step by step the quantum state shared between Alice and Bob. The following results have been found using Mathematica. Finally, we also calculate the yield of the SPADS and SPOTL schemes.

GENERATION OF ELEMENTARY LINKS

Single-photon scheme on Alice's side. The application of the single-photon scheme on Alice's side leads Alice and the quantum repeater to share a state given in Eq. (3.54). This state can be rewritten as

$$\rho_{\text{A-QR}^e} = A_1 |\Psi^+\rangle\langle\Psi^+| + B_1 |\Psi^-\rangle\langle\Psi^-| + C_1 (|10\rangle\langle 10| + |01\rangle\langle 01|) + D_1 |11\rangle\langle 11| + E_1 |00\rangle\langle 00| , \tag{3.57}$$

with $A_1 = A(\theta_A, Y_A)$, $B_1 = B(\theta_A, Y_A)$, $C_1 = C(\theta_A, Y_A)$, $D_1 = D(\theta_A, Y_A)$ and $E_1 = E(\theta_A, Y_A)$. Here we have that

$$A(\theta, Y) = \frac{1}{Y} 2\cos^2(\theta)\sin^2(\theta)\eta(1-p_d)\left[(F_{\text{prep}}^2 + (1-F_{\text{prep}})^2)\lambda + 2F_{\text{prep}}(1-F_{\text{prep}})(1-\lambda)\right] ,$$

$$B(\theta, Y) = \frac{1}{Y} 2\cos^2(\theta)\sin^2(\theta)\eta(1-p_d)\left[(F_{\text{prep}}^2 + (1-F_{\text{prep}})^2)(1-\lambda) + 2F_{\text{prep}}(1-F_{\text{prep}})\lambda\right] ,$$

$$C(\theta, Y) = \frac{2}{Y}\cos^2(\theta)\sin^2(\theta)p_d(1-p_d)(1-\eta) ,$$

$$D(\theta, Y) = \frac{1}{Y}\cos^4(\theta)\left(2(1-\eta)\eta(1-p_d) + \eta^2(1-p_d) + 2(1-\eta)^2 p_d(1-p_d)\right) ,$$

$$E(\theta, Y) = \frac{2}{Y}\sin^4(\theta)p_d(1-p_d) .$$

In the above, Y denotes the yield or the probability of success of the single-photon scheme and is given by Eq. (3.54). Subscript A indicates that in that expression for the yield and for each of the above defined coefficients we use $\theta = \theta_A$. Moreover, we have made here the following change of notation with respect to the Appendix 3.8.5, $|\downarrow\rangle \to |0\rangle$ and $|\uparrow\rangle \to |1\rangle$.

SWAP GATE IN THE MIDDLE NODE

In the next step a SWAP gate is applied in the middle node to transfer the electron state to the nuclear spin of the NV centre. This causes a depolarising noise of parameter $F_{\text{swap}} = F_g^2$ (see Appendix 3.8.1). The resulting state can then be written as

$$\rho_{\text{A-QR}^C} = F_{\text{swap}}\rho_{\text{A-QR}^e} + (1-F_{\text{swap}})\,\text{tr}_{\text{QR}}[\rho_{\text{A-QR}^e}] \otimes \frac{\mathbb{I}_{2,\text{QR}}}{2} . \tag{3.58}$$

THE PROCEDURE ON BOB'S SIDE

We now use the electron spin of the quantum repeater to generate the second quantum state. Here the procedures for the SPADS and SPOTL schemes diverge.

In the procedure for the SPADS scheme, the quantum repeater generates a spin-photon entangled state where the photonic qubit is encoded in the time-bin degree of freedom. Since the spin-photon entangled state is imperfect, the electron and the photon share a state

$$\rho_{\mathrm{QR}^e-B} = F_{\mathrm{prep}} \left|\Psi^+\right\rangle\!\left\langle\Psi^+\right| + (1-F_{\mathrm{prep}}) \left|\Psi^-\right\rangle\!\left\langle\Psi^-\right| . \tag{3.59}$$

Here we use the following labeling for time-bin encoded early and late mode of the photon: $|e\rangle = |1\rangle, |l\rangle = |0\rangle$. This photon is then sent towards Bob's detector. The lossy channel acts on such a time-bin encoded qubit as an erasure channel and so the quantum spin-photon state of the successful events in which the photonic qubit successfully arrives at the detector is unaffected by the lossy channel.

For the SPOTL scheme the repeater's electron spin and Bob's quantum memory generate a second state of the form given in Eq. (3.54). We can rewrite this state as

$$\rho_{\mathrm{QR}^e-B} = A_2 \left|\Psi^+\right\rangle\!\left\langle\Psi^+\right| + B_2 \left|\Psi^-\right\rangle\!\left\langle\Psi^-\right| + C_2 \left(|10\rangle\langle 10| + |01\rangle\langle 01|\right) + D_2 |11\rangle\langle 11| + E_2 |00\rangle\langle 00| ,$$

with $A_2 = A(\theta_B, Y_B)$, $B_2 = B(\theta_B, Y_B)$, $C_2 = C(\theta_B, Y_B)$, $D_2 = D(\theta_B, Y_B)$ and $E_2 = E(\theta_B, Y_B)$.

DECOHERENCE IN THE QUANTUM MEMORIES

Decoherence of the carbon spin in the middle node can be modelled identically for both the SPADS and SPOTL scheme.

During the $n < n^*$ attempts to generate the state $\rho_{\mathrm{QR}^e\text{-B}}$, the carbon spin in the middle node holding half of the state $\rho_{\mathrm{A\text{-}QR}^C}$ will decohere. Using the decoherence model discussed in Appendix 3.8.2, decoherence of the carbon spin will thus give us

$$\begin{aligned}\rho'_{\mathrm{A\text{-}QR}^C} =& F_{T_1}\left(F_{T_2}\rho_{\mathrm{A\text{-}QR}^C} + (1-F_{T_2})(\mathbb{I}_2\otimes Z)\rho_{\mathrm{A\text{-}QR}^C}(\mathbb{I}_2\otimes Z)^\dagger\right)\\ &+(1-F_{T_1})\,\mathrm{tr}_{\mathrm{QR}}[\rho_{\mathrm{A\text{-}QR}^C}]\otimes\frac{\mathbb{I}_{2,\mathrm{QR}}}{2} .\end{aligned} \tag{3.60}$$

For key generation, Alice (SPADS and SPOTL schemes) and Bob (SPOTL scheme) can actually measure their electron spin(s) immediately after the generation of spin photon entanglement, preventing the effect of decoherence on these qubit(s).

NOISE DUE TO MEASUREMENTS

MEASUREMENT OF THE QUBITS OF ALICE AND BOB

In the SPADS scheme Alice performs a measurement on her electron spin immediately after each of the spin-photon entanglement generation events to prevent any decoherence with time of this qubit. This measurement causes an error modelled as a depolarising channel of parameter F_m. Bob on the other hand performs a measurement on a photonic qubit that is encoded in the time-bin degree of freedom. His measurement utilises the squashing map so that we can model the noise arising from this measurement as a depolarising channel with parameter α_B as described in Appendix 3.8.1. Hence the total state just before the Bell measurement is given by

$$\begin{aligned}\rho_{A-QR-B} =& F_m\alpha_B\rho'_{\mathrm{A\text{-}QR}^C}\otimes\rho_{\mathrm{QR}^e-B} + (1-F_m)\alpha_B\frac{\mathbb{I}_{2,A}}{2}\otimes\mathrm{tr}_A[\rho'_{\mathrm{A\text{-}QR}^C}]\otimes\rho_{\mathrm{QR}^e-B}\\ &+(1-\alpha_B)F_m\rho'_{\mathrm{A\text{-}QR}^C}\otimes\mathrm{tr}_B[\rho_{\mathrm{QR}^e-B}]\otimes\frac{\mathbb{I}_{2,B}}{2}\\ &+(1-F_m)(1-\alpha_B)\,\mathrm{tr}_{AB}[\rho'_{\mathrm{A\text{-}QR}^C}\otimes\rho_{\mathrm{QR}^e-B}]\otimes\frac{\mathbb{I}_{4,AB}}{4} .\end{aligned} \tag{3.61}$$

For the SPOTL scheme, both Alice and Bob perform a measurement on their electron spins immediately after each of the spin-photon entanglement generation events. This measurement causes an error modelled as a depolarising channel of parameter F_m on each qubit, which means that after both Alice and Bob succeeded in performing the single-photon scheme with the repeater, the total, four-qubit state just before the Bell-measurement and including the noise of the measurements of Alice and Bob will be given by

$$\begin{aligned}\rho_{A-QR-B} = {} & F_m^2 \rho'_{\text{A-QR}^C} \otimes \rho_{\text{QR}^e-B} \\ & + (1-F_m)F_m \left[\frac{\mathbb{I}_{2,A}}{2} \otimes \text{tr}_A[\rho'_{\text{A-QR}^C}] \otimes \rho_{\text{QR}^e-B} + \rho'_{\text{A-QR}^C} \otimes \text{tr}_B[\rho_{\text{QR}^e-B}] \otimes \frac{\mathbb{I}_{2,B}}{2} \right] \\ & + (1-F_m)^2 \,\text{tr}_{AB}[\rho'_{\text{A-QR}^C} \otimes \rho_{\text{QR}^e-B}] \otimes \frac{\mathbb{I}_{4,AB}}{4} \,. \end{aligned} \tag{3.62}$$

BELL STATE MEASUREMENT

Before the entanglement swapping, we have a total state ρ_{A-QR-B}. We now perform a Bell state measurement on the two qubits in the middle node. The error coming from this measurement is modelled by concatenation of depolarising channels (see Appendix 3.8.1) which means that the measurement is actually performed on

$$\rho_{\text{fin}} = F_g^2 F_m^2 \rho_{A-QR-B} + F_g^2(1-F_m^2)\,\text{tr}_{QR^e}[\rho_{A-QR-B}] \otimes \frac{\mathbb{I}_{2,QR^e}}{2} + (1-F_g^2)\,\text{tr}_{QR}[\rho_{A-QR-B}] \otimes \frac{\mathbb{I}_{4,QR}}{4} \,.$$

While $\rho'_{\text{A-QR}^C}$ is not Bell diagonal for the SPADS scheme, ρ_{QR^e-B} is, and so we find that taking into account the classical correction (which will be performed on the measured bit-value by Alice and Bob) the four cases corresponding to different measurement outcomes are equivalent. This means that if we model the correction to be applied to the quantum state rather than the classical bit, then the four post-measurement bipartite states shared between Alice and Bob are exactly the same.

For the SPOTL scheme, both $\rho'_{\text{A-QR}^C}$ and ρ_{QR^e-B} are not Bell diagonal which means that the resulting state of qubits of Alice and Bob after the Bell state measurement depends on the outcome of this Bell measurement and those four corresponding states are not equivalent under local unitary corrections. In fact, the two states corresponding to the $\Phi^\pm$ outcomes and the two states corresponding to the $\Psi^\pm$ outcomes are pairwise equivalent under local Pauli corrections. Hence, we will derive two different QBER corresponding to the following different resulting states shared between Alice and Bob,

$$\rho_{\Phi,AB} = (\mathbb{I}_A \otimes U_{\Phi^\pm,B})\,\text{Tr}_{QR}\left[\frac{(\mathbb{I}\otimes|\Phi^\pm\rangle\langle\Phi^\pm|\otimes\mathbb{I})\rho_{\text{fin}}(\mathbb{I}\otimes|\Phi^\pm\rangle\langle\Phi^\pm|\otimes\mathbb{I})^\dagger}{\text{Tr}\big(\rho_{\text{fin}}(\mathbb{I}\otimes|\Phi^\pm\rangle\langle\Phi^\pm|\otimes\mathbb{I})\big)} \right](\mathbb{I}\otimes U_{\Phi^\pm,B})^\dagger \,,$$

$$\rho_{\Psi,AB} = (\mathbb{I}_A \otimes U_{\Psi^\pm,B})\,\text{Tr}_{QR}\left[\frac{(\mathbb{I}\otimes|\Psi^\pm\rangle\langle\Psi^\pm|\otimes\mathbb{I})\rho_{\text{fin}}(\mathbb{I}\otimes|\Psi^\pm\rangle\langle\Psi^\pm|\otimes\mathbb{I})^\dagger}{\text{Tr}\big(\rho_{\text{fin}}(\mathbb{I}\otimes|\Psi^\pm\rangle\langle\Psi^\pm|\otimes\mathbb{I})\big)} \right](\mathbb{I}\otimes U_{\Psi^\pm,B})^\dagger \,.$$

Here $U_{\Phi^\pm,B}$ and $U_{\Psi^\pm,B}$ denote the four Pauli corrections implemented by Bob after the corresponding outcome of the Bell measurement. Note that for the SPADS scheme $\rho_{\Phi,AB} = \rho_{\Psi,AB}$.

THE YIELD AND QBER

YIELD

For both SPADS and SPOTL scheme we calculate the yield as the inverse of the number of channel uses required to generate one bit of raw key, $Y = 1/\mathbb{E}[N]$, where $\mathbb{E}[N]$ is given by Eq. (3.34). For the

SPOTL scheme in that formula we use $p_{A/B} = Y_{A/B}$, where $Y_{A/B}$ denotes the yield of the single-photon scheme on Alice's/Bob's side given by Eq. (3.54). For the SPADS scheme p_A takes the same form as for the SPOTL scheme (but is now calculated for two thirds of the total distance between Alice and Bob rather than half), while p_B is the probability of registering a click in Bob's optical detection setup as in the SiSQuaRe scheme.

EXTRACTION OF THE QUBIT ERROR RATES

By projecting these final corrected states onto the correct subspaces, we can obtain the qubit error rates e_z and e_{xy} (with our model we find that for both SPADS and SPOTL schemes the error rates in X and Y bases are the same). The state shared between Alice and Bob after the Pauli correction will always be the same for the SPADS scheme. Thus, there is only a single QBER e_z and e_{xy} independently of the outcome of the Bell measurement. For the SPOTL scheme that is not the case, there will be two set of QBER corresponding to the states $\rho_{\Phi,AB}$ and $\rho_{\Psi,AB}$.

$$e_{z,\Phi} = \mathrm{Tr}\big((|00\rangle\langle 00| + |11\rangle\langle 11|)\rho_\Phi\big)\,, \tag{3.63}$$

$$e_{z,\Psi} = \mathrm{Tr}\big((|00\rangle\langle 00| + |11\rangle\langle 11|)\rho_\Psi\big)\,, \tag{3.64}$$

$$e_{xy,\Phi} = \mathrm{Tr}\big((|{+}{-}\rangle\langle{+}{-}| + |{-}{+}\rangle\langle{-}{+}|)\rho_\Phi\big) = \mathrm{Tr}\big((|0_y1_y\rangle\langle 0_y1_y| + |1_y0_y\rangle\langle 1_y0_y|)\rho_\Phi\big)\,, \tag{3.65}$$

$$e_{xy,\Psi} = \mathrm{Tr}\big((|{+}{-}\rangle\langle{+}{-}| + |{-}{+}\rangle\langle{-}{+}|)\rho_\Psi\big) = \mathrm{Tr}\big((|0_y1_y\rangle\langle 0_y1_y| + |1_y0_y\rangle\langle 1_y0_y|)\rho_\Psi\big)\,. \tag{3.66}$$

Again, for the SPADS scheme $e_{z,\Phi} = e_{z,\Psi} = e_z$ and $e_{xy,\Phi} = e_{xy,\Psi} = e_{xy}$.

AVERAGING THE QUBIT ERROR RATES

We have now derived the qubit error rates as a function of the experimental parameters. For the SPOTL scheme we now average the QBER over the two outcomes to get the final average QBER

$$\langle e_z\rangle = \langle p_\Psi e_{z,\Psi} + p_\Phi e_{z,\Phi}\rangle\,, \tag{3.67}$$

$$\langle e_{xy}\rangle = \langle p_\Psi e_{xy,\Psi} + p_\Phi e_{xy,\Phi}\rangle\,, \tag{3.68}$$

where p_Ψ (p_Φ) is the probability of measuring one of the $|\Psi\rangle$ ($|\Phi\rangle$) states in the Bell measurement and $\langle ...\rangle$ is found by averaging the expression over the number of Bob's attempts n with the geometric distribution within the first n^* trials. For the SPADS scheme $\langle e_z\rangle$ and $\langle e_{xy}\rangle$ can be averaged directly. The dependence on n arises from the decoherence terms F_{T_1} and F_{T_2}. Indeed, those terms correspond to the decoherence in the middle node during the attempts on Bob's side. Denoting by p_B the probability that in a single attempt Bob generates entanglement with the quantum repeater using the single-photon scheme for the SPOTL scheme and using direct transmission of the time-bin encoded qubit from the repeater to Bob for the SPADS scheme, we have that the exponentials in those expressions can be averaged as follows (see preceding chapter),

$$\langle e^{-cn}\rangle = \frac{p_B e^{-c}}{1-(1-p_B)^{n^*}}\,\frac{1-(1-p)^{n^*}e^{-cn^*}}{1-(1-p_B)e^{-c}}\,. \tag{3.69}$$

3.8.7. SECRET-KEY FRACTION AND ADVANTAGE DISTILLATION

In this section we review the formulas for the secret-key fraction for the QKD protocols used in our model as a function of the QBER. In particular, we focus on the case where the QBER is not the same for all bases. In such a scenario, it is sometimes possible to extract more secret-key using the six-state protocol with advantage distillation by changing the basis in which one extracts key.

ONE-WAY BB84 PROTOCOL

For the fully asymmetric BB84 protocol with standard one-way post-processing, the secret-key fraction is given by [95, 148]:

$$r = 1 - h(e_x) - h(e_z)\,, \tag{3.70}$$

where $h(x)$ is the binary entropy function. Note that this formula is symmetric under the exchange of e_x and e_z - that is, the secret-key fraction is the same independently of whether we extract the key in the Z- or X-basis. As we will see later in this section, this is not the case for the six-state protocol with advantage distillation.

SIX-STATE PROTOCOL WITH ADVANTAGE DISTILLATION

In Appendix 2.9.4 from the preceding chapter we considered the secret-key that can be extracted using the advantage distillation scheme presented in [171]. It is important to note that for the advantage distillation scheme considered in this chapter, the amount of generated secret key depends on the basis in which it is extracted, as has been shown in [114]. This should be contrasted with the previous chapter, where the used secret-key fraction assumed the key was extracted in the Z-basis. This thus allows for a potential increase in the extractable secret-key. Let us now have a look at the amount of key that can be extracted in the X- and Y-bases. The secret-key fraction in these cases is also given by Eq. (2.39) but now the Bell coefficients depend on QBER in the following way [114],

$$\begin{aligned} p_{00} &= 1 - \frac{e_z}{2} - e_{xy}\,, \\ p_{10} &= e_{xy} - \frac{e_z}{2}\,, \\ p_{01} &= p_{11} = \frac{e_z}{2}\,. \end{aligned} \tag{3.71}$$

And so

$$\begin{aligned} P_{\tilde{\mathrm{X}}}(0) &= 1 - 2e_{xy} + 2e_{xy}^2\,, \\ P_{\tilde{\mathrm{X}}}(1) &= 2(1 - e_{xy})e_{xy}\,. \end{aligned} \tag{3.72}$$

We note that we have assumed here that in the case of key extraction in Y-basis, either Alice or Bob applies a local bit flip in the Y-basis to the shared state, as the target state $|\psi(0,0)\rangle$ is anti-correlated in that basis.

In [114] it has been also observed that in the considered case of having the QBER in the X- and Y-bases being equal, the six-state protocol with advantage distillation allows us to extract more key if it is extracted in the basis with higher QBER. This observation determines the basis that we use for extracting key for the single-photon and the SPOTL schemes that use fully asymmetric six-state protocol with advantage distillation. Specifically, for the single-photon scheme we observe higher QBER in the Z-basis, while for the SPOTL scheme the QBER is higher in the X- and Y-bases. Therefore these are the bases that we choose to use for extracting key for those schemes.

For the SiSQuaRe and SPADS schemes the symmetric six-state protocol is used, hence for those schemes we group the raw bits into three groups corresponding to three different key-extraction bases and we extract the key separately for each of these bases. Finally, to obtain the final secret-key fraction, we note that for the symmetric six-state protocol we also need to include sifting, that is only one third of all the raw bits were obtained by Alice and Bob measuring in the same basis (the raw bits for the protocol runs in which they measured in different bases are discarded). Hence, if we denote by r_i the secret-key fraction obtained from the group of raw bits in which both Alice

and Bob measured in the basis i, the final secret-key fraction for the six-state protocol for those schemes is given by

$$r = \frac{1}{3}\left(\frac{1}{3}r_x + \frac{1}{3}r_y + \frac{1}{3}r_z\right). \tag{3.73}$$

Clearly in our case we have $r_x = r_y = r_{xy}$.

ONE-WAY SIX-STATE PROTOCOL

In Figure 3.6 we have also plotted the secret-key fraction for the one-way six-state protocol. For the fully asymmetric protocol and the case in which the key is extracted in the Z-basis, it is given by [148]

$$r = 1 - e_z h\left(\frac{1+(e_x-e_y)/e_z}{2}\right) - (1-e_z)h\left(\frac{1-(e_x+e_y+e_z)/2}{1-e_z}\right) - h(e_z). \tag{3.74}$$

Although this formula does not appear to be symmetric under the permutation of e_x, e_y, e_z, it is in fact invariant under this permutation [137]. This means that for the symmetric one-way six-state protocol, in our case the final secret-key fraction is given by the expression in Eq. (3.74) multiplied by the sifting efficiency of one-third.

3.8.8. RUNTIME OF THE EXPERIMENT

In this section we will detail how to perform an experiment that will be able to establish that a setup can surpass the capacity of a quantum channel modeling losses in a fibre (see Eq. (2.15)). This experiment can validate a setup to qualify as a quantum repeater, without explicitly having to generate secret-key. We show then that, for the listed parameters in the main text, the single-photon scheme can be certified to be a quantum repeater within approximately twelve hours.

The experiment is based on estimating the yield of the scheme and the individual QBER of the generated states. More specifically, here we will calculate the probability that, assuming our model is accurate and each individual run is independent and identically distributed, the observed estimate of the yield and the individual QBER are larger and smaller, respectively, than some fixed threshold values. If, with these threshold values for the yield and QBER, the calculated asymptotic secret-key still surpasses the capacity, we can claim a working quantum repeater. The experiment consists of first performing n attempts at generating a state between Alice and Bob, from which the yield can be estimated by calculating the ratio of the successful attempts and n. Then, the QBER in each basis is estimated by Alice and Bob measuring in the same basis in each of the successful attempts.

Central to our calculation is the fact that, for n instances of a Bernoulli random variable with probability p, the probability that the number of observed successes $S(n)$ is smaller or equal than some value k is equal to

$$P(S(n) \leq k) = \sum_{i=0}^{k} \binom{n}{i} p^i (1-p)^{n-i}. \tag{3.75}$$

Assuming the outcomes of our experiment are independent and identically distributed, the observed yield $\bar{Y}$ satisfies

$$P\left(\bar{Y} \leq (Y - t_Y)\right) = P\left(n\bar{Y} \leq n(Y - t_Y)\right) = \sum_{i=0}^{\lfloor n(Y-t_Y)\rfloor} \binom{n}{i} Y^i (1-Y)^{n-i}, \tag{3.76}$$

where $Y - t_Y$ is the lower threshold. Let us make this more concrete with a specific calculation. For a distance of $17L_0$ the yield is equal to $\approx 5.6 \cdot 10^{-6}$. Setting the maximum deviation in the yield to

$\bar{Y} = Y - t_Y$ with $t_Y = 2.0 \cdot 10^{-7}$ and the number of attempts to $n = 5 \cdot 10^9$ (which corresponds to approximately a runtime of twelve hours assuming a single attempt takes $8.5 \cdot 10^{-6}$ s, corresponding to t_{prep} and a single-shot readout lasting $2.5 \cdot 10^{-6}$ s), we find that

$$P\left(\bar{Y} \leq (Y - t_Y)\right) \leq 9.2 \cdot 10^{-10} . \tag{3.77}$$

Similarly, for the individual errors $\{e_k\}_{k\in\{x,y,z\}}$ in the three bases we have that

$$P\left(\bar{e}_k \geq \left(e_k + t_k\right)\right) = P\left(m \cdot \bar{e}_k \geq m\left(e_k + t_k\right)\right) = \sum_{i=\lceil m(e_k+t_k)\rceil}^{m} \binom{m}{i} \left(e_k\right)^i \left(1 - e_k\right)^{m-i} .$$

Here we set $m = \left\lfloor \frac{n}{3} (Y - t_Y) \right\rfloor$, which is an estimate for the number of raw bits that Alice and Bob obtain from measurements in each of the three bases, for the total n attempts of the protocol. All the raw bits from those three sets are then compared to estimate the QBER in each of the three bases. Note that we gather the same amount of samples for each basis, even when an asymmetric protocol would be performed. Setting $t_i = t = 0.015$, $\forall i \in \{x, y, z\}$ and, as before, $n = 5 \cdot 10^9$, we find, at a distance of $17L_0$ where $e_z \approx 0.171$ and $e_y = e_x \approx 0.141$, that

$$P\left(\bar{e}_z \geq (e_z + t)\right) \leq 9.0 \cdot 10^{-5} , \tag{3.78}$$

$$P\left(\bar{e}_y \geq \left(e_y + t\right)\right) = P\left(\bar{e}_x \geq (e_x + t)\right) \leq 2.7 \cdot 10^{-5} . \tag{3.79}$$

Then, with probability at least

$$\begin{aligned} &\left(1 - P\left(\bar{e}_x \geq (e_x + t)\right)\right) \cdot \left(1 - P\left(\bar{e}_y \geq \left(e_y + t\right)\right)\right) \cdot \left(1 - P\left(\bar{e}_z \geq (e_z + t)\right)\right) \cdot \left(1 - P\left(\bar{Y} \leq (Y - t_Y)\right)\right) \\ &\qquad \geq 1 - 1.5 \cdot 10^{-4} , \end{aligned} \tag{3.80}$$

none of the observed QBER and yield exceed their threshold conditions. The corresponding lowest secret-key rate for these parameters (with a yield of $Y - t_Y$ and QBER of $e_x + t_x$, $e_y + t_y$, $e_z + t_z$) is $\approx 1.97 \cdot 10^{-7}$, which we observe is greater than the secret-key capacity by a factor ≈ 3.29 (see Eq.(2.15)) at a distance of $17L_0$, since the secret-key capacity equals $-\log_2\left(1 - e^{-17}\right) \lesssim 5.97 \cdot 10^{-8}$.

Thus, with high probability we can establish that the single-photon scheme achieves a secret-key rate significantly greater than the corresponding secret-key capacity for a distance of $17L_0 \approx 9.2$ kilometer within approximately twelve hours.

3.8.9. MDI QKD

We note here that the single-photon scheme for generating key is closely linked to the measurement device independent (MDI) QKD protocol [96]. In particular it is a version of a scheme in which Alice and Bob prepare and send specific photonic qubit states to the heralding station in the middle, where the qubits are encoded in the presence/absence of the photon. We note that in the ideal case of the single-photon scheme, the spin-photon state is given in Eq. (3.38). For the six-state protocol the spin part of this state is then measured in the X-, Y- or Z- basis at random according to a fixed probability distribution (this probability distribution dictates whether we use symmetric or asymmetric protocol). Considering the probabilities of the individual measurement outcomes, this is equivalent to the scenario in which Alice and Bob choose one of the three set of states at random according to the same probability distribution and prepare each of the two states from that set with the probability equal to the corresponding measurement outcome probability. These sets do not form bases, as the two states within each set are not orthogonal. We will therefore refer to these sets here as "pseudo-bases". Depending on the chosen pseudo-basis they prepare one of the six states encoding the bit value of "0" or "1" in that pseudo-basis. These states and the corresponding preparation probabilities are

- pseudo-basis 1: $\{|0\rangle, |1\rangle\}$ with probabilities $\{\sin^2\theta, \cos^2\theta\}$,
- pseudo-basis 2: $\{\sin\theta|0\rangle + \cos\theta|1\rangle, \sin\theta|0\rangle - \cos\theta|1\rangle\}$ with probabilities $\{\frac{1}{2}, \frac{1}{2}\}$,
- pseudo-basis 3: $\{\sin\theta|0\rangle + i\cos\theta|1\rangle, \sin\theta|0\rangle - i\cos\theta|1\rangle\}$ with probabilities $\{\frac{1}{2}, \frac{1}{2}\}$.

These states are then sent towards the beam splitter station. The station performs the standard photonic Bell-state measurement and sends the outcome to both Alice and Bob. Alice and Bob discard all the runs for which the beam splitter station measured A_2 (recall the measurement operators in Eq. (3.46)). They then exchange the classical information about their pseudo-basis choice and keep only the data for the runs in which they both used the same basis. For those data they apply the following post-processing in order to obtain correlated raw bits

- pseudo-basis 1: for both outcomes A_0 and A_1 Bob flips the value of his bit.
- pseudo-basis 2: for the outcome A_0 they do nothing, for the outcome A_1 Bob flips the value of his bit.
- pseudo-basis 3: for the outcome A_0 they do nothing, for the outcome A_1 Bob flips the value of his bit.

In this way Alice and Bob have generated their strings of raw bits.

We note here that the direct preparation of the six states from the three pseudo-bases described above in the photonic presence/absence degree of freedom is experimentally hard. This is related to the fact that linear optics does not allow to easily perform single qubit rotations necessary to prepare these states. The use of memory-based NV-centres offers a great advantage here, as in these schemes the rotations that allow us to obtain the required amplitudes of the photonic states are performed on the electron spins rather than the photons themselves. There has also been proposed an alternative scheme that also benefits from single photon detection events in which Alice and Bob send coherent pulses to the heralding station [98, 157].

4

OPTIMISING REPEATER SCHEMES FOR THE QUANTUM INTERNET

Kenneth Goodenough*, David Elkouss and Stephanie Wehner

The rate at which quantum communication tasks can be performed using direct transmission is fundamentally hindered by the channel loss. Quantum repeaters allow, in principle, to overcome these limitations, but their introduction necessarily adds an additional layer of complexity to the distribution of entanglement. This additional complexity - along with the stochastic nature of processes such as entanglement generation, Bell swaps, and entanglement distillation - makes finding good quantum repeater schemes non-trivial. We develop an algorithm that can efficiently perform a heuristic optimisation over a subset of quantum repeater schemes for general repeater platforms. We find a strong improvement in the generation rate in comparison to an optimisation over a simpler class of repeater schemes based on BDCZ repeater schemes. We use the algorithm to study three different experimental quantum repeater implementations on their ability to distribute entanglement, which we dub information processing *implementations,* multiplexed elementary pair generation *implementations, and combinations of the two. We perform this heuristic optimisation of repeater schemes for each of these implementations for a wide range of parameters and different experimental settings. This allows us to make estimates on what are the most critical parameters to improve for entanglement generation, how many repeaters to use, and which implementations perform best in their ability to generate entanglement.*

This chapter has been adapted from the following publication: Phys. Rev. A 103, 032610.

4.1. INTRODUCTION

In the two preceding Chapters 2 and 3 we performed a thorough analysis of proof-of-principle repeater schemes, consisting of one or two nodes. Such setups will still be limited in the maximum distance at which shared entanglement can be generated (at a reasonable rate). Our aim for this chapter is to consider multiple such nodes, leading to *linear quantum repeater chains*. The corresponding quantum repeater schemes are built on the concept of breaking the total length between Alice and Bob into several shorter (elementary) *links*. Depending on the scheme, the nodes have different requirements ranging from storage of quantum states to full-fledged quantum computation. We note here that we will consider more complex schemes than those in chapters 2 and 3., but at the same time our analysis will not be as fine-grained as the analysis performed before.

By generating and storing entanglement over the elementary links and performing Bell state measurements on the locally held states, the distance over which entanglement is present can be increased, until the two parties at the end are entangled [20, 29, 30, 48].

However, the imperfect operations during this process lower the quality of the entanglement, potentially ruining the benefits of utilising quantum repeater nodes. The effects of noise can be counteracted by using *entanglement distillation*, which can (possibly probabilistically) turn multiple entangled pairs of lower fidelity into a smaller amount of pairs with higher fidelity [9, 11, 83].

An entanglement generation scheme between two spatially separated parties Alice and Bob consists of the generation of entanglement over elementary links, entanglement swaps and distillation. Our goal is to find schemes that minimise the generation time of the entanglement between Alice and Bob for a given fidelity to the maximally entangled state in a suitable experimental model. However, finding optimal schemes is non-trivial for two reasons. First, the amount of schemes that can be performed grows super-exponentially in the number of elementary links/nodes, making a full systematic optimisation infeasible (see [77] and Appendix 4.6.1). Second, entanglement generation, Bell state measurements, and distillation are all processes that are in general probabilistic. Finding the corresponding probability distributions is believed to be computationally intensive [17, 147, 150, 168].

For the reasons mentioned above, it seems necessary to either approximate or simplify the problem. Notably, in [77], an algorithm based on dynamical programming was proposed capable of efficiently optimising repeater schemes over the full parameter space. Under the heuristic approximation that all processes finish at the average time and there is no decoherence over time in the quantum memories, the algorithm constructs the scheme for a large chain combining the optimal solutions over smaller (multi-hop) links.

We take a different route. Instead of approximating the behaviour of the schemes by the mean, we simplify the problem by considering a relevant subset of schemes. In particular, we consider schemes that succeed at all levels near-deterministically. Such schemes have the benefit of having a small variance of the fidelity and generation time. We note that the requirement of being near-deterministic does not imply that our algorithm cannot handle non-deterministic processes. High success probabilities can be enforced even when certain processes are not deterministic - in that case, the probability of a single success can be increased by repeating the process a number of times, ensuring that the whole process can be made near-deterministic, see Section 4.2 for further details. Furthermore, this allows us to calculate the success probability of a scheme exactly, even when more complicated protocols such as distillation and probabilistic swapping are performed. Finally, this approach also allows us to calculate the average noise experienced during storage, in contrast to [77], see Appendix 4.6.4.

In this chapter, we detail an algorithm (publicly available as a Python script at [58]) that performs a heuristic optimisation over the set of near-deterministic schemes. The optimisation runs in $\mathcal{O}\left(n^2 \log(n)\right)$ time, and $\mathcal{O}\left(n \log(n)\right)$ time if all the nodes have the same parameters and are

equidistant, where n is the number of elementary links. Concretely, the input to our algorithm is given by the experimental parameters of the nodes and connecting fibres, the distances between adjacent nodes, the possible protocols for elementary pair generation, swapping and distillation, and a set of algorithm-specific parameters, see Section 4.2.4.3. The algorithm returns a collection of optimised schemes for generating entanglement between Alice and Bob.

We exploit the fact that our algorithm is not specific to any particular experimental setup, which allows for the optimisation over repeater schemes for several types of platforms.

The experimental platforms that we consider can be split up into three types:

- *Information processing platforms* - Information processing (IP) implementations have the ability to store quantum states and perform operations on them, such that it is possible to perform distillation. However, the number of quantum states that can be processed at the same time is presently limited to a small number. Examples of information processing implementations include systems such as trapped ions [15, 73, 109], nitrogen-vacancy centres in diamond [28, 163], neutral atoms [134, 151, 174], and quantum dots [64, 179].

- *Multiplexed elementary pair generation platforms* - Multiplexed elementary pair generation (MP) implementations lack the ability to properly perform operations on the stored states, prohibiting distillation. However, a large number ($10^4 - 10^7$) of states can potentially be generated, transmitted and stored simultaneously with such implementations, effectively increasing the success probability for the elementary pair generation. Examples of such implementations include the different types of atomic ensembles [31, 90, 104].

- *A combination of IP and MP platforms* - Multiplexed elementary pair generation platforms can overcome the effects of losses over the elementary links more easily than information processing platforms, but suffer from the lack of control and long coherence times available to information processing platforms. This motivates a combination of the two. That is, the elementary pair generation is performed with an MP implementation, after which the quantum state is transferred into an information processing system. Such a combined setup benefits from the high success probability of the generation of the elementary pairs, together with the ability to perform entanglement distillation and longer coherence times.

We find that the optimisation returns schemes that outperform a simplified optimisation over more structured schemes, similar to those in [20, 29, 30, 48]. This highlights the complexity of repeater protocols for realistic repeater chains and the non-trivial nature of the optimisation problem. With such optimised schemes in hand we use our algorithm to study a range of questions, such as which setups hold promise for near-term quantum networks, how many nodes should be implemented, and which experimental parameters are the most important to improve upon.

In Section 4.2 we detail the basics of our algorithm, which takes as input an arbitrary repeater chain configuration, and returns a collection of heuristically optimised schemes which generate entanglement between two specified nodes, i.e. the schemes have an optimal trade-off between the fidelity and generation time (over the set of considered schemes). This section also contains the heuristics we use to reduce the search space/complexity of the algorithm in Section 4.2.4 (with further details in Appendices 4.6.1 and 4.6.3 regarding the complexity/runtime) and closes with the pseudocode of our algorithm in Section 4.2.4.3. Section 4.3 contains an overview of how we model the three experimental platforms considered in this chapter, namely information processing (Section 4.3.1) implementations, multiplexed (Section 4.3.2) implementations, and a combination of the two (Section 4.3.3). We then use the algorithm to heuristically optimise over repeater schemes for each of the implementations for several different scenarios in Section 4.4. We close with a discussion of the results and the algorithm in Section 4.5.

Figure 4.1: Elementary pair generation (EPG) between adjacent nodes QR_i and QR_{i+1}. The schemes take a number of rounds $r = r^*$, even if entanglement is generated at an earlier round. See main text for further details.

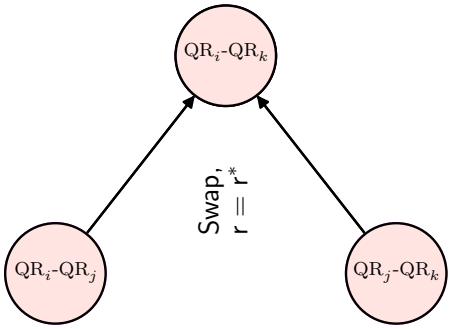

Figure 4.2: Entanglement swapping between two entangled pairs between (multi-hop) links (QR_i, QR_j) and (QR_j, QR_k), indicated by a circle node. By performing a Bell state measurement on the two local states at QR_j, the two entangled states turn into one entangled state between (QR_i, QR_k). The schemes take a number of rounds $r = r^*$ even if the scheme succeeds at an earlier round, see main text for further details. Note that the distances over which the entanglement has been generated for the (multi-hop) links (QR_i, QR_j) and (QR_j, QR_k) need not be the same.

4.2. ALGORITHM DESCRIPTION

In this section we first explain the general structure of quantum repeater schemes (Section 4.2.1). We then focus on the construction of so-called near-deterministic schemes (Section 4.2.2). Afterwards, we first detail a non-scalable brute-force algorithm for optimising over such near-deterministic schemes (Section 4.2.3), after which we provide a feasible algorithm by implementing certain heuristics into the brute-force algorithm (Section 4.2.4). Appendices 4.6.1 and 4.6.3 contain a more explicit discussion regarding the complexity/runtime with and without the heuristics implemented.

4.2.1. STRUCTURE OF QUANTUM REPEATER SCHEMES

The goal of a quantum repeater scheme is to distribute an entangled state between two remote parties Alice and Bob. Quantum repeater schemes are built up from smaller schemes. Schemes are constructed by performing *connection* and *distillation* protocols on pairs of smaller schemes.

Connection protocols extend the range over which entanglement exists. This can be done by elementary pair generation and entanglement swapping. Elementary pair generation (EPG) creates entanglement over elementary links, see Fig. 4.1. Entanglement swapping transforms two entangled states over two shorter (multi-hop) links to an entangled state over a longer multi-hop link using a Bell state measurement, see Fig. 4.2.

Distillation protocols allow to (possibly probabilistically) convert two entangled states to a single, more entangled state using only local operations and classical communication [11, 41]. There

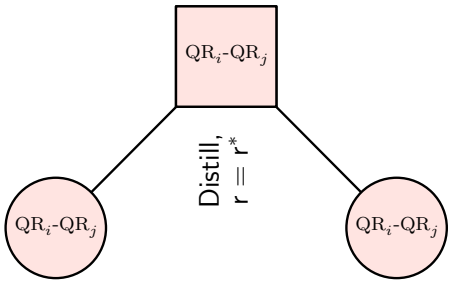

Figure 4.3: Example of a generic entanglement distillation protocol, transforming (possibly probabilistically) two entangled states to a single, more entangled state between nodes QR$_i$ and QR$_j$, using only local operations and classical communication. Distillation is indicated by a square node. The schemes take a number of rounds $r = r^*$ even if distillation succeeds at an earlier round, see main text for further details. Note that QR$_i$ and QR$_j$ do not have to be directly connected by a fibre.

exist more complicated protocols, where an arbitrary number of entangled states are converted to a smaller number of entangled states [39, 69]. Here, we only consider distillation protocols taking two states to a single one [1]. See Fig. 4.3 for an illustration of a distillation protocol.

4.2.2. NEAR-DETERMINISTIC SCHEMES

Entanglement generation schemes should preferably minimise the average generation time for a given fidelity F. However, the generation and distribution of entanglement is typically a *stochastic* process, greatly complicating the optimisation over such schemes. Here, we simplify the problem by demanding that every step of the entanglement generation scheme is *near-deterministic*. This requirement can be enforced even when some of the processes are not deterministic, such as elementary pair generation or Bell swaps. The probability of having at least a single success can be increased by repeating the whole scheme up until that point for multiple attempts [2]. Near-deterministic schemes deliver a state with high probability at a specific time T, and it is this generation time T that we use as our metric in this work [3].

Let us exemplify this idea through a process for elementary pair generation (EPG). This process might have a very small probability p to succeed in a single attempt, which takes a time T_{attempt} to perform. The probability of having at least a single success after r attempts, is

$$p_{\text{single success}} = 1 - (1-p)^r \,. \tag{4.1}$$

Thus, the probability of having at least one success can be increased to no less than $p_{\min}$ by trying for $r = \left\lceil \frac{\log(1-p_{\min})}{\log(1-p)} \right\rceil$ attempts. We now consider protocols where the state is stored until a total time $r \cdot T_{\text{attempt}}$ has passed, even if a success occurs before r attempts have passed. This ensures that a state can be delivered near-deterministically (i.e. with probability at least $p_{\min}$) at a prespecified time $T = r \cdot T_{\text{attempt}}$. However, it comes at the cost of increased decoherence, since the state might have to be stored for a longer time (see [147] for a related concept).

Consider now the success probability of distillation protocols and (optical) Bell state measurements. Both protocols require the two states to be present, which holds with probability equal to the product of the probabilities of the two individual schemes having succeeded for near-deterministic schemes. Furthermore, distilling and swapping typically have a non-zero failure

[1] As discussed later, we actually do consider distillation protocols taking three or more states to a single one, but these are composed of several distillation protocols taking two states to a single one.

[2] Or, in the case of MP platforms, the probability of having a single success for elementary link generation can be increased by having more modes.

[3] Since the success of a scheme follows a geometric distribution, the average generation time can be computed from the success probability and the generation time of one attempt.

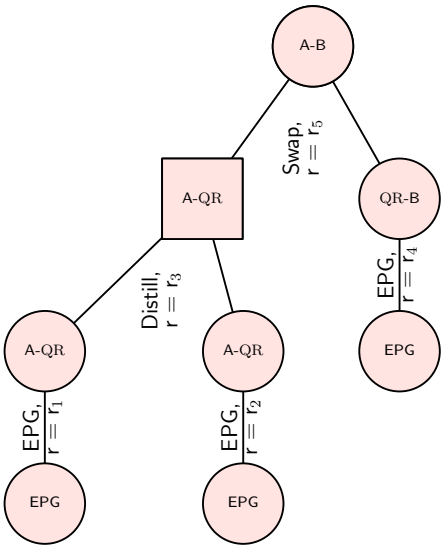

Figure 4.4: Schematic description of how near-deterministic schemes are constructed from the protocols shown in Figs. 4.1, 4.2 and 4.3. Here entanglement is generated between the nodes A and B, using an intermediate node labelled by QR. The overall structure is that of a binary tree (modulo the leaves indicating elementary pair generation, indicated by EPG), since swapping and distillation is always performed between exactly two schemes. Each sub-tree is required to succeed with probability at least $p_{\min}$, which can be enforced by repeating the whole sub-tree for a number of attempts r. Here, the specific number of attempts is indicated by r_b, $b \in \{1,2,3,4,5\}$. The circular nodes indicate either elementary pair generation or swapping, and the square nodes indicate distillation.

probability, potentially decreasing the success probability even further. However, we can use the same strategy used previously to increase the total success probability. That is, by repeating the whole scheme up to that point, it is possible to increase the success probability to at least the threshold $p_{\min}$. Let us consider this concept for the example of a swap operation between two elementary pairs. The total success probability can now be increased by repeating the whole process of generating both elementary pairs and performing the swap operation.

This concept can be extended to more complex repeater schemes, ensuring that each step in the repeater scheme succeeds with high probability. A repeater scheme can thus be constructed by combining protocols from the ground up, where the average state, generation time T, and success probability p of each scheme are only a function of the number of attempted rounds r, the protocol used, the parameters of the repeater chain, and the used schemes. We show an example of how such schemes can be constructed in Fig. 4.4.

We note here that such near-deterministic schemes require us to keep states stored for some time, even if the underlying process has already succeeded, similar to the approaches in [72, 147]. This evidently comes at the cost of increased storage times, and thus a greater amount of average decoherence. Near-deterministic schemes also have benefits, however. Firstly, with near-deterministic schemes it is possible to make the variance of the resultant probability distributions arbitrary small by increasing $p_{\min}$. Thus, near-deterministic protocols are able to deliver entanglement at a pre-specified time with high probability, which may be important for quantum information protocols consisting of multiple steps [72], such as entanglement routing [123, 149]. Secondly, it is possible to calculate exactly the generation times and fidelities of near-deterministic schemes with relative ease (see for example Appendix 4.6.4), allowing for the optimisation over such schemes.

Let us now compare near-deterministic schemes with two similar frameworks considered in [77] and [147]. Both near-deterministic schemes and the schemes considered in [77] take as building blocks a similar set of probabilistic protocols. In [77] however, the protocols are freely combined which makes challenging to estimate the average time they take to generate entanglement. This

problem is sidestepped in [77] by heuristically assuming that all protocols take average time. In contrast, in our framework, we combine protocols in blocks that have high success probability and take a fixed amount of time. This reduces the class of schemes but allows us to estimate exactly the generation time and the fidelity of the state generated.

The framework from [147] considers combining schemes in a similar fashion as done in this work, but from a more analytical perspective. That is, at each 'nesting level' there is a maximum number of time in which attempts can be made, after which a state is only made available at exactly this time. The optimisation performed there is only over a limited set of schemes. Namely, distillation is not considered, and swapping only occurs between schemes that have been generated in the same way. In particular, the optimisation from [147] can only be performed for repeater chains where the number of elementary links is given by a power of two, and all nodes/fibres share the same experimental parameters. Furthermore, the only noise present was assumed to be dephasing noise. On the other hand, the optimisation in [147] was not restricted to near-deterministic schemes.

4.2.3. BRUTE-FORCE ALGORITHM

We now introduce a brute-force algorithm to optimise entanglement distribution over the set of near-deterministic schemes between two distant nodes Alice and Bob. The algorithm takes as input the experimental parameters of the nodes and connecting fibres, the distances between adjacent nodes, a set of protocols for elementary pair generation, swapping and distillation, a minimum success probability and a limit on the maximum number of attempts and the maximum number of distillation rounds. The output consists of a data structure containing the schemes that minimise generation time parametrised by success probability and fidelity.

The brute-force algorithm generates and stores every possible scheme that can be created from the input conditions. Then for each achieved fidelity, it walks over the stored schemes to find the scheme minimising the generation time achieving at least that fidelity. In the following we sketch only the first part, as this is enough to argue that such an approach is non-scalable.

First, the algorithm takes the set $\mathscr{E}$ of protocols for elementary pair generation, together with the different number of attempts considered (of which there are at most r_{discr}), and explores all possible combinations of elementary pair generation protocols and number of attempts for each elementary link. Each of these combinations is stored if the success probability is larger than a specified $p_{\min}$.

Next, the algorithm takes the set of distillation protocols $\mathscr{D}$ and a maximum number of distillation rounds m. For each elementary link, the algorithm loops over the number of distillation rounds: $1,\ldots,m$. For each number of rounds, the algorithm explores all combinations of pairs of schemes, number of attempts and distillation protocols and stores the resulting scheme if the success probability is larger than $p_{\min}$.

The algorithm then proceeds iteratively over multi-hop links of length $i \in \{2,3,\ldots,n\}$, where n is the total number of elementary links between the target nodes. Each iteration is divided into a swapping and a distillation step.

In the swapping step the algorithm considers all adjacent (multi-hop) links of lengths i_1, i_2 such that $i_1 + i_2 = i$. For each valid pair of adjacent links and for each pair of schemes stored over the adjacent links, the algorithm explores all combinations of number of attempts and protocols in the set of swapping protocols $\mathscr{S}$. It stores a resulting scheme if the success probability is larger than $p_{\min}$.

In the distillation step, the algorithm proceeds analogously to the description above for distillation over elementary links. The output of the brute-force algorithm is then a collection of schemes. Each of these schemes is built up from smaller schemes, similar to the scheme shown in Fig. 4.4.

While the approach just described might work for a very small chain, the number of schemes grows too quickly. In particular, the number of schemes to consider in the brute-force approach is lower bounded by

$$\mathcal{O}\left(\left(\left(r_{\text{discr}}\right)^2 \cdot |\mathscr{E}| \cdot |\mathscr{S}|\right)^n\right) \tag{4.2}$$

when distillation protocols are not considered and by

$$\mathcal{O}\left(\left(r_{\text{discr}} \cdot |\mathscr{E}| \cdot |\mathscr{S}| \cdot |\mathscr{D}|\right)^{2^{m \cdot n}}\right) \tag{4.3}$$

when distillation is considered. Here n is the number of elementary links, $|\mathscr{E}|$ is the number of ways elementary pairs can be generated (due to for example varying a parameter over some set of values), $|\mathscr{S}|$ the number of swapping protocols, $|\mathscr{D}|$ the number of distillation protocols, r_{discr} the different number of attempts considered, and m the number of distillation rounds (see Appendix 4.6.1).

4

4.2.4. A HEURISTIC ALGORITHM

Now we introduce an efficient heuristic optimisation algorithm. The heuristic algorithm takes as starting point the brute-force algorithm presented before and incorporates a number of modifications that reduce the search space, thus overcoming the fast-growing complexity of the brute-force algorithm. We divide the modifications into heuristics for the pruning of schemes and heuristics for good schemes and detail them in the following. In the following we first discuss the modifications to the brute-force algorithm before presenting the pseudocode of the algorithm and analysing its complexity.

HEURISTICS FOR THE PRUNING OF SCHEMES

The brute-force algorithm explores a grid of parameters at each step and stores all schemes with success probability above p_{min} independently of their quality. Instead, we can identify schemes that either are unlikely to combine into good schemes at subsequent steps or are very similar to existing schemes and not store them.

A first strategy is to only store schemes that deliver a state with fidelity above the threshold $F_{\text{threshold}} \geq \frac{1}{2}$.

A second strategy is to *coarse-grain* the fidelity and success probabilities. For this, the algorithm rounds the fidelity F and success probability p of each scheme to $\tilde{F}$ and $\tilde{p}$, the closest values in the sets $\left[F_{\text{threshold}},\ F_{\text{threshold}} + \varepsilon_F,\ F_{\text{threshold}} + 2\varepsilon_F,\ \ldots,\ 1\right]$ and $[p_{\text{min}},\ p_{\text{min}} + \varepsilon_p,\ p_{\text{min}} + 2\varepsilon_p,\ \ldots,\ p_{\text{max}}]$ (see Appendix 4.6.3).

If no scheme with the same $\tilde{F}$ and $\tilde{p}$ exists, the scheme is stored. Otherwise, we compare the two generation times of the two schemes. If the old scheme has a lower generation time, the new scheme is not stored. Otherwise, the new scheme replaces the old one. We note here that the actual values of F and p are stored, and not the values $\tilde{F}$ and $\tilde{p}$.

The third strategy consists in pruning sub-optimal protocols after having considered all protocols over a given (multi-hop) link. A scheme is sub-optimal if there exists another scheme over that same (multi-hop) link with the same $\tilde{p}$ and has a lower generation time but equal or higher fidelity. We detail the implementation of the above pruning heuristics in Algorithm 3.

HEURISTICS FOR GOOD SCHEMES

Pruning reduces the amount of sub-optimal schemes that are kept stored. This prevents those schemes from being combined with other schemes, reducing the algorithm runtime. However, it would be preferable if those schemes would not even be considered in the first place. For this

reason, we use *heuristics* on what schemes to consider. The heuristics that we use are *banded distillation*, *banded swapping*, and the *bisection heuristic*, which we will detail in what follows.

Many distillation protocols acting on two states yield states of fidelity larger than the input states only when the input states have fidelities that are relatively close to each other [47]. This motivates restricting distillation to states that have fidelities F_1 and F_2 separated at most by some threshold $\varepsilon_{\text{distill}}$,

$$|F_1 - F_2| \leq \varepsilon_{\text{distill}} \,. \tag{4.4}$$

This heuristic, first considered in [166] is called *banded distillation*.

Inspired by banded distillation we introduce a similar heuristic for entanglement swapping that we dub *banded swapping*. A naive extension of banded distillation to swapping would be to require that the absolute difference of the fidelities of the two swapped states be small. However, by investigating the heuristically optimised schemes, our numerical exploration (see Appendix 4.6.3) suggests that the number of nodes over which the entanglement is generated also plays a role. In particular, we find that it is sufficient to restrict swapping to states that satisfy,

$$|i_1 - i_2| \leq 2\log(i_1 + i_2 - 1), \tag{4.5}$$

and

$$\left|\frac{\log(F_1)}{i_1} - \frac{\log(F_2)}{i_2}\right| \leq \varepsilon_{\text{swap}} \tag{4.6}$$

where $\varepsilon_{\text{swap}}$ controls the granularity of the heuristic, F_1, F_2 are the fidelities of the two states, and i_1, i_2 is the number of elementary links over which the entanglement was generated, e.g. the number of elementary links between QR_i-QR_j and QR_j-QR_k in Fig. 4.2, respectively. We note that the first condition was already present in [77].

The third heuristic - which we call the *bisection heuristic* - is inspired by the BDCZ scheme [20]. Similarly to the BDCZ scheme, it applies to *symmetric repeater chains*. That is, repeater chains where all nodes have the same parameters and are connected by identical elementary links. However, unlike the BDCZ scheme which is only applicable if the number of elementary links is equal to a power of two, the bisection heuristic is applicable independent of the number of elementary links.

The heuristic works as follows. Factorisation allows us to write the total number of elementary links as $n = 2^j \cdot h$, where j is the number of times n is divisible by 2, and h is the odd part of n. First, an optimisation is performed over a link of length h. From then on, similar to the BDCZ scheme, swapping only occurs between entanglement that has been generated over a total number of elementary links equal to a multiple of h. This heuristic has the possibility of dramatically reducing the algorithm runtime for certain values of n. However, for odd n this heuristic provides no speedup.

PSEUDOCODE OF THE HEURISTIC ALGORITHM

We now present the pseudocode of the heuristic algorithm. The general algorithm is described in Algorithm 4, while the subroutines for storing the schemes and for the pruning heuristic are given in Algorithm 2 and Algorithm 3.

The algorithm takes as input an additional number of parameters on top of the parameters already discussed for the brute-force algorithm. These parameters regard the heuristics and were described in the previous section. These parameters are ε_F, ε_p (the discretisation used for the pruning of schemes for the fidelity and success probability, respectively), $F_{\text{threshold}}$ and p_{max} (the minimum values required to consider a scheme for the fidelity and success probability, respectively). A software implementation requires also a number of experimental parameters for characterising the hardware and estimating the output of each scheme, however we leave the explicit

description of the hardware parameters out of the pseudocode. For details of the actual implementation, please refer to the repository [58].

Algorithm 2: STORESCHEME, subroutine for storage of the schemes. Here, 'link' refers to either an elementary or multi-hop link.

Input: scheme, store, p_{min}, $F_{\text{threshold}}$, link, ε_F, ε_p, T
Output: store with scheme possibly added

$F \leftarrow$ fidelity stored in scheme ;
$p \leftarrow$ probability stored in scheme ;
$n_{\varepsilon_p} \leftarrow \text{argmin}_{n\in\mathbb{N}}$ s.t. $p < p_{\text{min}} + n\cdot\varepsilon_p$;
$n_{\varepsilon_F} \leftarrow \text{argmin}_{n\in\mathbb{N}}$ s.t. $F < F_{\text{threshold}} + n\cdot\varepsilon_F$;
if $F \geq F_{\text{threshold}}$ **then**
  if store[link][n_{ε_p}][n_{ε_F}] *already exists* **then**
    $T' \leftarrow$ generation time of store[link][n_{ε_p}][n_{ε_F}];
    if T<T' **then**
      store[link][n_{ε_p}][n_{ε_F}] $\leftarrow$ scheme
    end if
  else
    store[link][n_{ε_p}][n_{ε_F}] $\leftarrow$ scheme
  end if
end if
return store

COMPLEXITY AND RUNTIME OF THE HEURISTIC ALGORITHM

As we show in Appendix 4.6.1, the heuristics allow us to go from a number of considered schemes that grows super-exponentially in the number of elementary links, to a number of schemes that is upper bounded by

$$\mathcal{O}\left(2\cdot r_{\text{discr}}\left(\frac{(1-F_{\text{threshold}})(1-p_{\text{min}})}{\varepsilon_F\varepsilon_p}\right)^2 n^2\log(n)\right), \tag{4.7}$$

implying that the number of considered schemes is now only on the order of $n^2\log(n)$, as opposed to super-exponential in n. Here r_{discr} is the maximum number of values allowed for the number of attempts r, $F_{\text{threshold}}$ the minimum fidelity we allow a scheme to have, p_{min} the minimum accepted success probability, ε_F, ε_p, are the discretisation used for the coarse-graining and n the number of elementary links. Furthermore, in the case of a symmetric repeater chain (i.e. every node has the same parameters and the nodes are equidistant), the optimisation can be further simplified. As we show in Appendix 4.6.1, the number of schemes to consider in the symmetric case is upper bounded by

$$\mathcal{O}\left(r_{\text{discr}}\left(\frac{(1-F_{\text{threshold}})(1-p_{\text{min}})}{\varepsilon_F\varepsilon_p}\right)^2 n\log(n)\right). \tag{4.8}$$

Algorithm 3: PRUNE, prunes the sub-optimal schemes stored for a given link. Here, 'link' refers to either an elementary or multi-hop link.

Input: store, $p_{\min}$, link, ε_p
Output: store with sub-optimal schemes over link pruned

```
for n ≥ 0 s.t. p_min + n·ε_p ≤ 1 do
    orderedSchemes ← store[link][n], ordered by fidelity from high to low;
    N ← size of orderedSchemes;
    maxTime ← generation time of orderedSchemes[0];
    for i ← 1,...,N do
        if maxTime ≤ generation time of orderedSchemes[i] then
            Remove orderedSchemes[i] from store[link][n]
        else
            maxTime ← generation time of orderedSchemes[i]
        end if
    end for
end for
return store
```

In practice, we find that our algorithm runtime ranges from approximately 100 seconds to approximately 100 minutes, when considering 1 and 35 intermediate nodes for a symmetric repeater chain, respectively. We investigate the effects of the heuristics on the algorithm runtime in more detail in Appendix 4.6.3, where we perform an experimental analysis of the algorithm runtime and its 'accuracy' when varying ε_F, ε_p, $\varepsilon_{\text{swap}}$, and $\varepsilon_{\text{distill}}$. We use these results to settle on the values for ε_F, ε_p, $\varepsilon_{\text{swap}}$, and $\varepsilon_{\text{distill}}$. We only investigate the bisection heuristic when going to a larger number of nodes in Section 4.4.3.

Algorithm 4: Heuristic optimisation over near-deterministic schemes for a repeater chain of n elementary links. Here, 'link' refers to either an elementary or multi-hop link.

Input: n : number of elementary links n in repeater chain
$\varepsilon_F, \varepsilon_p$: coarse-graining parameters for the fidelity and probability
$\varepsilon_{\text{distill}}, \varepsilon_{\text{swap}}$: parameters for the heuristics for distillation and swapping
m : maximum number of distillation rounds
$p_{\min}, p_{\max}$: minimum and maximum scheme success probabilities
r_{discr} : number of different attempt values
$F_{\text{threshold}}$: minimum fidelity for schemes to be stored
$\mathscr{E}, \mathscr{S}, \mathscr{D}$: sets of protocols for elementary pair generation, swapping and distillation
$\mathscr{L}_i, i \in [1, n]$: set of links of length i

Output: store: a data structure containing entanglement generation schemes with the minimum generation time parametrised by the link, coarse grained success probability and coarse grained fidelity.

```
Initialise store
for i ← 1 to n do
    for link in L_i do
        if i = 1 then
            // Loop over elementary pair generation protocols
            for EPGProtocol in E do
                for r such that (4.1) or (4.24) (IP/MP platforms, resp.) is between p_min and p_max in r_discr steps do
                    scheme ← EPGProtocol(r, link, n)
                    store ← StoreScheme(scheme, store, p_min, F_threshold, link, ε_F, ε_p)
                end for
            end for
        else
            // Loop over all schemes satisfying the swapping heuristic and over all swapping protocols
            for every link1 and link2 such that entanglement can be created over link by swapping between those links do
                for every pair (s_1, s_2) of stored schemes in store[link1] and store[link2] satisfying (4.5) and (4.6) do
                    for swapProtocol in S do
                        for r such that (4.1) or (4.24) (IP/MP platforms, resp.) is between p_min and p_max in r_discr steps do
                            scheme ← swapProtocol(s_1, s_2, r, link, n)
                            store ← StoreScheme(scheme, store, p_min, F_threshold, link, ε_F, ε_p)
                        end for
                    end for
                end for
            end for
        end if
        for j ← 1 to m do
            // Loop over all schemes satisfying the distillation heuristic and over all distillation protocols
            for every pair (s_1, s_2) of stored schemes in store[link] and satisfying (4.4) do
                for distillationProtocol in D do
                    for r such that (4.1) or (4.24) (IP/MP platforms, resp.) is between p_min and p_max in r_discr steps do
                        scheme ← distillationProtocol(s_1, s_2, r, link, n)
                        store ← StoreScheme(scheme, store, p_min, F_threshold, link, ε_F, ε_p)
                    end for
                end for
            end for
        end for
        store ← Prune(store, p_min, F_threshold, link, ε_p)
    end for
end for
return store
```

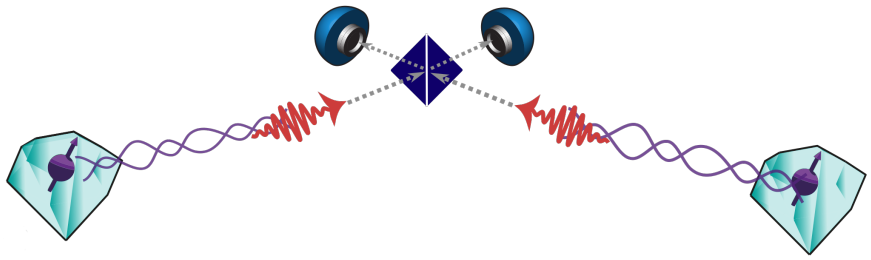

Figure 4.5: An example of an elementary link implemented with an information processing platform. The two nodes are connected by a fibre with a beamsplitter in the middle and two detectors. For the case considered in this figure, the two nodes are nitrogen-vacancy centres in diamond. For both protocols, the two nodes both send one-half of an entangled state to the middle, which after interference and successful detection leads to a shared state between the two nodes. Figure taken with permission from [143].

4.3. PLATFORM MODELS

The algorithm discussed is independent of the underlying physical implementation, and can thus be applied to several experimental platforms. We use our algorithm to study three different types of platforms encapsulating a large range of technologies. The three platforms share the capability to store quantum information but differ in their quantum information processing capabilities. We call these platforms: information processing platforms, multiplexed elementary pair generation platforms, and combined platforms. Information processing platforms have the ability to perform operations on the stored qubits, but are currently limited to a small number of qubits. Multiplexed elementary pair generation platforms, on the other hand, lack the ability to perform operations on stored states, but can generate and store a potentially very large number of different states simultaneously. Obviously, these platforms differ greatly, but both approaches have complementary qualities for long-distance entanglement generation. This motivates us to also compare a combination of the two. That is, a setup where the elementary pairs are generated with an MP platform, but swapping and distillation are performed by an IP platform.

In the rest of the section, we discuss the basics of each of the implementations and the modelling of the underlying processes.

4.3.1. QUANTUM REPEATERS BASED ON INFORMATION PROCESSING PLATFORMS

We call information processing (IP) platforms those that have the capability to perform gates on the stored states, thus enabling entanglement distillation. The number of quantum states that can be stored and processed is presently limited. Experimental information processing platforms that have demonstrated excellent control over storage qubits include NV centres in diamond [1, 35, 65, 66, 72, 83, 163], neutral atoms [151, 174], non-NV color centres in diamond [116, 117], quantum dots [22, 27, 70], and trapped ions [15, 73, 109].

In this work we consider two protocols for the generation of elementary pairs for information processing platforms. These protocols are the single-click- [23, 98, 143] and double-click protocol [7]. We give an example based on nitrogen-vacancy centres in diamond in Fig. 4.5. We stress that this is just one example of an information processing platform, and that our algorithm can be applied to other platforms.

The setup for both the single-click and the double-click protocols consists of two nodes with at least one memory qubit. The two nodes are connected via an optical channel to an intermediate

beamsplitter station with a detector at each of the output ports (see Fig. 4.5).

Let us now detail first the single-click protocol. The qubits at the nodes are prepared in a superposition of the ground state ($|\downarrow\rangle$) and the first excited state ($|\uparrow\rangle$): $\sin(\theta)|\downarrow\rangle + \cos(\theta)|\uparrow\rangle$. Upon receiving an appropriate excitation signal, the memory emits a photon ($|1\rangle$) if it is in the excited state, and no photon ($|0\rangle$) otherwise. Since the memory qubit is in a superposition, this results in a memory-photon entangled state $\sin(\theta)|\downarrow\rangle|0\rangle + \cos(\theta)|\uparrow\rangle|1\rangle$. The two photons are then directed to and interfered on the intermediate beamsplitter. One experimental complication here is that the phase picked up by the photons as they travel through the fibre is unknown unless the fibres are stabilised. However, if this is the case, upon the detection of a single photon (single-click) at the beamsplitter station, the creation of an entangled pair can be heralded to the two nodes.

The double-click protocol on the other hand does not rely on phase-stabilisation. For the double-click protocol, each node prepares a qubit in a uniform superposition of the ground and first excited state [7]. By applying specific pulses to the qubits, a photon will be coherently emitted in the early or late time-bin, depending on the state of the qubit at the node. The photons are then interfered at the beamsplitter station. The entanglement between the two qubits is heralded to the two nodes upon the detection of two consecutive clicks at the beamsplitter station. While the double-click protocol does not require phase-stabilisation, it has a lower success rate in comparison to the single-click protocol for experimentally relevant distances.

The parameter θ is tuneable, which allows for a trade-off between the success probability and the fidelity of the heralded state for the single-click protocol [24, 83, 143]. For the double-click protocol there is no such trade-off however.

For the single-click protocol we use the error model from [143]. For the double-click protocol we use the error model from [7].

Entanglement distillation across two separated matter qubits has been achieved with an NV-centre setup [83], where a specific entanglement distillation protocol [24] was implemented. This distillation protocol is optimal when the involved states are correlated in a particular manner [141]. In general however, the states that we consider are not of this form. For this reason, we will consider here only the DEJMPS protocol [41], which was originally designed to work well for maximally entangled states with depolarising noise. In this protocol, we first apply a local rotation on each of the qubits, then two local CNOT operations, and measure the targets of the CNOT operations in the computational basis. We deem the distillation to be a success when the measurement outcomes are equal.

We now sketch the underlying abstract error models and the various experimental parameters.

State preparation for the generation of elementary pairs takes some time t_{prep}, performing the gates for distillation takes time t_{distill}, and performing a Bell state measurement takes time t_{swap}. State preparation is also imperfect, which we model as dephasing with parameter F_{prep}. States stored in the memories for a time t are subject to decoherence. We model this decoherence as joint depolarising and dephasing noise, see Appendix 4.6.4 for details on the decoherence model.

The fibre has a refractive index of n_{ri} and an attenuation length L_0. The attenuation length is defined such that $\eta = e^{-L/L_0}$, where η is the transmissivity and L the length of the fibre. There are three other sources of photon loss that we model, similar to Chapters 2 and 3. The probability of successfully emitting a photon p_{em}, the probability of emitting a photon with the correct frequency and it not being filtered out (conditioned on having emitted the photon) p_{pps} and the probability of the detector successfully clicking when a photon is incident p_{det}.

Applying gates induces noise on the states. Performing a Bell state measurement induces depolarising and dephasing with parameters $\lambda_{\text{BSM, depol}}$ and $\lambda_{\text{BSM, deph}}$, respectively. Performing the CNOT operations for distillation also leads to depolarising and dephasing with parameters $\lambda_{\text{CNOT, depol}}$ and $\lambda_{\text{CNOT, deph}}$, respectively. Furthermore, we model measurement errors by applying depolarising noise with parameter $\lambda_{\text{meas. depol.}}$ before measuring a state. Finally, the uncer-

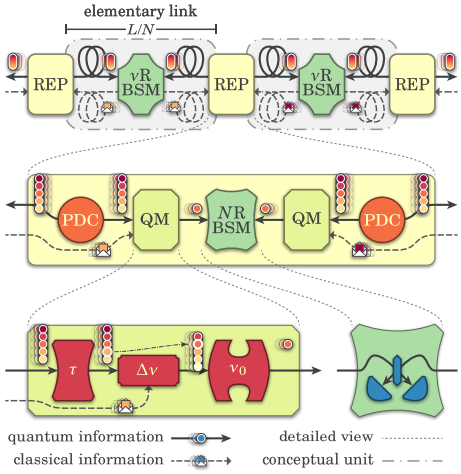

Figure 4.6: Schematic description of an MP implementation. Top: The total distance L is split into N elementary links, each with a spectrally-resolving BSM (indicated by νRBSM) in the middle, and with two nodes (each indicated by REP) at the end point of the elementary links. Middle: Zoom in of a node. Each node contains two PDC sources of multiplexed bipartite entanglement, two quantum memories (indicated by QM) and a number-resolving Bell state measurement station (indicated by NRBSM). Bottom: Detailed view of QM and NRBSM. Each quantum memory not only stores (in the unit indicated by τ), but can also perform a frequency-shift (in a unit indicated by $\Delta\nu$) and a frequency filter (indicated by the unit ν_0), while each NRBSM contains a beamsplitter and two single-photon detectors, which performs a Bell state measurement on the frequency-shifted photons. Illustration taken with permission from [90].

tainty in the phase stabilisation $\Delta\phi$ induces dephasing in the state preparation for the single-click protocol (see [143]).

4.3.2. QUANTUM REPEATERS BASED ON MULTIPLEXED ELEMENTARY PAIR GENERATION PLATFORMS

Multiplexed elementary pair generation (MP) platforms are a promising candidate for quantum repeater implementations [31, 90, 104, 146]. Such implementations generate elementary pairs with a potentially large number of modes at the same time. While multiplexed elementary pair generation platforms lack the ability to perform gates on the states stored in the memories, they have the potential to process a large number of states simultaneously, which can dramatically increase the probability at which elementary pairs can be generated. Here we discuss the basics of a model for the quantum repeater scheme proposed in [90] (see Appendix 4.6.5). This repeater scheme uses photon-number and spectrally resolving detectors, frequency-multiplexed multimode memories, and parametric down conversion (PDC) sources.

An elementary link consists of two PDC sources, each located at one of the two nodes. The PDC sources emit entangled states for a large set of frequencies. One half of each entangled state is sent towards a jointly collocated quantum memory, which can store a large number of modes simultaneously. The other half is sent to an intermediate station between the two nodes, where it interferes on a spectrally-resolving beamsplitter with the corresponding state sent from an adjacent node. If at least one successful click pattern is detected at the output of the beamsplitter, the information of the corresponding mode is sent to the nodes. The information is used to filter out the other modes, after which frequency conversion is performed to a predetermined frequency at each of the nodes. The frequency conversion to a predetermined frequency ensures that at each node the successful modes from the two adjacent links can interfere at a local beamsplitter station. Photon-number resolving detectors are collocated at the output of the local beamsplitter to identify and discard multiphoton events. A schematic description can be found in Fig. 4.6.

Let us now investigate the parameters underlying the scheme we have just described. Consider a PDC source emitting entangled states with time-bin encoding. An ideal source would emit states of the form $\frac{1}{\sqrt{2}}(|10,01\rangle + |01,10\rangle)$, where the notation $|nm,mn\rangle$ indicates n/m photons in the 'early/late' bin in one half of the state and m/n photons in the 'early/late' bin in the other

half. However, realistic PDC sources include additional terms. The resulting state can be approximated [90] by a state of the form

$$|\psi_{N_s}\rangle = \sqrt{p_0}\,|00,00\rangle + \sqrt{\frac{p_1}{2}}\,(|10,01\rangle + |01,10\rangle) + \sqrt{\frac{p_2}{3}}\,(|20,02\rangle - |11,11\rangle + |02,20\rangle)\,, \tag{4.9}$$

with

$$\begin{aligned} p_0 &= \frac{1}{(N_s+1)^2}, \\ p_1 &= \frac{2N_s}{(N_s+1)^3}, \\ p_2 &= 1 - p_0 - p_1\,. \end{aligned} \tag{4.10}$$

Here N_s is the mean photon number present in the state and is a tuneable parameter. Increasing the mean photon number N_s increases the probability of detecting two clicks at the middle station (as can be seen from the decrease in the parameter p_0), while at the same time lowering the fidelity of the state conditioned on detecting two clicks.

Note that (4.10) is a truncated version of the state derived in [86], i.e. all the higher order terms are included in p_2. As described in [90], the multiphoton components limit the ability to generate entanglement without the use of photon-number resolving detectors.

The number of modes N_{modes} increases the success probability of elementary pair generation. If the success probability of the creation of a single elementary pair is given by p_{el}, the success probability of generating at least one elementary pair is given by $1-(1-p_{\text{el}})^{N_{\text{modes}}}$. Thus, N_{modes} should be on the order of $\frac{1}{p_{\text{el}}}$, since $\lim_{p_{\text{el}}\to 0} 1-(1-p_{\text{el}})^{\frac{\alpha}{p_{\text{el}}}} = 1-e^{-\alpha}$. Finally, while a purely deterministic Bell state measurement is impossible using only linear optics [100, 164], there are theoretical workarounds to increase the success probability [50, 62, 91–93, 120, 178]. We consider the approach introduced in [62], where the success probability of the Bell state measurement can be increased to $1-\frac{1}{2^{N+1}}$ by using $2^{N+1}-2$ ancillary photons.

We assume the states can be retrieved from the memories on-demand. On-demand retrieval is necessary for our algorithm to work, since the storage times are not fixed. This is due to the uncertainty in which attempt entanglement will be generated. On-demand retrieval can be achieved with rare-earth ion ensembles by, for example, switching coherence from electronic levels to spin levels, as done in [68, 161]. Besides allowing for on-demand recall, this also has the added benefit of increased memory life-time [145].

We consider the same type of noise for operations as we did for information processing platforms. This means that measurements have an associated amount of depolarising and dephasing. Finally, 'decoherence' over time for the memory manifests as an exponential decay in the output efficiency of the memory, not in a reduction of the fidelity of the state [2, 145]. Thus, the longer a state is stored, the smaller the probability it can be retrieved for measuring or further processing.

4.3.3. COMBINING THE TWO SETUPS

An information processing implementation has the benefit of long coherence times and control over the memory qubits, which allows for distillation. On the other hand, multiplexed elementary pair generation platforms do not support distillation, but have the benefit of emitting and storing a large number of modes, increasing the success probability of the elementary pair generation significantly. Optimistically, one could imagine a futuristic setup which combines the strengths of

the two setups. That is, elementary pair generation is performed by a multiplexed elementary pair generation platform, after which the successfully generated pairs are frequency-converted into a frequency that can be stored in an information processing platform. The state is then stored in a memory, which can be done using, for example, a reflection-based heralded transfer [82, 115]. For simplicity, we assume that the transfer and frequency conversion do not introduce any further noise or losses.

4.4. RESULTS

In this section, we study information processing platforms, multiplexed elementary pair generation platforms and the combination thereof with the algorithm that we introduced in Section 4.3. In order to compare different optimisation results, we have chosen four sets of parameters for both platforms. With these sets, we first investigate the performance of information processing platforms for short (≈ 15-50 km), intermediate (50-200 km) and large (i.e. ≈ 200-800 km) distances. We then perform a similar investigation for multiplexed elementary pair generation platforms, after which we investigate the combination of the two. In order to get an understanding of the necessary parameters to generate remote entanglement with each platform or combination, the four sets of parameters for each platform are strictly ordered, with set 4 having the best parameters. We begin each three of the investigations with a specification of the input to our algorithm, which consists of the used elementary pair generation, swapping and distillation protocols, experimental parameters and the parameters specific to the algorithm discussed previously.

In order to investigate longer repeater chains, we consider only *symmetric repeater chains* (see Section 4.2.4) in this section unless specified otherwise.

4.4.1. SCHEME OPTIMISATION RESULTS FOR IP PLATFORMS

In the following we discuss the heuristic optimisation results for information processing platforms. Let us first briefly discuss the protocols that we include in the optimisation.

We consider two protocols for elementary pair generation: the single- and double-click protocol, see Section 4.3.1. The single-click protocol has an additional parameter θ, which modulates the weight of the zero and one photon component [143]. We optimise over all single-click protocols with θ taking values between $\frac{1}{2}$ and π, equally spaced in 300 steps, thus $|\mathscr{E}| = 301$.

Both for swapping and distillation we consider a single protocol, i.e. $|\mathscr{S}| = |\mathscr{D}| = 1$. For swapping we perform a deterministic Bell state measurement on matter qubits while for distillation we implement the DEJMPS protocol. For swapping and distillation, we optimise over all pairs of schemes that satisfy the banded swapping and distillation heuristics, see Section 4.2.4.

For all of the schemes, r ranges from r_{min} to r_{max} in (at most) $r_{\text{discr}} = 200$ steps, where r_{min} and r_{max} are chosen such that the success probabilities are at least p_{min} and p_{max}, respectively.

We set $\varepsilon_{\text{swap}} = \varepsilon_{\text{distill}} = 0.05$, $\varepsilon_F = 0.01$ and $\varepsilon_p = 0.02$. These parameters were settled on by investigating the trade-off between the accuracy of the algorithm and its runtime, see Section 4.6.3 of the Appendix for a detailed analysis. We only consider $m = 2$ distillation rounds. Finally, we set $p_{\text{min}} = 0.9$.

We now specify four sets of parameters for information processing platforms. We fix the parameters in Table 4.1 as a baseline common to all sets. We then choose sets of parameters for the efficiency coherence times, efficiencies and gate fidelities, which can be found in Table 4.2.

ENTANGLEMENT GENERATION FOR SHORT DISTANCES WITH IP PLATFORMS

Small-scale experiments relevant for entanglement distribution with information processing platforms have already been performed [14, 35, 65, 66, 72, 83, 135], demonstrating the potential of such platforms for quantum networks. It is therefore of interest to understand what is within reach for

t_{prep} (entanglement preparation time)	6 μs [65]
F_{prep} (dephasing for state preparation)	0.99 [65]
DcS (dark count rate)	10 Hz [65]
L_0 (attenuation length)	22 km [125]
n_{ri} (refractive index of the fibre)	1.44 [125]
$\Delta\phi$ (optical phase uncertainty)	14.3° [72]
$F_{\text{gates, deph}}$ (dephasing for all gates)	1

Table 4.1: Base parameters used for information processing platforms.

	Set 1	Set 2	Set 3	Set 4
T_{deph} (dephasing with time)	3 s	10 s	50 s	100 s
T_{depol} (depolarising with time)	3 s	10 s	50 s	100 s
p_{em} (probability of emission)	0.8	0.9	0.95	0.99
p_{ps} (probability of post-selection)	0.8	0.9	0.95	0.99
F_{gates} (depolarisation of all gates)	0.98	0.99	0.995	0.999

Table 4.2: Four different sets of example parameters considered for information processing platforms.

information processing platforms, and what are the relevant parameters to improve. Thus, in this section we investigate how well we can perform entanglement generation with a small number of nodes and near-term parameters over short distances with information processing platforms. In particular, we are interested in when the introduction of a node becomes useful. To this end, we first consider entanglement generation over a distance of 50 kilometres with parameter set 1. We show the results from our heuristic optimisation in Fig. 4.7, where we consider the scenarios with no node, a single node, and two intermediate nodes. Furthermore, we plot the results where we include only the single-click protocol, and both the single- and double-click protocol.

First off, the double-click protocol provides only a benefit for higher fidelities and for the scenarios with one and two intermediate nodes/three hops. This can be attributed to the fact that the double-click protocol is inherently less noisy if there are no losses, but is more sensitive to losses than the single-click protocol. However, this does not necessarily imply that all the elementary pairs have been generated with the double-click protocol. As we will see in later results, we will find schemes where elementary pairs are generated using both the single and double-click protocol, indicating the importance of considering such complex schemes in our optimisation.

Secondly, we observe that there is a cross-over point for $F \approx 0.7$ below which adding a node allows for a shorter generation time. Thus, implementing a quantum node over a modest distance of less than 50 kilometres, can in fact increase the generation rate by a moderate amount for low fidelities ($\lesssim 0.7$). However, increasing the total distance does not shift this cross-over point, since the maximum achieved fidelity with a single node also drops down if the parameters do not change.

Next, we explore the impact of a single parameter in the performance of implementations expected in the longer term. To this end, in Fig. 4.8a we investigate how the minimum generation time for several fixed target fidelities ($F = 0.7, 0.8, 0.9$) scales, when varying the gate fidelities and coherence times and using parameter set 2. More specifically, we vary the gate fidelities from 0.98 to 1 and the coherence times T_{deph} and T_{depol} from 1 to 100 seconds. We perform a similar investigation in Fig. 4.8b, where instead of varying the coherence times, we vary the success probabilities of the detector successfully clicking (p_{det}), successfully emitting a photon from a node (p_{em}), and the probability of emitting a photon of the correct frequency (p_{pps}) simultaneously from 0.8 to 1.

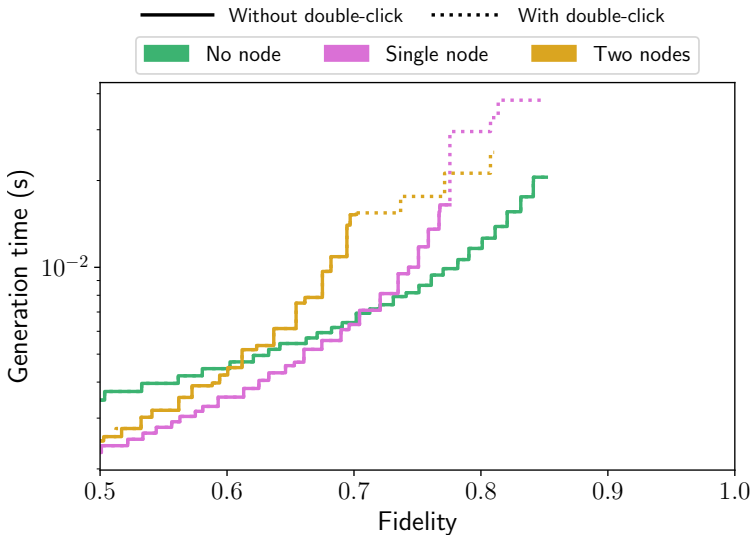

Figure 4.7: Results of the achieved fidelity and generation time for a total distance of 50 kilometre using parameter set 1 (see Table 4.2) for information processing nodes, where we consider having 0 (green), 1 (purple), or 2 (yellow) of such intermediate nodes. The solid line corresponds to a heuristic optimisation where we have excluded the double-click protocol, and the dotted line corresponds to a heuristic optimisation with both the single- and double-click protocol. The double-click protocol does not provide a benefit for direct transmission, since the double-click protocol suffers more strongly from losses than the single-click protocol.

From Fig. 4.8a we observe that increasing the gate fidelities has a bigger impact on the ability to generate entanglement than increasing the coherence times. In the bottom plot of Fig. 4.8a we choose two points, indicated by a blue ring and a yellow diamond. The schemes corresponding to those two points are visualised in Fig. 4.9. The non-monotonicity of the maximum generation rate most noticeable in Fig. 4.8a is an artefact from the heuristics occasionally leading to worse protocols, even with improved experimental parameters.

We make two observations about the algorithm from Fig. 4.9. First, the two schemes in Fig. 4.9 require swaps and distillation on states that have been created in different ways. This shows that already for only a single node entanglement distribution benefits from combining schemes in asymmetric fashion, even if the repeater chain itself is symmetric. Secondly, the algorithm is sensitive to parameter changes. We see that a small change in the parameters allows the diamond scheme to achieve a generation rate approximately four times as large as the ring scheme.

The trade-off between the success probability and the gate fidelities in Fig. 4.8b appears more complex. Not surprisingly, we observe that increasing the success probabilities has the greatest effect on the generation time and the ability to generate entangled states. In contrast to the previous scenario where only varying the gate fidelities leads to jumps in the generation time, we do not observe a similar phenomenon when varying the success probabilities. This is due to the fact that changing the success probabilities changes the generation time primarily by reducing the required number of attempts. Thus, if the minimal number of attempts $r_{\min}$ is well approximated by a continuous function $\hat{r}_{\min}(p)$ in p, we expect to see no jumps in the generation time as we vary p. More formally, we say that $r_{\min}$ approximates $\hat{r}_{\min}$ well if $\frac{r_{\min}(p)-\hat{r}_{\min}(p)}{r_{\min}(p)} \approx 0$. Since $r_{\min}(p) = \left\lceil \frac{\log(1-p_{\min})}{\log(1-p)} \right\rceil$, an obvious choice for $\hat{r}_{\min}$ is $\frac{\log(1-p_{\min})}{\log(1-p)}$. Note that we then have that

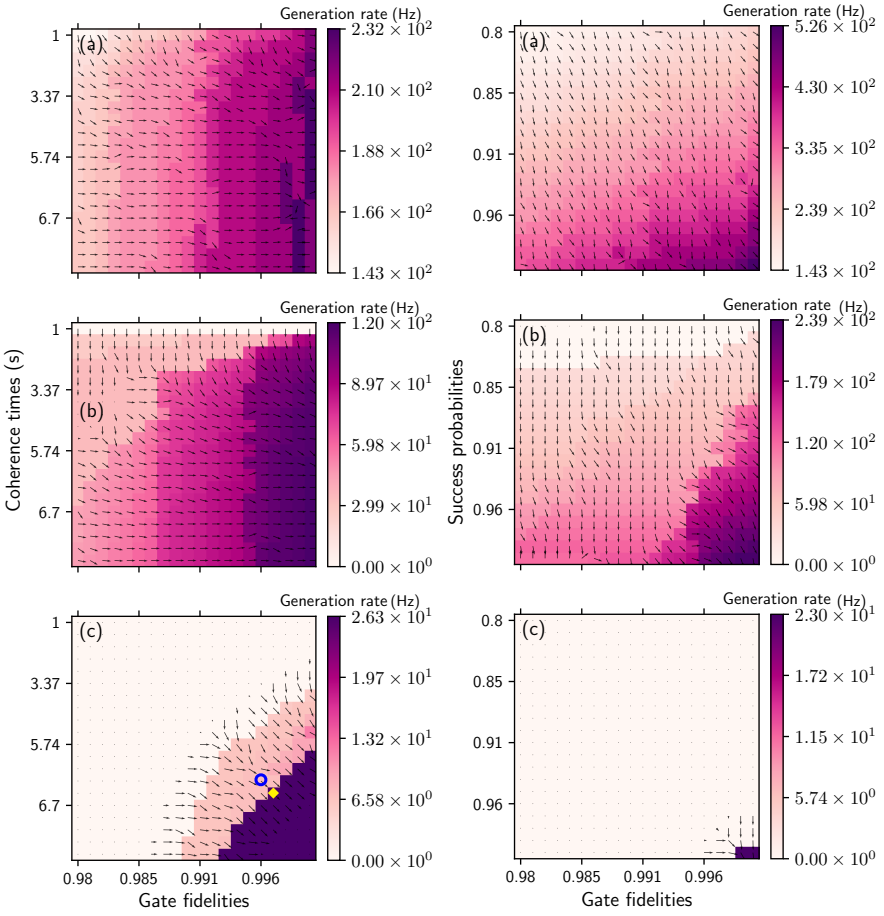

(a) Varying the coherence times (1-10s) and gate fidelities (0.98 to 1). The blue ring and yellow diamond indicate the schemes we investigate in Fig. 4.9.

(b) Varying the success probabilities (i.e. we vary $p_{\text{det}} = p_{\text{em}} = p_{\text{ps}}$ simultaneously from 0.8 to 1) and gate fidelities (0.98 to 1).

Figure 4.8: Maximum generation rates for several different values of the coherence times (1-10s) and gate fidelities (0.98 to 1) (left), and success probabilities (right) and for several different target fidelities, for a distance of 50 kilometre and a single information processing node. Down and to the right in the plots indicate better parameters. All the other parameters are fixed to those of set 2 (Table 4.2) or the base parameters (Table 4.1). The target fidelities are (a) $F = 0.7$, (b) $F = 0.8$, (c) $F = 0.9$, respectively. We also plot the gradient, indicating the direction and magnitude of steepest ascent.

$\left|r_{\min}(p) - \hat{r}_{\min}(p)\right| \leq 1$, and that for p small enough, $\left\lceil \frac{\log(1-p_{\min})}{\log(1-p)} \right\rceil \gg 1$. Since the total success probability of establishing an elementary pair is small, we have indeed that $\frac{r_{\min}(p) - \hat{r}_{\min}(p)}{r_{\min}(p)} \approx 0$, explaining the lack of sudden jumps. Furthermore, we find from Fig. 4.8b (c) that, for almost all values of success probabilities and gate fidelities, it is impossible to generate a state with a fidelity of 0.9.

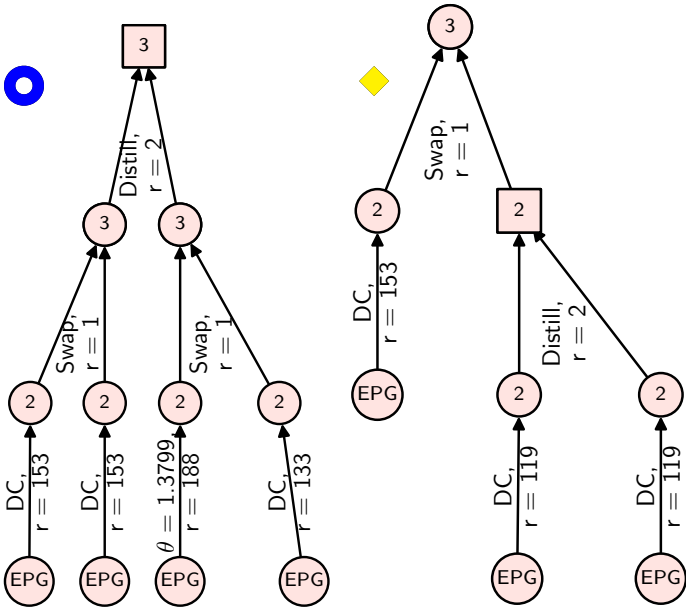

Figure 4.9: Visualisation of the two schemes indicated in the bottom of Fig. 4.8a by the blue ring (left) and the yellow diamond (right). The numbers indicate the number of nodes over which entanglement has been established, or elementary pair generation (EPG) has been performed. The 'DC' indicates the double-click protocol, and the '$\theta = \theta^*$' indicates a single-click protocol with the θ parameter set to θ^*. The 'r' here indicates the number of rounds the corresponding subtree is attempted. Note the necessity of combining disparate schemes - in both cases the EPG protocols used are not the same, and the yellow diamond scheme requires a swap on a distilled and undistilled pair.

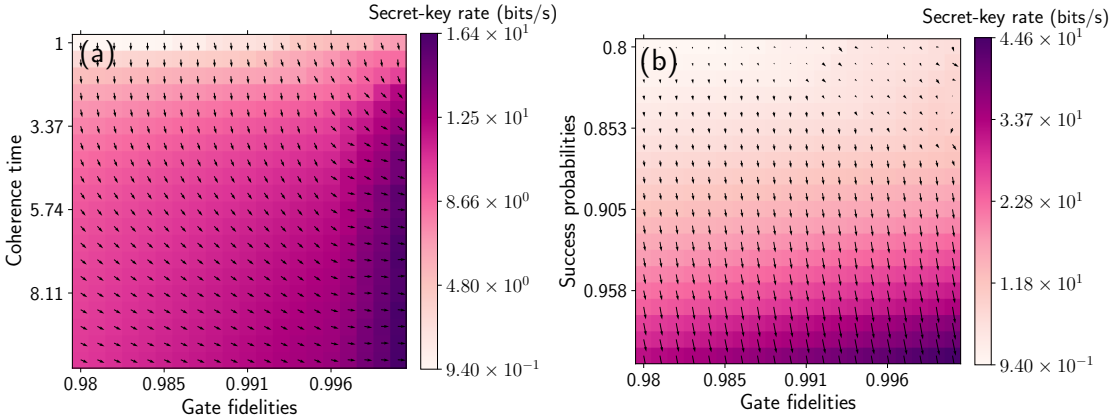

Figure 4.10: Secret-key generation using the six-state protocol, for several different values of (a) the coherence times (1-10 s) and gate fidelities (0.98 to 1), and (b) the success probabilities (i.e. we vary $p_{\text{det}} = p_{\text{em}} = p_{\text{ps}}$ simultaneously from 0.8 to 1) and gate fidelities (0.98 to 1) for a distance of 50 kilometre and a single intermediate node for information processing platforms. Down and to the right in the plots indicate better parameters. All the other parameters are fixed to those of set 2 (Table 4.2) and the base parameters (Table 4.1). We also plot the gradient, indicating the direction and magnitude of steepest ascent.

One of the near-term applications of a quantum repeater chain is the generation of secret-key. This motivates investigating the rate at which secret-key can be generated per unit time for several parameter ranges. Concretely, in Fig. 4.10 (a) and (b) we investigate the same experimental settings

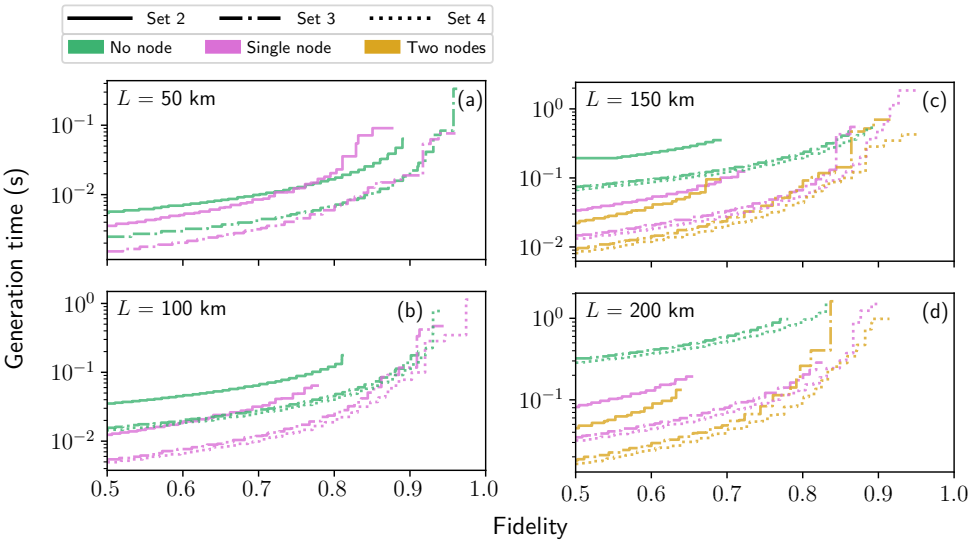

Figure 4.11: Results of the achieved fidelity and generation time for total distances of 50 (a), 100 (b), 150 (c) and 200 (c) kilometre using parameter sets 2 (solid), 3 (dashed-dotted) and 4 (dashed) (see Table 4.2) for information processing nodes, where we consider having 0 (green), 1 (purple), or 2 (yellow) of such intermediate nodes.

and parameters as in Fig. 4.8a and Fig. 4.8b. Each point corresponds to the maximum achieved secret-key per unit time generated using a six-state protocol with advantage distillation [171] for each of the schemes in the output of our algorithm.

As in Fig. 4.8a, we find in Fig. 4.10(a) that for both increasing the generation rate or secret-key rate, increasing the coherence times is most beneficial only up to a certain point, after which the gate fidelities become more important. As in Fig. 4.8b, we observe in Fig. 4.10(b) that almost always the success probabilities are more critical than the gate fidelities for increasing the secret-key rate.

INTERMEDIATE-DISTANCE ENTANGLEMENT GENERATION USING IP PLATFORMS

We expect the addition of nodes to become more beneficial as the distance over which entanglement is generated increases, conditioned on the fact that the experimental parameters are sufficiently high. In this section, we aim to quantify how good the experimental parameters need to be for this to be true. This motivates us to perform the heuristic optimisation for the entanglement generation for greater distances, and with improved parameter sets. More concretely, we investigate the achieved generation times and fidelities for intermediate distances (i.e. 50 to 200 kilometre) for the different experimental parameters proposed in Table 4.2. We start with Fig. 4.11(a), where we re-examine the scenario of Fig. 4.7 of a total distance of 50 kilometre. We now perform the heuristic optimisation with parameter sets 2 and 3, where we consider implementing either no or a single intermediate node. It is clear from Fig. 4.11(a) that introducing a node over a distance of 50 kilometre only improves the generation time by a modest amount for low fidelities, even with increased parameters. If we increase the total distance to 100 kilometre, where we now also include parameter set 4, we find in Fig. 4.11(b) that a single node proves advantageous for almost all fidelities over all three considered parameter sets. In Fig. 4.11(c) and (d) we consider greater distances of 150 and 200 kilometre, where we also include the heuristic optimisation with two intermediate nodes. We observe that while having no node is clearly inferior to having at least one,

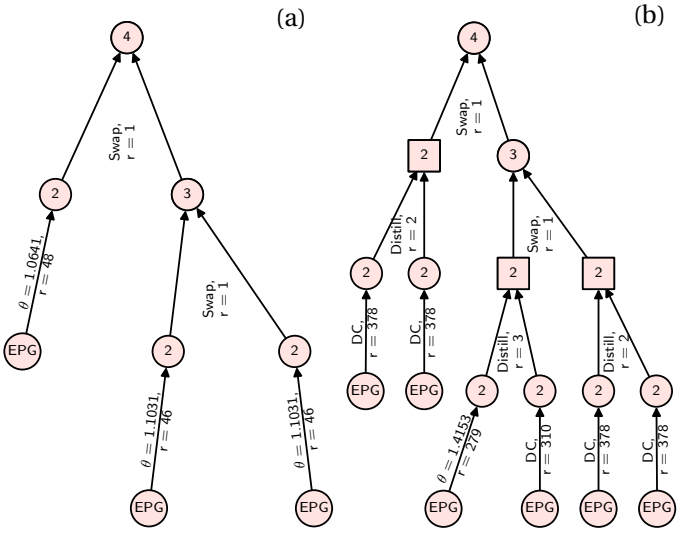

Figure 4.12: Visual representation of the schemes with the lowest non-trivial fidelity (a) and highest fidelity (b), for a distance of 200 kilometres with information processing platforms using parameter set 4 (see Table 4.2) and two intermediate nodes/three hops. The numbers in the vertices indicate the number of nodes over which entanglement has been established. The '$\theta = \theta^*$' indicates a single-click protocol with the θ parameter set to θ^*. The 'r' indicates the number of rounds the corresponding subtree is attempted. We find that the second scheme performs distillation between two elementary pairs generated with a single- and double-click protocol, demonstrating the benefit of including such distillation protocols in our optimisation.

introducing two nodes also outperforms a single node for most fidelities and sets of parameters for these distances. This suggests that the values of parameter set 3 (see Table 4.2) are a relevant objective to reach for fast near-deterministic entanglement generation with information processing platforms.

We investigate the schemes for the above scenario of 200 kilometres in Fig. 4.12, where we depict the schemes that achieve the lowest (non-trivial) fidelity and the highest fidelity. Interestingly, the scheme that achieves the highest fidelity requires that the different elementary pairs are generated both with the double- and single-click protocol. This exemplifies the need for including such asymmetric schemes in our optimisation, which appears to become more important for higher fidelities.

The numerical investigation until this point has been dedicated to symmetric repeater chains. However, realistic quantum networks will be inhomogeneous and nodes will not be equally separated. In Fig. 4.25 in Appendix 4.6.8 we show the optimisation results when considering an asymmetric repeater chain over 200 kilometres with three intermediate nodes equally separated. The asymmetry is thus only in the parameters used for the nodes, not the distance between them. The parameters used are: parameter set 4 for the three intermediate nodes, and parameter set 2 for the nodes corresponding to Alice and Bob (see Table 4.2). Such a situation can arise if the end users have access to different technology than the network operator. In this setting, we compare the results of a full optimisation with an optimisation over BDCZ schemes, a class of schemes similar to the ones proposed in [20, 48]. In particular, we include under the BDCZ class schemes that only combine identical pairs of schemes for connection and distillation. This class is different than the one in [77] as it allows optimisation over the elementary pair generation protocols but, on the other hand, it does not include distillation schemes based on pumping [77]. We find that the full optimisation gives an increased generation rate of up to a factor of 10 over BDCZ schemes.

LONG-DISTANCE ENTANGLEMENT GENERATION USING IP PLATFORMS

Generating near-deterministic entanglement over larger distances requires excellent experimental control. It is not clear how the number of nodes and the experimental parameters affect our ability to generate entanglement. To this end, we consider here the generation of high fidelity entanglement over distances of 200, 400, 600 and 800 kilometre. To gain an understanding of the relevant parameters, we study the effects of increasing gate fidelities and the memory coherence separately in Fig. 4.28 in Appendix 4.6.8. We observe in Fig. 4.28 that increasing the coherence times yields a greater benefit than increasing the gate fidelities for these distances and parameters. In particular, increasing the coherence times allows for the generation of entanglement over larger distances, while increasing the gate fidelities effectively extends the ranges of fidelity over which entanglement is generated with the same generation time. We note here that the parameters p_{em}, p_{pps} and p_{det} (corresponding to the probability of emitting a photon from the memory, emitting in the correct mode/frequency, and the probability of detecting a photon successfully, respectively) remain fixed, which inhibits the potential benefits of including more nodes.

We have found that information processing platforms, with sufficiently high parameters are a good candidate for near-term entanglement generation. In particular the success probabilities are an important factor for the generation of entanglement. However, even with multiple nodes, the maximum fidelity that can be reached is limited when attempting entanglement generation at large distances.

4.4.2. OPTIMISATION RESULTS FOR MP PLATFORMS

Having investigated the performance of information processing platforms with regards to entanglement generation, we now explore entanglement generation with multiplexed elementary pair generation platforms. Not only are we interested in how well entanglement can be generated with a repeater chain built using a multiplexed implementation, but also in how the performance differs from information processing platforms. As explained in Section 4.1, we expect that multiplexed elementary pair generation platforms perform better than information processing platforms for larger distances, provided the experimental parameters are high enough. Our aim for this section is thus to investigate for which parameters and network configurations this becomes true.

First, let us discuss the set of protocols, the algorithm parameters and the hardware parameters we will consider.

We consider one protocol for elementary pair generation, one for swapping and no protocol for distillation.

The elementary pair generation protocol (see Section 4.3.2) has one free parameter, the mean photon number N_s. Similar to information processing platforms, we also optimise over values of the mean-photon number by considering a range of values of N_s. In this case, the range is from $2\cdot10^{-4}$ to $\frac{1}{2}\left(\sqrt{5+\frac{2\sqrt{F_{\text{threshold}}(F_{\text{threshold}}+3)}}{F_{\text{threshold}}}}-3\right)$, in steps of 10^{-4}. The lowest value of $2\cdot10^{-4}$ was empirically found from the simulations to be a good conservative lower bound, while the upper bound corresponds to achieving a fidelity of the elementary pair with fidelity equal to $F_{\text{threshold}}$ when $\eta\to0$ [4], see Eq. 4.23.

The swapping protocol is a photonic Bell state measurement with fixed efficiency depending on the number of ancillary photons, see Table 4.4. Similar to the optimisation with information processing platforms, to reduce the parameter space, we implement the banded swapping heuristic, see 4.2.4.

[4] Obviously, when $\eta\to0$ the probability of getting a successful click pattern is zero. However, here we are only interested in the worst-case scenario/upper bound, which corresponds to detecting a successful click pattern as $\eta\to0$.

We use the same algorithm parameters as with the information processing platform optimisation. For all of the schemes r ranges from $r_{\min}$ to $r_{\max}$ in (at most) $r_{\text{discr}} = 200$ steps, where $r_{\min}$ and $r_{\max}$ are chosen such that the success probabilities are equal to $p_{\min}$ and $p_{\max}$, respectively. We set $\varepsilon_{\text{swap}} = \varepsilon_{\text{distill}} = 0.05$, $\varepsilon_F = 0.01$, $\varepsilon_p = 0.02$ and $p_{\min} = 0.9$. We consider only symmetric repeater chains, i.e. all the node have the same parameters and are equidistant.

Regarding the hardware parameters, the base parameters are given in Table 4.3, while the four sets of parameters are given in Table 4.4.

t_{prep} (entanglement preparation time)	6 μs
DcS (dark count rate)	10 per second
L_0 (attenuation length)	22 km
n_{ri} (refractive index of the fibre)	1.44 [125]

Table 4.3: Base parameters used for the multiplexed elementary pair generation platforms considered in this chapter.

	Set 1	Set 2	Set 3	Set 4
T_{coh} (efficiency coherence times)	10^{-2} s	10^{-1} s	10^{0} s	10^{1} s
N_{modes} (number of modes)	10^4	10^5	10^6	10^7
p (success probabilities)	0.9	0.95	0.99	0.999
p_{BSM} (BSM efficiency)	$\frac{1}{2}$	$\frac{3}{4}$	$\frac{7}{8}$	$\frac{15}{16}$

Table 4.4: The different sets of parameters considered for multiplexed elementary pair generation platforms in this chapter.

ENTANGLEMENT GENERATION FOR SHORT DISTANCES WITH MP PLATFORMS

We expect that multiplexed elementary pair generation platforms provide mostly a benefit over information processing platforms for larger distances. However, it is still of interest to investigate the performance of multiplexed elementary pair generation platforms for shorter distances. This is to gain an understanding of what can be done experimentally in the very near-term. Thus, as in Section 4.4.1.1, we first explore entanglement generation with MP platforms for short distances. We performed the heuristic optimisation with parameter set 1 for distances of 15, 25 and 50 kilometre, with 0, 1 or 2 intermediate nodes. We found that, except for a distance of 15 kilometres with no nodes, no entanglement could be generated. Even in the scenario of 15 kilometres with no nodes, the maximum fidelity that could be generated was approximately 0.56. It is thus clear that, at least with the used parameters, IP platforms are better than MP platforms for entanglement generation over short distances. We now investigate what are the relevant parameters to increase for MP platforms for entanglement generation over short distances. To this end, we perform a parameter exploration for a distance of 15 kilometres. In particular, we vary the success probabilities and the efficiency coherence times from the values of parameter set 1 to those of set 2 in Table 4.4, see Fig. 4.13.

We observe that with modest increases in the efficiency coherence times and success probabilities, entanglement generation becomes significantly more efficient. In particular, parameter set 1 (i.e. top left corner of the parameter plots) is only good enough for the generation of entanglement of very low fidelity (~ 0.56), while already a secret-key rate of ~ 500 bits per second can be achieved for parameter set 2, see Table 4.4. We conclude from the plots that, for current and near-term parameters and short distances, increasing the success probabilities is more important than increasing the efficiency coherence times.

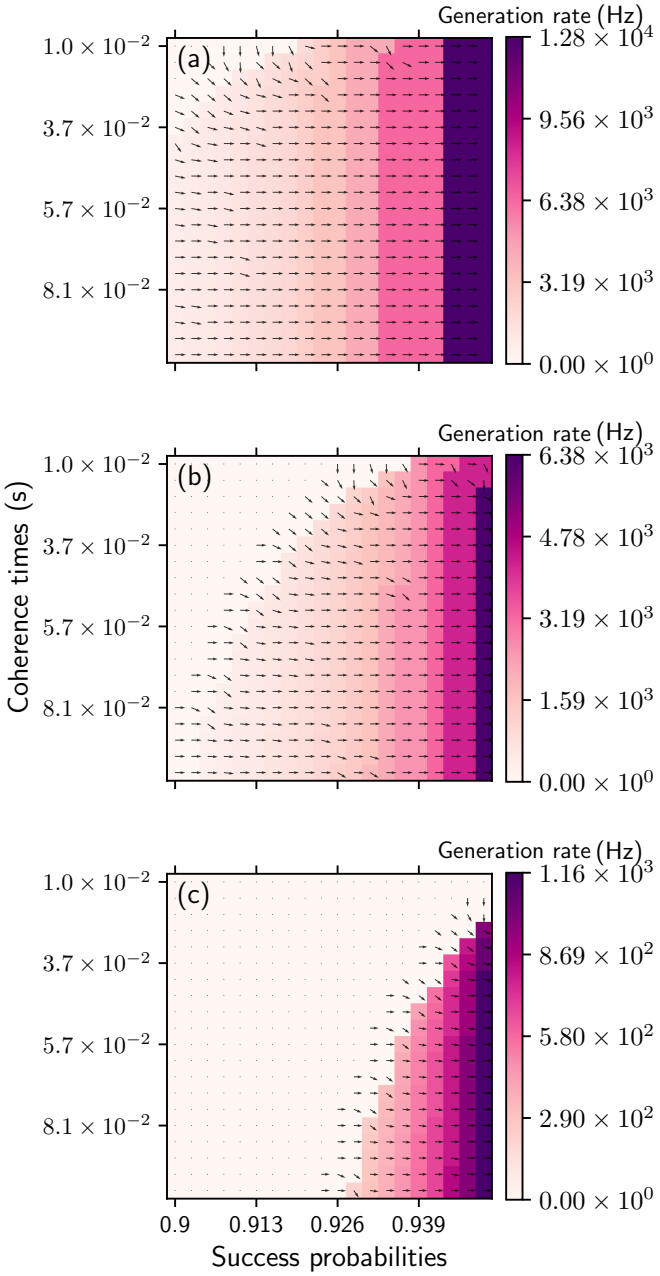

Figure 4.13: Maximum generation rates for several different values of the success probabilities (i.e. we vary $p_{\text{det}} = p_{\text{em}} = p_{\text{ps}}$ simultaneously) and efficiency coherence times, and for several different target fidelities, for a distance of 15 kilometre and a single node for MP platforms. All the other parameters are fixed to those of set 2 (Table 4.4) or the base parameters (Table 4.3). The target fidelities are (a) $F = 0.7$, (b) $F = 0.8$, (c) $F = 0.9$, respectively. We also plot the gradient, indicating the direction and magnitude of steepest ascent.

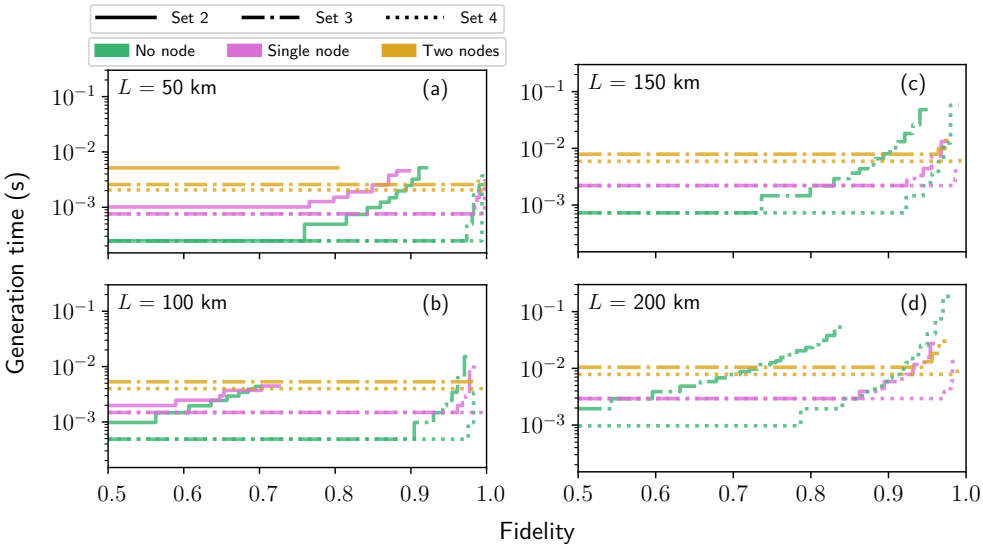

Figure 4.14: Results of the achieved fidelity and generation time for total distances of 50 (a), 100 (b), 150 (c) and 200 (c) kilometre using parameter sets 2 (solid), 3 (dashed-dotted) and 4 (dashed) (see Table 4.4) for multiplexed elementary pair generation platforms, where we consider having 0 (green), 1 (purple), or 2 (yellow) of such intermediate nodes.

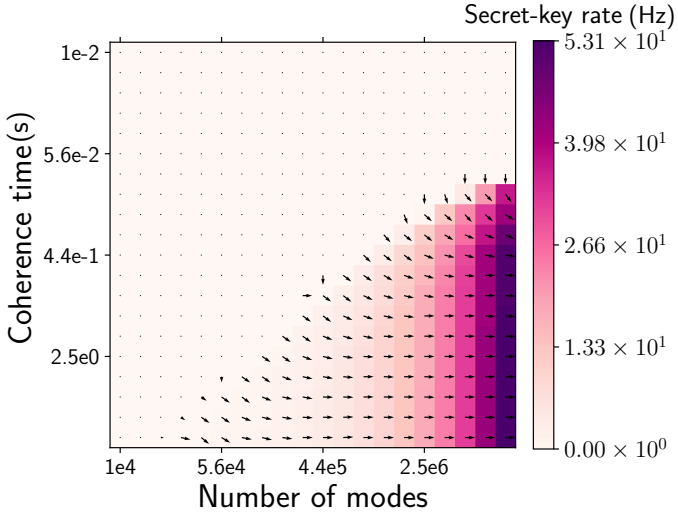

Figure 4.15: Secret-key generation using the six-state protocol, for several different values of the efficiency coherence times (10^{-2}-10^{0} s) and number of modes (10^{4}-10^{7}), for a distance of 200 kilometre and a single node for MP platforms. All the other parameters are fixed to those of set 2 (Table 4.4) and the base parameters (Table 4.3). We also plot the gradient, indicating the direction and magnitude of steepest ascent.

INTERMEDIATE-DISTANCE ENTANGLEMENT GENERATION USING MP PLATFORMS

In the previous section we have found that at short distances MP platforms do not fare as well as information processing platforms. This motivates us to investigate for which parameters and distances this does become the case. We thus investigate here entanglement distribution over distances of 50, 100, 150 and 200 kilometre, where we consider the improved parameters found in sets 2, 3, and 4 in Tables 4.3 and 4.4 in Fig. 4.14.

We find that, for most target fidelities in Fig. 4.14(a), (b) and (c), that the generation time is relatively independent of the desired fidelity. We now explain this behaviour. The fidelity is most strongly controlled by the parameter N_s - lowering N_s allows us to increase the fidelity, but lowers the success probability p of the elementary pair generation. However, the *total* success probability of generating at least one elementary pair $1-(1-p)^{N_{\text{modes}}}$ does not decrease significantly, due to the large number of modes N_{modes}. In Appendix 4.6.6 we investigate how the minimum number of modes changes, as a function of the desired fidelity of the elementary pairs. We find that the required number of modes scales at least as $\frac{\exp\left(\frac{L}{L_0}\right)}{(1-F)^2}$, where L is the distance between nodes and L_0 the attenuation length.

Since MP platforms are expected to have an advantage over information processing platforms for longer distances, we investigate the secret-key rate per unit time for a total distance of 200 kilometre (instead of 50 kilometres for information processing platforms, see Fig. 4.8a and 4.10), where we vary the number of modes and the efficiency coherence time. In Fig. 4.15 we find that for most parameters the secret-key rate per unit time is zero. As in the previous parameter explorations performed, we observe that increasing the efficiency coherence times is only (strongly) beneficial up to a certain point (which depends on the number of modes in this case), after which increasing the efficiency coherence times further does not help. Interestingly, increasing the number of modes has the greatest effect on the secret-key per unit time. Increasing the number of modes allows us to push the mean photon number to smaller numbers, effectively increasing the fidelity that can be generated within the same time-window.

LONG-DISTANCE ENTANGLEMENT GENERATION WITH MP PLATFORMS

We observe by comparing Figs. 4.11 and 4.14 that MP platforms start to outperform information processing platforms for distances of around ~200 km. Here we are interested in whether multiplexed elementary pair generation platforms still perform well for even greater distances, which is the relevant scenario for large-scale quantum networks.

Let us first focus on the effect of the efficiency coherence times and Bell state measurement efficiency on long distance entanglement generation. In Fig. 4.16 we investigate a repeater chain with 10 nodes with the parameters from set 2, the success probabilities of the Bell state measurements given by $\frac{3}{4}$, $\frac{7}{8}$ or $\frac{15}{16}$ (corresponding to a number of ancillary photons 2, 6 and 14, respectively), and the efficiency coherence time T_{coh} set to 1 or 10. We find that, even with the most optimistic parameters it is not possible to generate entanglement for distances of 800 kilometre with ten nodes.

This leads to our results shown Fig. 4.17, where we plot the heuristic optimisation results using parameter set 4, for distances of 200, 400, 600 and 800 kilometre, and the number of nodes running from one to four. We find indeed that, even for a distance of 800 kilometres, entanglement can still be generated at a high fidelity (~ 0.95). This, combined with the fact that entanglement generation for the same distance is not possible in Fig. 4.16, suggests that it is essential to also increase the number of modes and the success probabilities to generate entanglement over large distances.

We give more detail of two schemes found from the optimisation of Fig. 4.17. In particular, in Fig. 4.18 we give the schemes that achieve the lowest non-trivial fidelity and highest fidelity, respectively. As expected, the second scheme uses smaller values of the mean-photon number N_s for the elementary pair generation. This increases the fidelity of the elementary pairs, at the cost

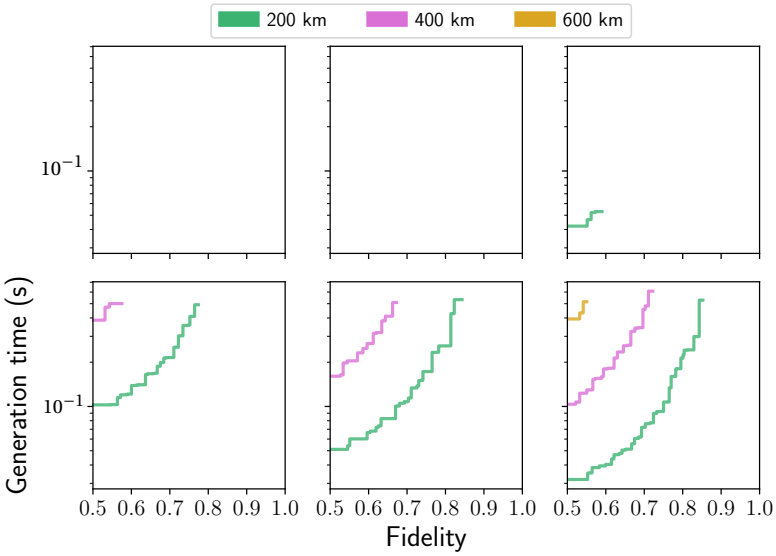

Figure 4.16: Results of the heuristic optimisation for total distances of 200, 400 and 600 kilometres, for MP platforms and ten intermediate nodes. We use parameter set 2 for MP platforms (see Table 4.3) as a baseline, where we set the success probability of the Bell state measurements to $\frac{3}{4}$, $\frac{7}{8}$, $\frac{15}{16}$ in the first, second, and third column, respectively. We set the efficiency coherence time T_{coh} to 1 and 10 in the first and second row, respectively.

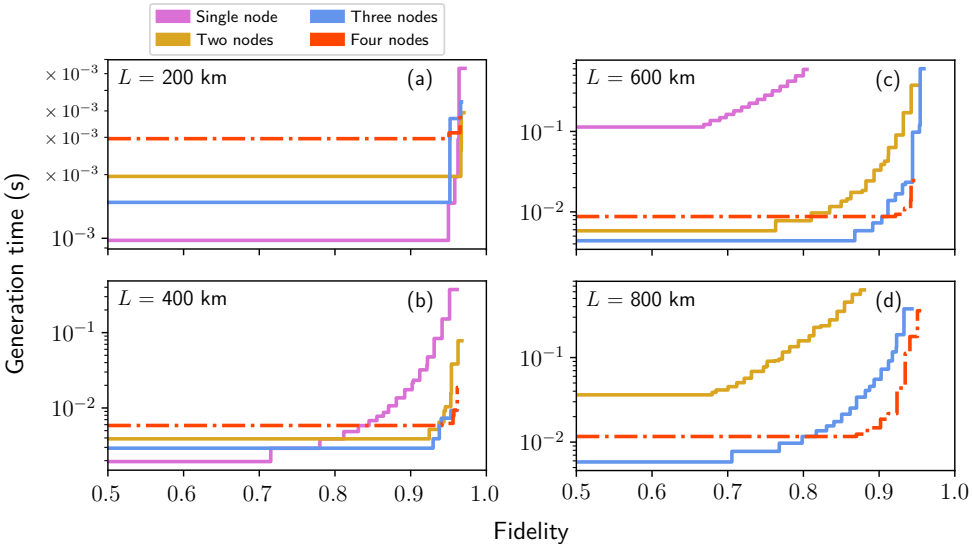

Figure 4.17: Results of the achieved fidelity and generation time for total distances of 200 (a), 400 (b), 600 (c) and 800 (c) kilometre using parameter set 4 (see Table 4.4) for MP platforms, where we consider having 1 (purple), 2 (yellow), 3 (blue) or 4 (orange) of such intermediate nodes.

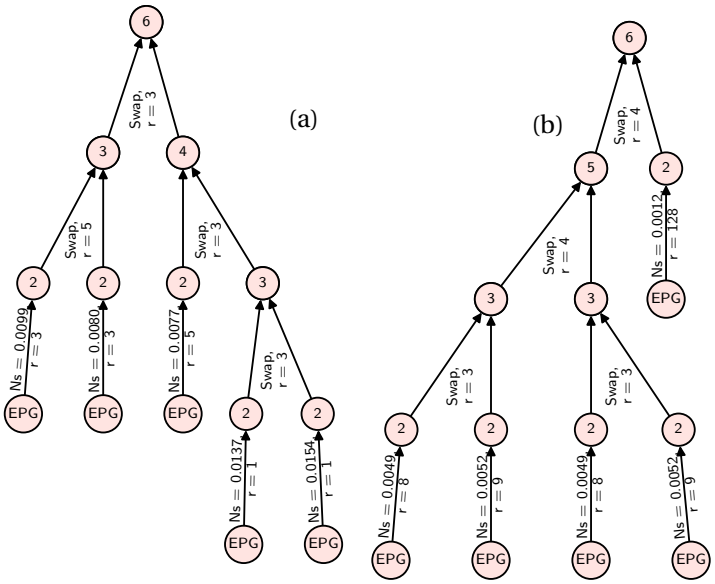

Figure 4.18: Visual representation of the schemes with the lowest non-trivial fidelity (a) and highest fidelity (b) respectively, for a distance of 800 kilometres with MP platforms using parameter set 4 (see Table 4.4) and four intermediate nodes/five hops. The 'N_s= N_s^*' indicates the elementary pair generation (EPG) protocol with mean photon number N_s^* used for MP platforms discussed in the main text. The '*r*' here indicates the number of rounds the corresponding subtree is attempted. Note that the second scheme requires a swap between schemes over multi-hop links of lengths five and two at the end.

of a lower success probability. Indeed, the number of attempts for the elementary pair generation range from 1 to 5 and from 8 to as high as 128, for the schemes in (a) and (b), respectively.

Here, we also note that there is a non-trivial interplay between the exponential decrease in output efficiency and performing more rounds (i.e. attempting more times to generate the elementary pairs) to increase the success probability. As we show in Appendix 4.6.7, the requirement that each step succeeds with probability at least $p_{\min}$ can lead to a scenario where under a slight change of the network/parameters, entanglement suddenly cannot be generated anymore.

Interestingly, we observe that the second scheme in Fig. 4.18 requires a swap between multi-hop links of lengths as five and two at the end. This shows that, as with information processing platforms, exploring more complex asymmetric schemes provides a benefit over more simplistic schemes.

4.4.3. LONG-DISTANCE ENTANGLEMENT GENERATION USING A COMBINATION OF IP AND MP PLATFORMS

Here we investigate combining the strengths of IP platforms with those of MP platforms. For this, we generate the elementary pairs with MP platforms, after which all the operations are performed with IP platforms. We optimise then over the same protocols as was done for IP and MP platforms, see Sections 4.4.1 and 4.4.2. We expect that, with sufficiently good parameters, the combination of the two outperforms the individual platforms, and that we can distribute entanglement over significantly larger distances.

Using the parameter set 4 of both platforms, we plot the results for 15, 25 and 35 nodes in Fig. 4.19, for a total distance of 4000 kilometre. Furthermore, we also plot a comparison here when the optimisation includes the bisection heuristic, see Section 4.2.4.

From Fig. 4.19 we observe that, by combining both the strengths from multiplexing and information processing platforms, it is possible to generate entanglement with a high fidelity near-

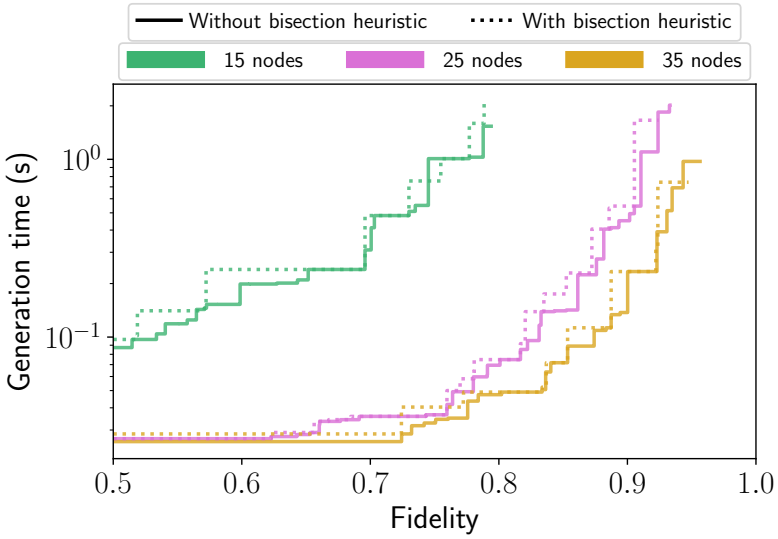

Figure 4.19: Optimisation results for a total distance of 4000 kilometres, using a combination of multiplexed and information processing platforms. We use parameter sets 4 from both the multiplexed and information processing part of the platform, see Tables 4.2 and 4.4. The solid lines are the optimisation without the bisection heuristic discussed in 4.2.4, while the dotted lines are with the bisection heuristic.

deterministically over large distances by using a large number of nodes. We find that the optimisation results with the bisection heuristic are similar to the results without, while being significantly faster to perform. We find for the cases of 15, 25, and 35 intermediate nodes that the algorithm runtime drops from an order of magnitude of ~100 minutes to $\lesssim$10 minutes. We thus find that the bisection heuristic allows for a faster heuristic optimisation, without the resultant schemes becoming significantly worse than without the bisection heuristic.

We conclude our results with a plot comparing entanglement generation with the three implementations considered in this chapter for a distance of 800 kilometres and five or ten intermediate nodes. We find in Fig. 4.20 that, for large distances, the combination of IP and MP platforms outperforms the individual platforms. In fact, it can generate target fidelities below ~ 0.9 an order of magnitude faster than the MP platform. We see that, as expected, information processing platforms perform significantly worse, where the maximum fidelity is limited to around ~ 0.6. This is due to the effects of losses during elementary pair generation becoming too strong. This can of course be counteracted by using more nodes, but this results in too much decoherence. This suggests that, for large distances, MP platforms outperform IP platforms for near-deterministic entanglement generation.

We depict the two schemes corresponding to the two crosses found in Fig. 4.20 in Figs. 4.26 and 4.27 in Appendix 4.6.8, respectively. The first of these (blue cross) corresponds to the lowest non-trivial fidelity achieved, while the second one (red cross) corresponds to a state with fidelity of $F = 0.9605$, generated in time $T = 17.7$ milliseconds. A higher fidelity was not chosen, due to those schemes becoming too big to fit on a page, demonstrating the non-trivial nature of the optimisation performed here.

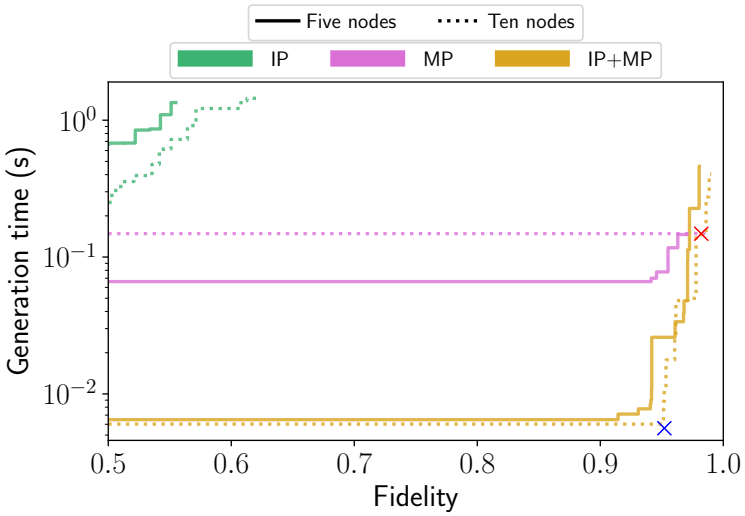

Figure 4.20: Results of the heuristic optimisation for a total distance of 800 kilometres, where we compare the three implementations considered in this chapter, using five (solid) or ten (dashed) intermediate nodes. We use parameter sets 4 from information processing (IP) platforms, multiplexed elementary pair generation (MP) platforms and the combination of the two (IP+MP). The two crosses in the plot indicate the schemes depicted in Figs. 4.26 and 4.27, respectively.

4.5. CONCLUSIONS

The future quantum internet has the potential to change our information society by enabling the implementation of quantum communication tasks. For many of these tasks the key resource is the availability of high fidelity entanglement at the necessary rates. However, given the complex relation between experimental parameters, entanglement distribution protocols and quantum network design, it is unclear what are the necessary parameters to distribute entanglement except for the most basic scenarios. Here, we develop an algorithm to partially answer this question. In particular, our algorithm optimises the near-deterministic distribution of entanglement over chains of quantum repeaters which are abstractly characterised by a small set of relevant parameters.

Even in this abstract setting, the number of possible protocols for a given quantum repeater chain is too large to attempt brute-force optimisation. To make optimisation feasible, we introduce a number of heuristics that render optimisation feasible by dramatically reducing the runtime of the algorithm. Moreover, the heuristics can also be interpreted as approximate rules for protocol design as numerical results show that optimal protocols follow the heuristics. We could expect these heuristics to apply to more dynamic schemes, where the information of the current present entanglement in the network is used to make decisions on the fly by the network.

Any realistic quantum repeater network will be asymmetric in the distances between the nodes and the experimental parameters. We have applied our algorithm to an asymmetric repeater chain, and have found that our optimisation results strongly outperform the results from a simplified optimisation over symmetric/hierarchical schemes, such as those presented in [20, 48].

We have used the algorithm not only for optimising entanglement distribution, but also for parameter exploration. In particular, we have optimised entanglement distribution for several parameter regimes investigating the most relevant parameters for both information processing and multiplexed elementary pair generation platforms. For both, we find that success probabilities (e.g. the emission probabilities, detector efficiencies, etc.) have a strong impact on performance.

In contrast with previous work, our focus on near-deterministic schemes allowed us to make

exact statements about the generation time and fidelities of the distributed states. The ability to deliver states with high probability at specific times could be of benefit for routing entanglement in a network.

In conclusion, here we have developed an algorithm that allows to efficiently optimise and explore the parameter space for near-deterministic entanglement distribution over repeater chains. We have investigated a number of representative platforms but the algorithm is not particular to these choices. We make the source code publicly available [58] to facilitate the investigation of other implementations, parameters and/or error models.

4.6. APPENDIX

4.6.1. COMPLEXITY OF THE ALGORITHM

Here we discuss the complexity of the algorithm. For this, first we bound from below the number of schemes that a brute-force approach without heuristics would need to explore. We then incorporate the heuristics and derive an upper bound on the number of schemes of the algorithm as described in Section 4.2.4. We finalise by deriving an upper bound on the number of schemes in the particular case of 'symmetric' repeater chains. That is, chains where each node has the same parameters and adjacent nodes are connected by identical elementary links.

4.6.2. A LOWER BOUND ON THE COMPLEXITY OF THE BRUTE-FORCE ALGORITHM

Here we derive two lower bounds on the number of schemes considered by a brute-force algorithm. The two lower bounds are given by $\mathcal{O}\left(\left(r_{\text{discr}}\cdot|\mathscr{E}|\cdot|\mathscr{S}|\cdot|\mathscr{D}|\right)^{2^{m\cdot n}}\right)$ and $\mathcal{O}\left(\left(\left(r_{\text{discr}}\right)^2\cdot|\mathscr{E}|\cdot|\mathscr{S}|\right)^n\right)$. These bounds correspond to the case with and without distillation protocols considered, respectively. Here, n denotes the number of elementary links in the repeater chain (i.e. one less than the number of nodes), m denotes the maximum number of distillation rounds, r_{discr} the maximum different values of the number of attempts, and $|\mathscr{E}|$, $|\mathscr{S}|$ and $|\mathscr{D}|$ denote the number of elementary pair generation, swapping and distillation protocols, respectively.

To make the analysis tractable, while still obtaining a strict lower bound on the number of schemes, we analyse a simpler algorithm that explores a reduced set of swapping schemes. At level i, instead of exploring all combinations of swapping between schemes for every pair of adjacent (multi-hop) links with a combined length i, this algorithm only considers swapping between schemes on the leftmost link of length $i-1$ and one of length 1, i.e. the entanglement is propagated by one elementary link at each level. Furthermore, we will assume the worst case scenario, where all generated schemes have success probability greater than p_{min}, meaning that all of them will be stored.

We first present a sketch of our derivation of the lower bound. We find two maps f_{swap} / f_{distill} which send the number of schemes ζ before the swap operation/m distillation round to a lower bound on the number of schemes after the swap/distillations. Denoting by ζ_{init} the number of schemes over an elementary link after distillation, a lower bound for the number of schemes after two hops is then given by:

$$\left(f_{\text{distill}}\circ f_{\text{swap}}\right)\left(\zeta_{\text{init}}\right)\,. \tag{4.11}$$

Similarly, after $n-1$ hops we find the following lower bound

$$\left(\left(f_{\text{distill}}\circ f_{\text{swap}}\right)\right)^{n-1}\left(\zeta_{\text{init}}\right)\,. \tag{4.12}$$

In what follows, we will find the maps f_{swap}, f_{distill} and ζ_{init}. The map f_{swap} will depend implicitly on r_{discr}, $|\mathscr{S}|$ and ζ_{init}, while f_{distill} depends implicitly on r_{discr}, $|\mathscr{D}|$ and m.

Let us start with f_{swap}. As mentioned above, for each multi-hop link of length i, the simplified algorithm combines each of the ζ schemes of the multi-hop link of length $i-1$ with each of the ζ_{init} schemes stored for an elementary link. We obtain the following map on the number of schemes,

$$\zeta \xrightarrow{f_{\text{swap}}} r_{\text{discr}} \cdot |\mathscr{S}| \cdot \zeta_{\text{init}} \cdot \zeta \,. \tag{4.13}$$

Let us now find the map f_{distill}. Assuming we start with ζ schemes, each of the $|\mathscr{D}|$ distillations will generate a new scheme for each possible pair of schemes, and for each of the r_{discr} possible values of attempts. Thus, after a single distillation round, we end up with $r_{\text{discr}} \cdot |\mathscr{D}| \cdot \zeta^2 + \zeta \geq r_{\text{discr}} \cdot |\mathscr{D}| \cdot \zeta^2$ schemes, where we have only kept the schemes which had a distillation step at the end.

We can now repeat the above for $m = 2$ distillation rounds, by setting $\zeta = r_{\text{discr}} |\mathscr{D}| \zeta^2$. We then find that for $m = 2$ distillation rounds that a lower bound is given by $r_{\text{discr}} \cdot |\mathscr{D}| \cdot \left(r_{\text{discr}} \cdot |\mathscr{D}| \cdot \zeta^2\right)^2$. In general, after m distillation rounds we find that

$$\zeta \xrightarrow{f_{\text{distill}}} \left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \zeta^{2^m} \,. \tag{4.14}$$

Thus, starting from a number ζ of schemes, the composition of swapping with a scheme on an elementary link (equation (4.13)) and then distilling (equation (4.14)) gives us the following map for the lower bound on the number of schemes

$$\begin{aligned} \zeta &\xrightarrow{f_{\text{distill}} \circ f_{\text{swap}}} \left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \left(r_{\text{discr}} \cdot |\mathscr{S}| \cdot \zeta_{\text{init}} \cdot \zeta\right)^{2^m} \\ &= \Omega \cdot \zeta^{2^m} \,, \end{aligned} \tag{4.15}$$

where we define $\Omega \equiv \left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \left(r_{\text{discr}} \cdot |\mathscr{S}| \cdot \zeta_{\text{init}}\right)^{2^m}$, which is independent of ζ.

Repeating the above map in equation (4.15) $n-1$ times on ζ_{init} (the number of schemes stored over an elementary link) yields the following lower bound,

$$\begin{aligned} \zeta_{\text{init}} &\xrightarrow{f_{\text{distill}} \circ f_{\text{swap}}} \Omega \cdot \left(\zeta_{\text{init}}\right)^{2^m} \\ &\xrightarrow{f_{\text{distill}} \circ f_{\text{swap}}} \Omega \cdot \left(\Omega \cdot \left(\zeta_{\text{init}}\right)^{2^m}\right)^{2^m} \xrightarrow{f_{\text{distill}} \circ f_{\text{swap}}} \cdots \\ &= \Omega^{1+2^m+2^{2\cdot m}+\ldots+2^{m\cdot(n-2)}} \cdot {\zeta_{\text{init}}}^{2^{m(n-1)}} \\ &= \Omega^{\frac{2^{m(n-1)}-1}{2^m-1}} \cdot \left(\zeta_{\text{init}}\right)^{2^{m(n-1)}} \,. \end{aligned} \tag{4.16}$$

The only ingredient missing from our analysis now is ζ_{init}, the number of schemes on an elementary link after m distillation rounds. First, for elementary pair generation there are r_{discr} different values of attempts per elementary pair generation protocol. In other words, for each of the $|\mathscr{E}|$ elementary pair generation protocols that can be performed, there are r_{discr} different choices of r, leading to a total of $|\mathscr{E}| \cdot r_{\text{discr}}$ schemes. To find the number of schemes after m distillation rounds we apply our map f_{distill} to $|\mathscr{E}| \cdot r_{\text{discr}}$ find that

$$\zeta_{\text{init}} \equiv \left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \left(|\mathscr{E}| \cdot r_{\text{discr}}\right)^{2^m} \,. \tag{4.17}$$

Inserting equation (4.17) into (4.16) and expanding gives

$$\begin{aligned}
&\Omega^{\frac{2^{m(n-1)}-1}{2^m-1}} \cdot \left(\zeta_{\text{init}}\right)^{2^{m(n-1)}} \\
&= \left(\left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \left(r_{\text{discr}} \cdot |\mathscr{S}| \cdot \left(\left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \left(|\mathscr{E}| \cdot r_{\text{discr}}\right)^{2^m}\right)^{2^m}\right)^{\frac{2^{m(n-1)}-1}{2^m-1}}\right. \\
&\quad \left. \cdot \left(\left(r_{\text{discr}} \cdot |\mathscr{D}|\right)^{2^m-1} \cdot \left(r_{\text{discr}} \cdot |\mathscr{E}|\right)^{2^m}\right)^{2^{m(n-1)}}\right. \\
&= \left(r_{\text{discr}}\right)^{\left(2^{m(n-2)}\right)\cdot\left(4^{m+1}-1\right)} \cdot |\mathscr{E}|^{\left(2^{m\cdot n+1}\right)\cdot\left(2^m+1\right)} \cdot |\mathscr{S}|^{2^{m(n-1)}} \cdot |\mathscr{D}|^{\left(2^{m(n-2)}\right)\cdot\left(2^{2m+1}-2^m-1\right)} \\
&= \mathscr{O}\left(\left(r_{\text{discr}} \cdot |\mathscr{E}| \cdot |\mathscr{S}| \cdot |\mathscr{D}|\right)^{2^{m\cdot n}}\right) . \qquad (4.18)
\end{aligned}$$

We note that this bound becomes trivial when no distillation is performed, i.e. $m = 0$. This is due to the fact that lower order terms were ignored in the number of schemes after distilling. We treat the $m = 0$ case separately here. For the case of no distillation, we perform $r_{\text{discr}} \cdot |\mathscr{S}|$ different swap protocols for $n-1$ times. Since we start with a total number of $r_{\text{discr}} \cdot |\mathscr{E}|$ schemes on the elementary links, the total number of schemes is then given by

$$\begin{aligned}
\left(r_{\text{discr}} \cdot |\mathscr{S}|\right)^{n-1} \cdot \left(r_{\text{discr}} \cdot |\mathscr{E}|\right)^n &= \left(r_{\text{discr}}\right)^{2n-1} \left(|\mathscr{E}| \cdot |\mathscr{S}|\right)^{n-1} \\
&= \mathscr{O}\left(\left(\left(r_{\text{discr}}\right)^2 \cdot |\mathscr{E}| \cdot |\mathscr{S}|\right)^n\right) . \qquad (4.19)
\end{aligned}$$

We see that with distillation (i.e. $m \geq 1$) the number of schemes to consider grows super-exponentially in the number of elementary links n, while without distillation $m = 0$ the number of schemes grows exponentially. It is clear that a brute-force optimisation becomes infeasible for any reasonable number of protocols (i.e. $|\mathscr{E}|$, $|\mathscr{S}|$, $|\mathscr{D}|$), number of distillation rounds m and elementary links n.

AN UPPER BOUND ON THE COMPLEXITY OF THE HEURISTIC ALGORITHM

In this section we consider the complexity with the heuristics implemented. The upper bounds we find scales as $\mathscr{O}\left(n^2 \log(n)\right)$ for an arbitrary repeater chain, where n is the number of elementary links in the repeater chain. As discussed in the main text, the optimisation can be simplified for the scenario of a repeater chain where every node has exactly the same parameters and the distance between each of the repeaters is equal. For such a symmetric repeater chain, we find a scaling of $\mathscr{O}\left(n \log(n)\right)$.

Let us first briefly discuss the effects the heuristics have on the complexity, before upper bounding the number of schemes. First off, we note here that in the worst-case scenario, all the schemes are incomparable, leading to no pruning. Secondly, the coarse-graining of the fidelity and probability imposes an upper limit on the considered schemes. The coarse-graining fixes the maximum stored schemes to be $\lceil \frac{(1-F_{\text{threshold}})}{\varepsilon_F} \rceil \lceil \frac{(1-p_{\min})}{\varepsilon_p} \rceil$ per (multi-hop) link. For instance, the number of schemes for elementary pair generation does not change with the heuristic, namely it remains $n|\mathscr{E}| \cdot r_{\text{discr}}$ in total. However, at most $\lceil \frac{(1-F_{\text{threshold}})}{\varepsilon_F} \rceil \lceil \frac{(1-p_{\min})}{\varepsilon_p} \rceil$ of these are stored per elementary link.

Let us now consider swapping. The algorithm restricts the creation of scheme on a multi-hop link of length i (equivalently, a link requiring i hops) to swapping two links of length $\frac{i}{2} \pm \log(i-1)$, leading to at most $2\lfloor \log(i-1) \rfloor + 1$ different options, see equation 4.5. The banded swapping heuristic further reduces swapping to schemes that verify equation 4.6 from the main text,

$$\left| \frac{\log(F_1)}{i_1} - \frac{\log(F_2)}{i_2} \right| \leq \varepsilon_{\text{swap}} .$$

This equation becomes in the asymptotic limit

$$\frac{F_1}{F_2} \leq \exp\left(\frac{i\cdot\varepsilon_{\text{swap}}}{2}\right),$$

since $i_1 \sim i_2 \sim \frac{i}{2}$, and where we have assumed without loss of generality that $F_1 \geq F_2$. Now note that $F_{\text{threshold}} \geq \frac{1}{2}$, which implies that $\frac{1}{2} \leq \frac{F_1}{F_2} \leq 2$. We thus have that, in the asymptotic limit, the banded swapping heuristic becomes void if $\varepsilon_{\text{swap}}$ is fixed. This means that asymptotically the algorithm considers the full $r_{\text{discr}} \cdot \frac{(1-F_{\text{threshold}})^2(1-p_{\min})^2}{(\varepsilon_F\varepsilon_p)^2}$ schemes for swapping.

The last heuristic is banded distillation. For a fixed distillation protocol, it reduces the number of schemes for performing distillation to at most $2\cdot\lceil\frac{\varepsilon_{\text{distill}}}{\varepsilon_F}\rceil\lceil\frac{1-p_{\min}}{\varepsilon_p}\rceil$. Since there are in total $\lceil\frac{(1-F_{\text{threshold}})}{\varepsilon_F}\rceil\lceil\frac{(1-p_{\min})}{\varepsilon_p}\rceil$ stored schemes, we find the following upper bound on the considered schemes for distillation for a single (multi-hop) link, $2\cdot r_{\text{discr}}\cdot|\mathscr{D}|\cdot m\cdot\frac{\varepsilon_{\text{distill}}(1-F_{\text{threshold}})(1-p_{\min})^2}{(\varepsilon_F\varepsilon_p)^2}$. We have removed here and in what follows the ceiling functions, since we are interested in the asymptotic complexity and increasing readability.

Combining the previous arguments, we find the following upper bound:

$$\begin{aligned}
&n|\mathscr{E}|\cdot r_{\text{discr}} + r_{\text{discr}}\sum_{i=2}^{n}(n-i+1)\cdot\left(\left(2\cdot\lfloor\log(i-1)\rfloor+1\right)\cdot|\mathscr{S}|\cdot\frac{\left(1-F_{\text{threshold}}\right)^2\left(1-p_{\min}\right)^2}{\left(\varepsilon_F\varepsilon_p\right)^2}\right.\\
&\qquad\left.+\,2\cdot m\cdot|\mathscr{D}|\cdot\frac{\varepsilon_{\text{distill}}\left(1-F_{\text{threshold}}\right)\left(1-p_{\min}\right)^2}{\left(\varepsilon_F\varepsilon_p\right)^2}\right)\\
=\ & r_{\text{discr}}\left(n|\mathscr{E}|+\frac{\left(1-F_{\text{threshold}}\right)\left(1-p_{\min}\right)^2}{\left(\varepsilon_F\varepsilon_p\right)^2}\sum_{i=2}^{n}(n-i+1)\cdot|\mathscr{S}|\cdot\left(\left(2\cdot\lfloor\log(i)\rfloor+1\right)\left(1-F_{\text{threshold}}\right)+2\cdot m\cdot|\mathscr{D}|\cdot\varepsilon_{\text{distill}}\right)\right)\\
\sim\ & r_{\text{discr}}\left(n|\mathscr{E}|+\frac{\left(1-F_{\text{threshold}}\right)\left(1-p_{\min}\right)^2}{\left(\varepsilon_F\varepsilon_p\right)^2}\left(2\left(1-F_{\text{threshold}}\right)\cdot|\mathscr{S}|\cdot n^2\log(n)+2\cdot m\cdot|\mathscr{D}|\cdot\varepsilon_{\text{distill}}\cdot\frac{n^2}{2}\right)\right)\\
=\ & r_{\text{discr}}\left(n|\mathscr{E}|+n^2\frac{\left(1-F_{\text{threshold}}\right)\left(1-p_{\min}\right)^2}{\left(\varepsilon_F\varepsilon_p\right)^2}\left(2\left(1-F_{\text{threshold}}\right)\cdot|\mathscr{S}|\cdot\log(n)+m\cdot|\mathscr{D}|\cdot\varepsilon_{\text{distill}}\right)\right).
\end{aligned}$$

We observe that the algorithm is $\mathscr{O}\left(n^2\log(n)\right)$, where the pre-factor is given by

$$2\cdot|\mathscr{S}|\cdot r_{\text{discr}}\left(\frac{\left(1-F_{\text{threshold}}\right)\left(1-p_{\min}\right)}{\varepsilon_F\varepsilon_p}\right)^2. \tag{4.20}$$

In the case of a symmetric repeater chain (i.e. every node has exactly the same parameters and the distance between each of the repeaters is equal) we can simplify the optimisation by exploiting the symmetry. That is, the optimisation done over a (multi-hop) link of length i only needs to be done once, as opposed to $n-i+1$ times in the general setting. Furthermore, there are only $\lfloor\log(i-1)\rfloor+1$ unique ways to perform swapping. The number of schemes is then upper bounded by

$$|\mathscr{E}|\cdot r_{\text{discr}} + r_{\text{discr}} \sum_{i=2}^{n} \left(\lfloor (\log(i-1) \rfloor + 1) \cdot |\mathscr{S}| \cdot \frac{(1-F_{\text{threshold}})^2 (1-p_{\text{min}})^2}{(\varepsilon_F \varepsilon_p)^2} \right.$$
$$\left. + 2 \cdot m \cdot |\mathscr{D}| \cdot \frac{\varepsilon_{\text{distill}} (1-F_{\text{threshold}}) (1-p_{\text{min}})^2}{(\varepsilon_F \varepsilon_p)^2} \right)$$
$$\sim r_{\text{discr}} \cdot |\mathscr{E}| + r_{\text{discr}} \cdot \frac{(1-F_{\text{threshold}}) (1-p_{\text{min}})^2}{(\varepsilon_F \varepsilon_p)^2} \left((1-F_{\text{threshold}}) \cdot |\mathscr{S}| \cdot n \log(n) + n \cdot m \cdot |\mathscr{D}| \cdot \varepsilon_{\text{distill}} \right).$$

We observe that for the symmetric scenario, the algorithm is $\mathscr{O}(n \log(n))$, where the pre-factor is given by $r_{\text{discr}} \cdot |\mathscr{S}| \cdot \left(\frac{(1-F_{\text{threshold}})(1-p_{\text{min}})}{\varepsilon_F \varepsilon_p} \right)^2$. As mentioned before, the algorithm developed supports both the general and symmetric case.

4

4.6.3. ANALYSIS OF THE HEURISTICS

The algorithm detailed in this chapter uses four different parameters to reduce the search space of the optimisation, namely ε_F, ε_p, $\varepsilon_{\text{swap}}$, and $\varepsilon_{\text{distill}}$, see 4.2.4. The parameters ε_F and ε_p are responsible for the coarse-graining in the algorithm, while the parameters $\varepsilon_{\text{swap}}$ and $\varepsilon_{\text{distill}}$ govern the restrictions on the states used for swapping an distillation, respectively. In this section we investigate the heuristics and how they affect the algorithm runtime and accuracy of the optimisation, which we use to settle on values for ε_F, ε_p, $\varepsilon_{\text{swap}}$, and $\varepsilon_{\text{distill}}$. The objective here is to find a good trade-off between the algorithm runtime and the accuracy of the algorithm. We first investigate the coarse-graining - i.e. we vary ε_F and ε_p. Afterwards, we investigate the effects of $\varepsilon_{\text{swap}}$ and $\varepsilon_{\text{distill}}$ on the optimisation results. Finally, we compare the banded swapping heuristic with the naive heuristic, i.e. where we require the two states to be close in fidelity.

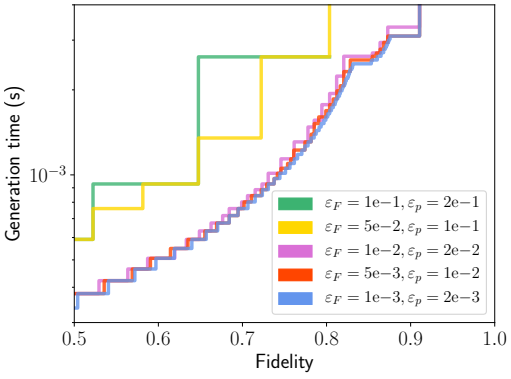

Figure 4.21: Optimised schemes for a distance of 15 kilometres using a single node with the IP parameter set 2 (see Table 4.2) for several different pairs of ε_F and ε_p. Note that as ε_F and ε_p approach zero, the curves converge.

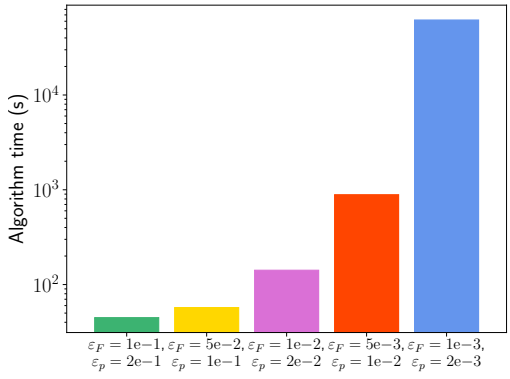

Figure 4.22: The runtime of the algorithm for the optimisations performed in Fig. 4.21. Notice the logarithmic scale, indicating the strong growth rate as ε_F and ε_p become smaller.

We first vary ε_F and ε_p simultaneously when optimising over schemes for a distance of 6 kilometres and a single repeater with the IP parameter set 2, of which the results can be seen in Fig. 4.21 and 4.22. As expected, there is a trade-off between the accuracy of the algorithm and

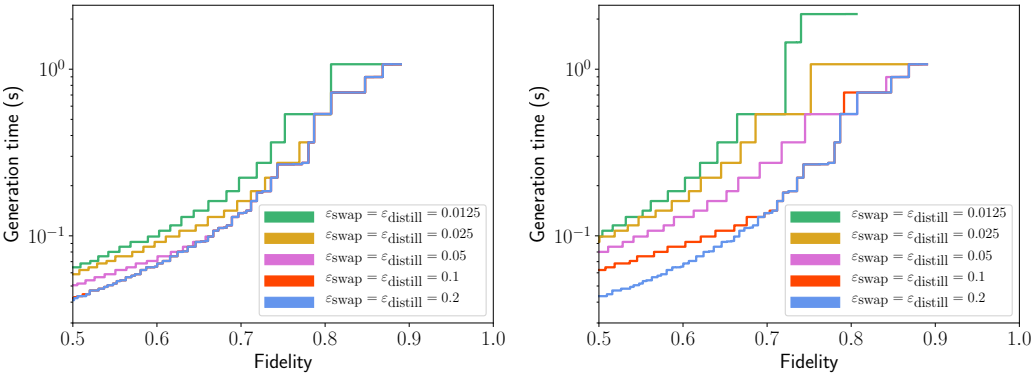

(a) Optimisation results with banded swapping heuristic. The difference between $\varepsilon_{\text{swap}} = \varepsilon_{\text{distill}} = 0.1$ and 0.2 is minimal, differing only slightly for very low fidelities.

(b) Optimisation results with naive swapping heuristic. Note that the curves converge in a significantly poorer fashion than in Fig. 4.23a.

Figure 4.23: Optimised schemes for a distance of 300 kilometres using four intermediate nodes with the IP parameter set 4 (see Table 4.2) for several different pairs of $\varepsilon_{\text{swap}}$ and $\varepsilon_{\text{distill}}$. Note that the curves converge in a significantly poorer fashion than in Fig. 4.23a.

its running time as ε_F and ε_p are varied. While a good trade-off between the accuracy and the runtime depends on each specific case, we use these results to settle in this chapter for $\varepsilon_F = 0.01$ and $\varepsilon_p = 0.02$. We settle for these parameters since the important characteristics of the generation time as a function of the fidelity appear to be similar when a more fine-grained optimisation is implemented, without the runtime becoming infeasible.

In Figs. 4.23a and 4.23b we perform an optimisation for several different values of $\varepsilon_{\text{swap}}$ and $\varepsilon_{\text{distill}}$, using parameter set 4 with four intermediate nodes for a distance of 300 kilometre. In Fig. 4.23a we use the banded swapping heuristic (see equation 4.6), while in Fig. 4.23b we only swap between states that are $\varepsilon_{\text{swap}}$-close in fidelity. We observe that the optimisation results in Fig. 4.23b are significantly worse than those in Fig. 4.23a, while Fig. 4.24 indicates that the runtimes are comparable for both heuristics. We use these results to settle on $\varepsilon_{\text{swap}} = \varepsilon_{\text{distill}} = 0.05$. Furthermore, we find that the heuristic plays primarily a role for smaller fidelities. This implies that only for small fidelities there is a benefit in swapping between states with disparate fidelities.

4.6.4. AVERAGE NOISE DUE TO STORAGE

Here we discuss the average noise induced when repeating a protocol with success probability p until success or until a maximum number of r attempts. Denote the quantum channel corresponding to storing for a single round by Λ. The average noise channel $\mathbb{E}[\Lambda]$ corresponds to having the channel $\underbrace{\Lambda \circ \Lambda \ldots \circ \Lambda}_{r-j}$ with probability $\frac{p(1-p)^{j-1}}{1-(1-p)^r}$ for $1 \leq j \leq r$. Note that we can calculate the average channel instead of the average density matrix at the output, due to the linearity of quantum channels.

We consider two types of noise in this chapter, depolarising and dephasing. These types of noise occur naturally in quantum information processing systems, and have the following exponential behaviour,

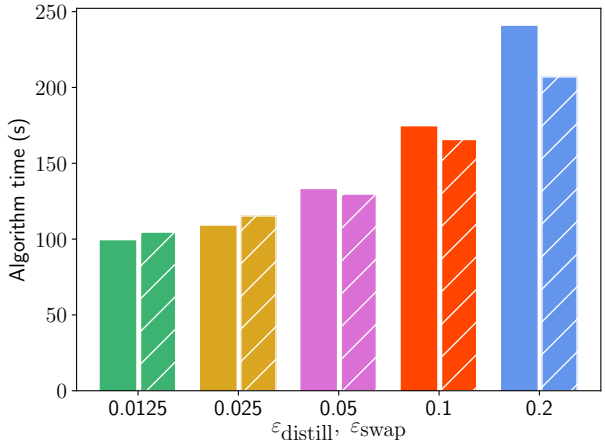

Figure 4.24: Algorithm runtimes for several different values of $\varepsilon_{\text{swap}}$ and $\varepsilon_{\text{distill}}$. The right, solid bars are for the optimisation with the heuristic for swapping found in Eq. 4.6, while the left, hatched bars are for the optimisation where we only swap between states that are $\varepsilon_{\text{swap}}$-close in fidelity. We observe that both heuristics lead to approximately the same runtime behaviour, while the results with the banded swapping heuristic are closer to optimal.

$$\mathcal{N}_d(\rho) = e^{-\frac{t}{\lambda_d}}\rho + \left(1 - e^{-\frac{t}{\lambda_d}}\right)\frac{\mathbb{I}}{2}, \text{ (depolarising)}$$

$$\mathcal{N}_Z(\rho) = \frac{1 + e^{-\frac{t}{\lambda_Z}}}{2}\rho + \frac{1 - e^{-\frac{t}{\lambda_Z}}}{2}\frac{\mathbb{I}}{2}, \text{ (dephasing)}$$

Thus, if we want to calculate the average amount of noise for depolarising and dephasing, it suffices to calculate the average of $e^{-c\cdot(k-j)}$ with probability distribution $\frac{p(1-p)^{j-1}}{1-(1-p)^r}$, $j = \{1,\ldots,r\}$, where $c \equiv \frac{T_{\text{attempt}}}{T_{\text{depol/deph}}}$ quantifies the noise experienced in a single attempt for depolarising and dephasing, respectively. We find thus that the average channels correspond to having the exponential terms in the above channels set to

$$\begin{aligned}\mathbb{E}\left[e^{-c\cdot(r-j)}\right] &= \sum_{j=1}^{r} \frac{p\left(1-p\right)^{j-1}}{1-\left(1-p\right)^r}\cdot e^{-c\cdot(r-j)} \\ &= \frac{pe^c\left(\left(1-p\right)^r - e^{-cr}\right)}{\left(1-\left(1-p\right)^r\right)\left(e^c\left(1-p\right)-1\right)}.\end{aligned} \tag{4.21}$$

Finally, the decay in the success probability for retrieving a state from a memory for MP platforms is given by $\mathbb{E}\left[e^{-c\cdot(r-j)}\right]$, where $c = \frac{T_{\text{attempt}}}{T_{\text{coh}}}$.

4.6.5. MODELLING OF ELEMENTARY PAIR GENERATION FOR MP PLATFORMS

Here we detail the calculations performed to derive the analytical form of the resultant state during elementary pair generation for MP platforms, and the success probability (see equation 4.22). We first discuss the effects of the losses on the state emitted by the PDC sources. Secondly, the Bell

state measurement and the resulting post-measurement state are discussed. We will close with a brief discussion on the post-selection of having zero photons. We model all losses in the setup as a pure-loss channel. Since we restrict ourselves to at most two-photon excitations in each mode, we truncate the Kraus operators from [74] to the $\{|0\rangle, |1\rangle, |2\rangle\}$ subspace and find the explicit matrix form of the truncated Kraus operators. They are

$$A_0 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \sqrt{1-\gamma} & 0 \\ 0 & 0 & 1-\gamma \end{bmatrix}, \; A_1 = \begin{bmatrix} 0 & \sqrt{\gamma} & 0 \\ 0 & 0 & \sqrt{2(1-\gamma)\gamma} \\ 0 & 0 & 0 \end{bmatrix}, \; A_2 = \begin{bmatrix} 0 & 0 & \gamma \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

where $\gamma = 1-\eta$ is the loss parameter. Note that, even after truncation, these Kraus operators still form a channel since $\sum_{i=0}^{2} A_i^\dagger A_i = \mathbb{I}$. We now let four such channels act on the state

$$|\psi_{N_s}\rangle = \sqrt{p_0}\,|00,00\rangle + \sqrt{\frac{p_1}{2}}\,(|10,01\rangle + |01,10\rangle) + \sqrt{\frac{p_2}{3}}\,(|20,02\rangle - |11,11\rangle + |02,20\rangle)\,,$$

where the early and late photonic modes in the direction towards the memory each evolve under a truncated pure-loss channel with parameter γ_1, and similarly for the two modes going towards the beamsplitter station with parameter γ_2. This results in a state $\rho(N_s,\gamma_1,\gamma_2)_{a_0a_1b_0b_1}$ between the memory and the photon just before the beamsplitter. The same situation holds for the other PDC source, such that the total state just before the beamsplitter is $\rho(N_s,\gamma_1,\gamma_2)_{a_0a_1b_0b_1} \otimes \rho(N_s,\gamma_1,\gamma_2)_{c_0c_1d_0d_1}$, where we have assumed the prepared states have equal mean photon number and experience equal losses. Instead of applying the unitary corresponding to the beamsplitter and then applying the POVMs for the detectors, we can apply the inverse of the beamsplitter unitary on the POVMs corresponding to success. Since we assume photon number resolving detectors, we find our POVM elements corresponding to success to be

$$\begin{aligned}
&\mathbb{I}_{a_0a_1} \otimes |\Psi^+\rangle\langle\Psi^+|_{b_0c_0} \otimes |\Psi^+\rangle\langle\Psi^+|_{b_1c_1} \otimes \mathbb{I}_{d_0d_1}\,, \\
&\mathbb{I}_{a_0a_1} \otimes |\Psi^-\rangle\langle\Psi^-|_{b_0c_0} \otimes |\Psi^+\rangle\langle\Psi^+|_{b_1c_1} \otimes \mathbb{I}_{d_0d_1}\,, \\
&\mathbb{I}_{a_0a_1} \otimes |\Psi^+\rangle\langle\Psi^+|_{b_0c_0} \otimes |\Psi^-\rangle\langle\Psi^-|_{b_1c_1} \otimes \mathbb{I}_{d_0d_1}\,, \\
&\mathbb{I}_{a_0a_1} \otimes |\Psi^-\rangle\langle\Psi^-|_{b_0c_0} \otimes |\Psi^-\rangle\langle\Psi^-|_{b_1c_1} \otimes \mathbb{I}_{d_0d_1}\,.
\end{aligned}$$

We find that the post-measurement states for each of these POVM element is equivalent up to local unitaries, such that we only have to consider the first one. While one could call the whole process described so far elementary pair generation, the state will have a fidelity equal to half or less for any $N_s > 0$. The reason for this is that there has not been any post-selection on detecting a valid click pattern on the detectors when performing, say, Bell state measurements. For this reason, we apply the following POVM to post-select on having non-zero photons at each side of the memory,

$$\mathbb{I}_{a_0a_1b_0b_1} - \left(|00\rangle\langle 00|_{a_0a_1} \otimes \mathbb{I} + \mathbb{I} \otimes |00\rangle\langle 00|_{d_0d_1} - |0000\rangle\langle 0000|_{a_0a_1d_0d_1}\right).$$

The resultant state is too cumbersome to report here, but can be found in the accompanying Mathematica and Python scripts. We find the success probability to be given by

$$\begin{aligned}
p_{\text{succ}} &= 4 \cdot p_{\text{bsm}} \cdot p_{\text{non-zero photons}} \\
&= 4 \cdot \eta^2 \frac{3p_1^2 - 4(4\eta-3)\,p_1p_2 + 4p_2\left(1+(3-8\eta+4\eta^2)\right)}{24} \\
&\cdot p_{\text{app}}^2 \frac{\left(p_1 + 4(\eta-1)\,p_2\,(p_{\text{app}}-2)\right)\left(3p_1 + 4(\eta-1)\,p_2\,(p_{\text{app}}-2)\right)}{4p_2 + \left(p_1 + (2-4\eta)\,p_2\right)\left(3p_1 + (6-4\eta)\,p_2\right)}\,.
\end{aligned} \tag{4.22}$$

4.6.6. THE INTERPLAY BETWEEN THE NUMBER OF MODES AND THE FIDELITY FOR MP PLATFORMS

Here we investigate the interplay between the number of modes, the fidelity, and the losses in the fibre for MP platforms. We assume here that the only source of noise is from the PDC source, and there are no or negligible losses locally. We take the state derived in the previous section (but which is too cumbersome to report here), and set $p_{\text{app}} = 1$. The fidelity of the resultant state is then calculated to be

$$F = \frac{3}{4(\eta-1)^2 N_s{}^4 + 24(\eta-1)^2 N_s{}^3 + 4\left(9\eta^2 - 20\eta + 11\right) N_s{}^2 - 24(\eta-1)N_s + 3}.$$

Solving for N_s, we find that

$$N_s = \frac{1}{2}\left(\sqrt{\frac{-9\eta F + 5F + 2\sqrt{F(F+3)}}{F - \eta F}} - 3\right).$$

Let us now input the above relation into equation 4.22 where we set p_1 and p_2 according to equations 4.10. We find a success probability of

$$p = \frac{32\eta^2\left(\sqrt{\frac{-9F\eta+5F+2\sqrt{F(F+3)}}{F-F\eta}} - 3\right)^2}{F\left(\sqrt{\frac{-9F\eta+5F+2\sqrt{F(F+3)}}{F-F\eta}} - 1\right)^6}. \tag{4.23}$$

Since we need on the order of $\frac{1}{p}$ modes, we find that we need on the order of

$$\frac{1}{p} = \frac{F\left(\sqrt{\frac{-9F\eta+5F+2\sqrt{F(F+3)}}{F-F\eta}} - 1\right)^6}{32\eta^2\left(\sqrt{\frac{-9F\eta+5F+2\sqrt{F(F+3)}}{F-F\eta}} - 3\right)^2} = \frac{F\left(\sqrt{\frac{2\sqrt{F(F+3)}}{F}} + 5 - 1\right)^6}{32\eta^2\left(\sqrt{\frac{2\sqrt{F(F+3)}}{F}} + 5 - 3\right)^2} + \mathcal{O}\left(\eta^{-1}\right)$$

modes to achieve a fidelity of F for $\eta \approx 0$. The $\eta^2 = \exp\left(-\frac{L}{L_0}\right)$ term in the denominator is given by the total losses of the fibre. The contribution due to the fidelity is then given by

$$\frac{F\left(\sqrt{\frac{2\sqrt{F(F+3)}}{F}} + 5 - 1\right)^6}{32\left(\sqrt{\frac{2\sqrt{F(F+3)}}{F}} + 5 - 3\right)^2} = \frac{32}{(1-F)^2} + \mathcal{O}\left((1-F)^{-1}\right).$$

We thus find that the number of minimum required modes scales as $\frac{e^{\frac{L}{L_0}}}{(1-F)^2}$, where L is the internode distance.

4.6.7. THE EFFECT OF EFFICIENCY DECOHERENCE FOR MP PLATFORMS

In this section we explore the effects the exponential decrease of the output efficiency has on the ability of performing schemes with probability greater than $p_{\min}$. While for information processing platforms it is always possible to achieve any success probability by performing as many attempts r as required, this is not the case for MP platforms due to the decrease in output efficiency over time. Here we derive conditions on the efficiency coherence times of the memories for a given $p_{\min}$, generation time T and success probability p of the underlying schemes, such that $p_{\min}$ can be achieved.

Since there are two memories used for state storage, the success probability of emitting both states again is modelled as given by

$$\begin{aligned} p_{\text{single success}} &= \left(1-(1-p)^r\right)\cdot \mathbb{E}\left[e^{-(c_1+c_2)\cdot(r-j)}\right] \\ &= \frac{pe^{(c_1+c_2)}\left((1-p)^r - e^{-(c_1+c_2)r}\right)}{e^{(c_1+c_2)}(1-p)-1}\,. \end{aligned} \tag{4.24}$$

We are interested in when the above quantity cannot be larger than $p_{\min}$. To this end, we take the derivative of equation 4.24 with respect to r and set it to zero to find the maximum value of success probability. Setting $c = c_1 + c_2$, we find

$$\begin{aligned} &\frac{e^c p\left(ce^{-cr} + (1-p)^r \log(1-p)\right)}{e^c(1-p)-1} = 0\,, \\ &\rightarrow r = \frac{c - \log\left(-\frac{e^c \log(1-p)}{c}\right)}{c+\log(1-p)}\,. \end{aligned} \tag{4.25}$$

However, since r needs to be an integer equal to or greater than one, we choose the ceiling or floor of equation 4.25, whichever maximises the resultant p_{succ}. Furthermore, since we cannot perform distillation, our main concern is the drop in success probability after performing a Bell state measurement. This motivates us to set $p = 1 - \frac{1}{2^{N+1}}$ [50, 62, 93, 100, 164]. Setting equation 4.24 equal to $p_{\min} = 0.9$, we numerically find that $N = 0$ gives $c \approx 0.023$, $N = 1$ gives $c \approx 0.053$, $N = 2$ gives $c \approx 0.101$ and $N = 3$ gives $c = \infty$. Obviously, the assumption here is that the initial success probability is given by $1 - \frac{1}{2^{N+1}}$, which is not true due to other losses in the system. However, it is clear that increasing N can increase the total time significantly during which entanglement can be generated in a near-deterministic fashion. In particular, we find that the sum of the reciprocals of the efficiency coherence times of the memories should be *at least* $\frac{1}{c}$ times the generation time of a scheme for MP platforms to successfully generate entanglement near-deterministically. This results in factors of approximately 43, 19 and 10 times the generation time for Bell state measurement success probabilities of $\frac{1}{2}$, $\frac{3}{4}$ and $\frac{7}{8}$, respectively.

4.6.8. ADDITIONAL OPTIMISATION RESULTS

This section contains the additional figures mentioned in the main text.

First, we compare the optimisation results of the full optimisation with an optimisation over BDCZ schemes only in Fig. 4.25. BDCZ schemes are those schemes that for each connection and distillation step only combines two schemes that have used the same sequence of protocols, as in [20, 48]. We consider an asymmetric repeater chain with three intermediate nodes over a distance of 200 kilometre. We model the behavior of the intermediate nodes with parameter set 4 and Alice and Bob with parameter set 2 (see Table 4.2). The full optimisation yields schemes that can

achieve faster generation rates (by approximately a factor of 10) than achievable with the BDCZ schemes.

Second, we consider two visualisations (Fig. 4.26 and 4.27) of the schemes found for a distance of 800 kilometres with a combination of IP and MP platforms using parameter sets 4 (see Tables 4.2 and 4.4) and ten intermediate nodes. Finally, Fig. 4.28 contains the results found while performing a parameter exploration for total distances of 200, 400, 600 and 800 kilometres, using ten intermediate nodes for information processing platforms.

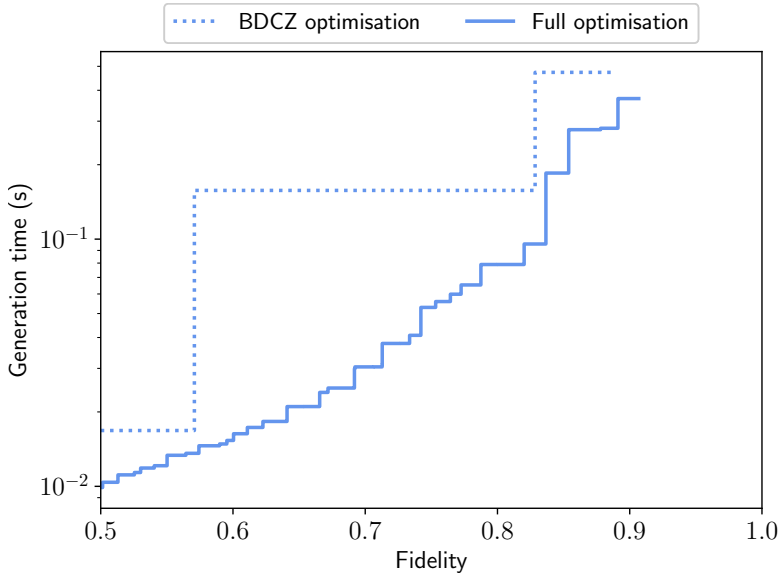

Figure 4.25: Comparison between the optimisation results of the full optimisation with the optimisation over BDCZ schemes. We consider a repeater chain over 200 kilometres with three intermediate nodes. The parameters used are parameter set 4 for the intermediate nodes and parameter set 2 for Alice and Bob. BDCZ schemes are those schemes that only perform swapping and distillation between two schemes that have used the same sequences of protocols. Contrary to the comparison with BDCZ schemes in [77] we allow for an optimisation over the different ways of generating elementary pairs, i.e. we vary the number of attempts r and the θ parameter. We observe that the schemes found with the full optimisation outperform BDCZ schemes, achieving a faster generation time by a factor of ~10, and extending the maximal achievable fidelity by a small margin.

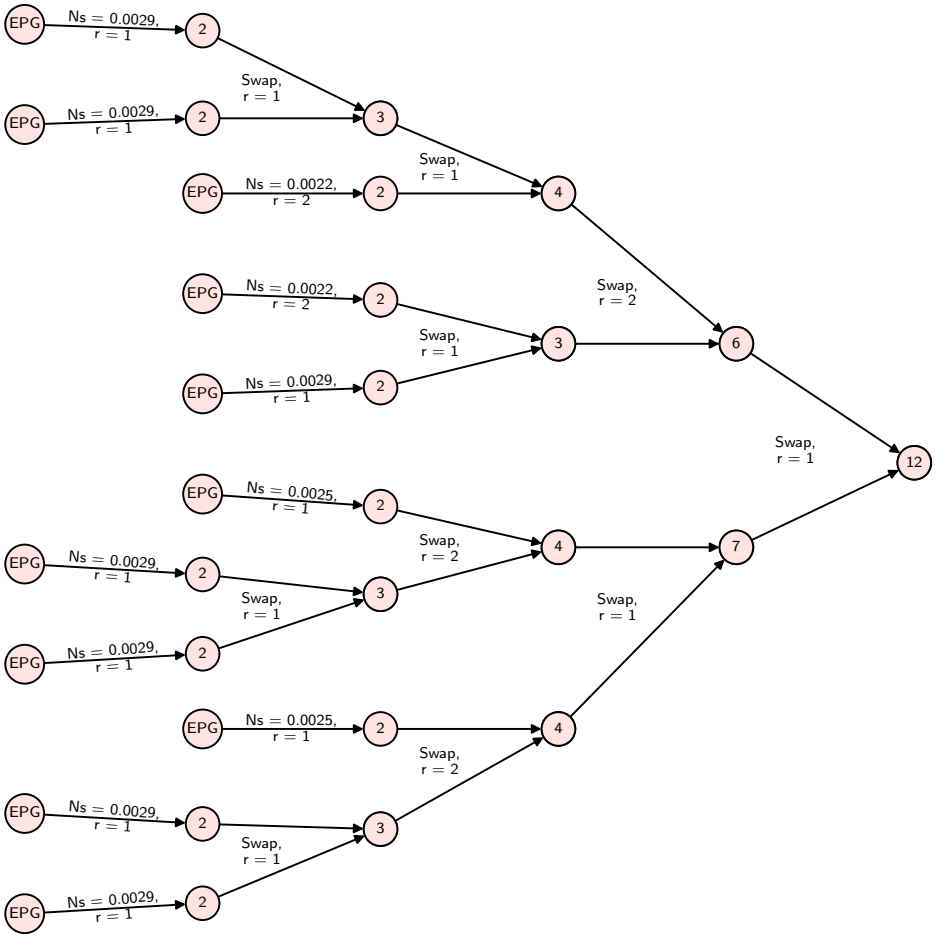

Figure 4.26: Visual representation of the scheme with the lowest non-trivial achieved fidelity for a distance of 800 kilometres with a combination of IP and MP platforms using parameter sets 4 (see Tables 4.2 and 4.4) and ten intermediate nodes/eleven hops. Elementary pair generation is indicated by EPG, and the mean photon number used is indicated by the N_s.

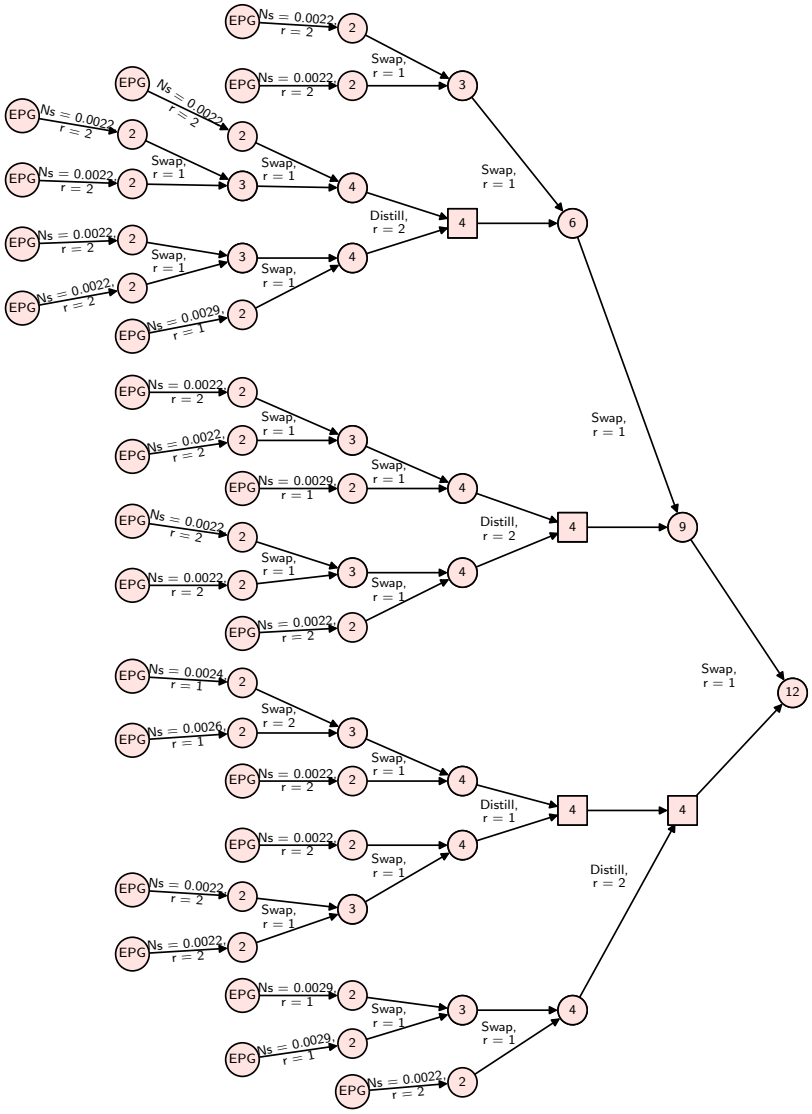

Figure 4.27: Visual representation of the scheme that achieves a fidelity of $F = 0.9605$ in time $T = 17.7$ milliseconds (indicated by the cross in Fig. 4.20), for a distance of 800 kilometres with a combination of IP and MP platforms using parameter sets 4 (see Tables 4.2 and 4.4) and ten intermediate nodes/eleven hops. Elementary pair generation is indicated by EPG, and the mean photon number used is indicated by the N_s. The structure of the scheme is non-hierarchical, which can most clearly be seen in the final swap operation, which happens between two multi-hop links of lengths four and nine.

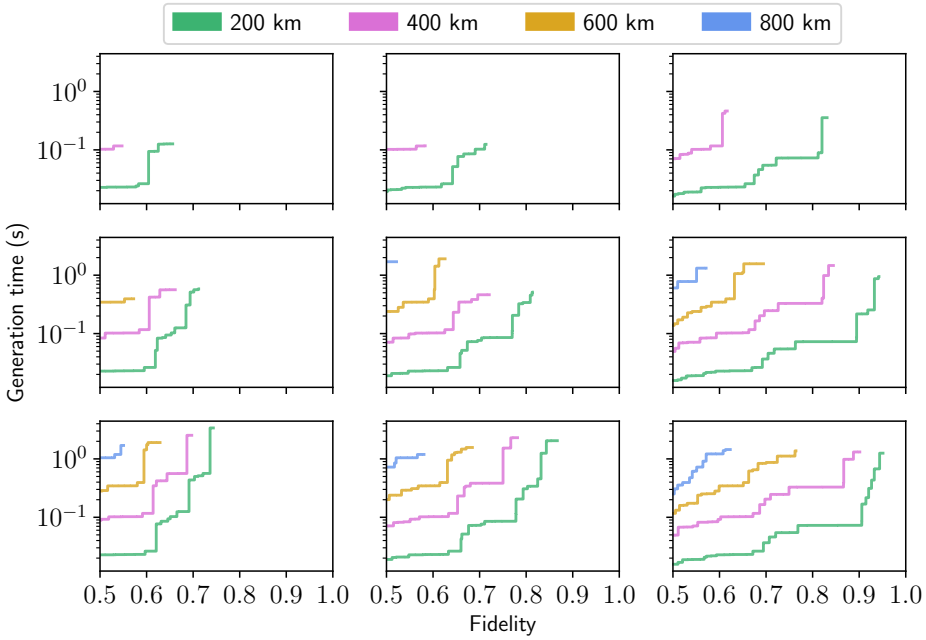

Figure 4.28: Optimisation results for total distances of 200, 400, 600 and 800 kilometres, using ten intermediate nodes. We use IP parameter set 2 as a baseline, where we set the gate fidelities to be 0.99, 0.995 and 0.999 in the first, second and third column respectively. We set the coherence times T_{deph}, T_{depol} to 10, 50 and 100 seconds in the first, second and third row, respectively.

5

ENUMERATING ALL BILOCAL CLIFFORD DISTILLATION PROTOCOLS THROUGH SYMMETRY REDUCTION

Sarah Jansen, Kenneth Goodenough, Sébastian de Bone, Dion Gijswijt and David Elkouss

Entanglement distillation is an essential building block in quantum communication protocols. Here, we study the class of near-term implementable distillation protocols that use bilocal Clifford operations followed by a single round of communication. We introduce tools to enumerate and optimise over all protocols for up to $n = 5$ (not necessarily equal) Bell-diagonal states using a commodity desktop computer. Furthermore, by exploiting the symmetries of the input states, we find all protocols for up to $n = 8$ copies of a Werner state. For the latter case, we present circuits that achieve the highest fidelity. These circuits have modest depth and number of two-qubit gates. Our results are based on a correspondence between distillation protocols and double cosets of the symplectic group, and improve on previously known protocols.

This chapter has been adapted from the following publication: arXiv:2103.03669. K. Goodenough contributed by setting up and supervising the corresponding bachelor project, co-writing the manuscript and implementing the code.

5.1. INTRODUCTION

The distribution of entanglement in Chapter 4 was in most cases dependent on the ability to *distill* entanglement. That is, the capability to turn a number of entangled states into (usually) a smaller number of states which are more strongly entangled (see [47] for a review). However, the only type of distillation protocols considered there were based on DEJMPS protocols [42]. A natural question arises, can one do better than DEJMPS protocols, especially with experimentally relevant constraints?

Our goal in this chapter is to answer the above question. In particular, we aim to find protocols where Alice and Bob use a small number of entangled states [51, 142], and require only a single round of communication after performing their local operations [142]. The above class of distillation protocols were first considered in [142], where they were called measure and exchange protocols. The semidefinite programming bounds found by Rozpedek et al. [142] allow to bound the optimal performance of measure and exchange protocols. Moreover, in some particular cases the existing protocols meet the bounds allowing to establish their optimality. Regarding the design of protocols, a heuristic procedure called the seesaw method allows to improve existing protocols [142]. More recently, Krastanov et al. investigated a genetic optimisation method for a subset of these protocols [87] and evaluated them including noisy operations.

5

Here, complementary to previous work, we find a systematic procedure to obtain good measure and exchange protocols.

To this end, we narrow down our investigation from general measure and exchange protocols to a practically relevant subset of protocols and states. Namely, we consider the distillation of Bell-diagonal states, where we use arbitrary *bilocal Clifford circuits* and measure out all but one of the qubit pairs. The measurement results are communicated between Alice and Bob, and the protocol is deemed successful if all pairs had correlated outcomes. We call this class of protocols *bilocal Clifford protocols* for short. This class of protocols includes a number of relevant protocols considered before in the literature [11, 20, 42, 46, 48, 54, 87, 144, 169].

The restriction to bilocal Clifford protocols and Bell-diagonal states allows us to reduce the finding of all bilocal Clifford protocols to enumerating all (double) cosets $\mathscr{D}_n \backslash \mathrm{Sp}(2n,\mathbb{F}_2)/\mathscr{K}_n$. Here, $\mathrm{Sp}(2n,\mathbb{F}_2)$ is the symplectic group over the field with two elements $\mathbb{F}_2$, $\mathscr{K}_n$ is the (possibly trivial) subgroup that preserves the input states and $\mathscr{D}_n$ is the distillation subgroup, which is the set of operations that leave both the success probability and fidelity invariant. One of our contributions in this work is to characterise this subgroup in terms of its generators and its order. We consider two cases for the input states - general input states (i.e. trivial symmetry group) and the n-fold tensor product of Werner states. For general input states, we find all protocols for up to $n = 5$ entangled pairs. For an n-fold tensor product of Werner states, we describe an algorithm that finds a complete set of double coset representatives. This allows us to optimise over all bilocal Clifford protocols when distilling an n-fold tensor product of a Werner state for n up to 8 pairs.

We find that for $n = 2,3$ copies of a Werner state, the highest fidelity out of all bilocal Clifford protocols is achieved by protocols studied before in the literature. For $n = 4$ to 8, we find increased fidelities over previously considered distillation schemes. Furthermore, we find explicit circuits achieving the highest fidelity out of all bilocal Clifford protocols, see Appendix 5.8.5. These circuits have comparable depth and number of two-qubit gates as previously studied protocols, highlighting also the practical feasibility of our findings.

This chapter is structured as follows. In Section 5.2 we describe the preliminaries and notation needed throughout the chapter. Section 5.3 explains bilocal Clifford distillation protocols and how the optimisation over such protocols can be rephrased as an optimisation over elements from the symplectic group $\mathrm{Sp}(2n,\mathbb{F}_2)$. In Section 5.4 we characterise the distillation subgroup $\mathscr{D}_n$. In Section 5.5 we prove a further reduction of our search space when the state to be distilled is an n-fold tensor product of a Werner state. In Section 5.6 we present our optimisation results. We

end with conclusions and discussions in Section 5.7.

5.2. PRELIMINARIES

We begin by setting some relevant notation. The field with two elements is denoted by $\mathbb{F}_2$. We use the notation U_i to denote a single-qubit operation on qubit i. The single-qubit operations that we use are the Pauli gates (I, X, Y and Z), the Hadamard gate (H) and the phase gate (S). Moreover, we denote by CNOT_{ij} a controlled-NOT operation with control qubit i and target qubit j, by CZ_{ij} a controlled-Z operation between qubits i and j and by SWAP_{ij} the operation that swaps qubits i and j.

5.2.1. PAULI GROUP AND CLIFFORD GROUP

The Pauli matrices are defined as

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \\ Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \qquad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{5.1}$$

The Pauli group with phases on n qubits $\overline{\mathscr{P}}_n$ consists of all $2^n \times 2^n$ matrices of the form $\lambda P_1 \otimes \cdots \otimes P_n$ with $\lambda \in \{\pm 1, \pm i\}$ and $P_i \in \{I, X, Y, Z\}$ for all $i \in \{1, \ldots, n\}$, together with standard matrix multiplication. Of particular interest to us is the Pauli group without any phase factors, $\mathscr{P}_n \cong \overline{\mathscr{P}}_n / \langle iI^{\otimes n} \rangle$. Here $\langle iI^{\otimes n} \rangle$ is the subgroup generated by $iI^{\otimes n}$. We will call this the Pauli group for short. An element of the group $\mathscr{P}_n$ is referred to as a Pauli string (of length n). The order of $\mathscr{P}_n$ equals $|\mathscr{P}_n| = 4^n$.

An important class of gates in quantum information theory are the so-called Clifford gates [60]. Circuits composed of Clifford gates are efficiently classically simulable, yet can be used to create complex quantum states, which are used for example in stabiliser error correction. The Clifford gates on n qubits form a group $\mathscr{C}_n$, and each $C \in \mathscr{C}_n$ induces an automorphism $f : \overline{\mathscr{P}}_n \to \overline{\mathscr{P}}_n$ on $\overline{\mathscr{P}}_n$ by conjugating each element with C, i.e. $f(P) = CPC^\dagger$. The Clifford group $\mathscr{C}_n$ is generated by Hadamard- (H_i) and phase (S_i) gates on each qubit ($1 \leq i \leq n$) and CNOT gates between every pair (i, j) of qubits. In matrix representation, these gates are given by

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \qquad S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \tag{5.2}$$

$$\mathrm{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{5.3}$$

5.2.2. BINARY REPRESENTATION OF THE PAULI AND CLIFFORD GROUP

The elements of the Pauli group and the Clifford group can be described in terms of binary vectors and matrices, respectively. To see this, we first introduce the following notation for the Pauli matrices.

$$\tau_{00} = I, \quad \tau_{10} = X, \quad \tau_{11} = iY, \quad \tau_{01} = Z. \tag{5.4}$$

We extend this notation to tensor products of Pauli matrices as follows.

$$\tau_a := \tau_{v_1 w_1} \otimes \cdots \otimes \tau_{v_n w_n}, \quad a = \begin{bmatrix} v \\ w \end{bmatrix}, \ v, w \in \mathbb{F}_2^n, \tag{5.5}$$

where $v = \begin{bmatrix} v_1 & v_2 & \dots & v_n \end{bmatrix}^\mathsf{T}$ and $w = \begin{bmatrix} w_1 & w_2 & \dots & w_n \end{bmatrix}^\mathsf{T}$. As mentioned in Section 5.2.1, the global phase factors are not important in the context of this chapter, so an element $\lambda\tau_a$, $\lambda \in \{\pm 1, \pm i\}$, of $\overline{\mathscr{P}}_n$ can be represented by the binary vector $a \in \mathbb{F}_2^{2n}$. The multiplication of the elements of $\overline{\mathscr{P}}_n$ corresponds then to the addition of the binary vectors.

For any $C \in \mathscr{C}_n$, the conjugation map f corresponds to a linear map on the set of binary vectors (and thus on $\overline{\mathscr{P}}_n$). The map f is an automorphism, and thus preserves the commutation relations of the elements of $\overline{\mathscr{P}}_n$. To see what this implies for the linear transformation in the binary picture, let $a, b \in \mathbb{F}_2^{2n}$. Then

$$\tau_a \tau_b = (-1)^{b^\mathsf{T}\Omega a} \tau_b \tau_a, \tag{5.6}$$

where $\Omega = \begin{bmatrix} 0 & I_n \\ I_n & 0 \end{bmatrix}$. A proof of this formula can be found in Appendix 5.8.1.

Let M denote the linear transformation corresponding to conjugation by C. It follows from Equation 5.6 that

$$\tau_{Ma} \tau_{Mb} = (-1)^{(Mb)^\mathsf{T}\Omega Ma} \tau_{Mb} \tau_{Ma}. \tag{5.7}$$

By Equation 5.6, we know that τ_a and τ_b commute iff $b^\mathsf{T}\Omega a = 0$ and anti-commute iff $b^\mathsf{T}\Omega a = 1$. In order to preserve the commutation relations, it must then hold that $b^\mathsf{T} M^\mathsf{T}\Omega M a = b^\mathsf{T}\Omega a$ for all $a, b \in \mathbb{F}_2^{2n}$, so $M^\mathsf{T}\Omega M = \Omega$. The matrices M that satisfy this condition thus preserve the so-called symplectic inner product $\omega(a,b) \equiv a^\mathsf{T}\Omega b$ between any two $a, b \in \mathbb{F}_2^{2n}$. These matrices form a group known as the symplectic group over $\mathbb{F}_2$, denoted by $\mathrm{Sp}(2n, \mathbb{F}_2)$. The order of the symplectic group over $\mathbb{F}_2$ is well-known [3] to be equal to

$$|\mathrm{Sp}(2n, \mathbb{F}_2)| = 2^{n^2} \prod_{j=1}^{n} (4^j - 1). \tag{5.8}$$

The symplectic complement of a subspace V of $\mathbb{F}_2^{2n}$ is defined as the set of elements of $\mathbb{F}_2^{2n}$ that have zero symplectic inner product with all elements from V,

$$V^\perp = \{a \in \mathbb{F}_2^{2n} \mid \omega(a,b) = 0 \ \forall b \in V\}. \tag{5.9}$$

The symplectic complement satisfies the following property,

$$\left(V^\perp\right)^\perp = V. \tag{5.10}$$

Calculations involving a symplectic matrix M can often be simplified by writing it as a block matrix $M = \begin{bmatrix} A & B \\ C & D \end{bmatrix}$, with $A, B, C, D \in M_{n\times n}(\mathbb{F}_2)$. From the condition $M^\mathsf{T}\Omega M = \Omega$ it follows that the blocks satisfy

$$\begin{aligned} B^\mathsf{T} D + D^\mathsf{T} B &= 0, \\ A^\mathsf{T} C + C^\mathsf{T} A &= 0, \\ A^\mathsf{T} D + C^\mathsf{T} B &= I_n. \end{aligned} \tag{5.11}$$

Moreover, the inverse of M is given by

$$M^{-1} = \begin{bmatrix} D^\mathsf{T} & B^\mathsf{T} \\ C^\mathsf{T} & A^\mathsf{T} \end{bmatrix}. \tag{5.12}$$

Let $\phi : \mathscr{C}_n \to \mathrm{Sp}(2n, \mathbb{F}_2)$ be the function that maps every Clifford gate to the corresponding symplectic matrix. This map is a surjective group homomorphism [38]. The symplectic group

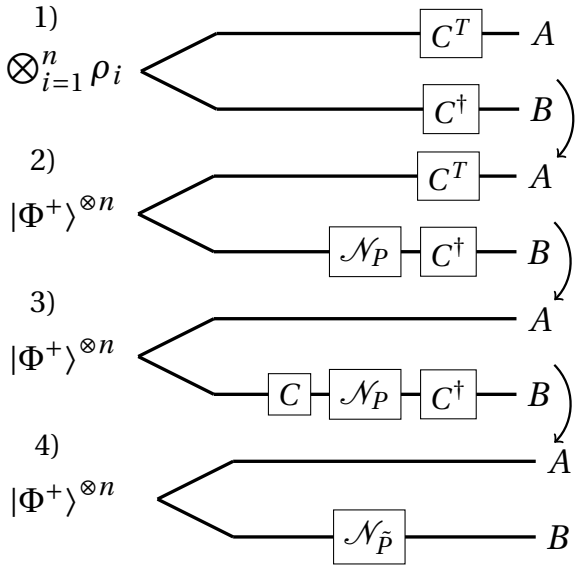

Figure 5.1: Schematic description of how bilocal Clifford circuits map n-qubit bipartite systems to n-qubit bipartite systems. From 1) to 2), we rewrite the state as $\otimes_{i=1}^{n}\rho_i = (I\otimes\mathcal{N})\left(|\Phi^+\rangle^{\otimes n}\right)$, where $\mathcal{N}(\cdot) = \sum_{P\in\mathscr{P}_n} p_P P(\cdot)P^\dagger$. In 3), we use the fact that $A^\mathsf{T}\otimes I|\Phi^+\rangle^{\otimes n} = I\otimes A|\Phi^+\rangle^{\otimes n}$ for any $2^n\times 2^n$ matrix A [175]. For 4), we use the fact the Cliffords act on the group of Pauli strings $\mathscr{P}_n$.

$\mathrm{Sp}(2n,\mathbb{F}_2)$ is thus generated by the images of a generating set of the Clifford group $\mathscr{C}_n$ under ϕ. The symplectic forms of the Hadamard, phase and CNOT gates, which generate the Clifford group, can be calculated using the following lemma [103].

Lemma 5.2.1. *For the following Clifford gates C, multiplying any $M\in\mathrm{Sp}(2n,\mathbb{F}_2)$ from the right by $\phi(C)$ has the following effect on M:*

$C = H_i$ *Swapping columns i and $n+i$*

$C = S_i$ *Adding column $n+i$ to column i*

$C = \mathrm{CNOT}_{ij}$ *Adding column j to column i and adding column $n+i$ to column $n+j$*

Similarly, multiplication from the left has the following effect on M:

$C = H_i$ *Swapping rows i and $n+i$*

$C = S_i$ *Adding row i to row $n+i$*

$C = \mathrm{CNOT}_{ij}$ *Adding row i to row j and adding row $n+j$ to row $n+i$*

5.3. BILOCAL CLIFFORD PROTOCOLS

This section covers the structure of the distillation protocols that are considered in this chapter. We consider a system consisting of two parties, Alice and Bob, that share n entangled two-qubit states. We focus on states that are diagonal in the Bell basis. Bell-diagonal states naturally arise with realistic noise models such as dephasing and depolarising. Moreover, any bipartite state can

be *twirled* into a Bell-diagonal state while preserving the fidelity [12]. Bell-diagonal states can be written as

$$\begin{aligned}\rho = p_I |\Phi^+\rangle\langle\Phi^+| + p_X |\Psi^+\rangle\langle\Psi^+| \\ + p_Y |\Psi^-\rangle\langle\Psi^-| + p_Z |\Phi^-\rangle\langle\Phi^-|.\end{aligned} \tag{5.13}$$

The indices of the probabilities arise from the following correspondence between the Bell states and the Pauli matrices.

$$\begin{aligned}|\Phi^+\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) = (I\otimes I)|\Phi^+\rangle, \\ |\Psi^+\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) = (I\otimes X)|\Phi^+\rangle, \\ |\Psi^-\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle) = (I\otimes iY)|\Phi^+\rangle, \\ |\Phi^-\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle) = (I\otimes Z)|\Phi^+\rangle.\end{aligned} \tag{5.14}$$

Equation (5.14) gives rise to a bijective mapping from the Bell states $|\Phi^+\rangle$, $|\Psi^+\rangle$, $|\Psi^-\rangle$ and $|\Phi^-\rangle$ to the Pauli matrices I, X, Y and Z, respectively. We denote a tensor product of n Bell-diagonal states by a tensor product of Pauli matrices, e.g. $|\Phi^+\rangle\otimes|\Psi^+\rangle\otimes|\Psi^-\rangle\otimes|\Phi^-\rangle$ is denoted by $I\otimes X\otimes Y\otimes Z$.

We generalise the notation of equation (5.13) and denote by p_P the probability that the system is in the state described by $P\in\mathscr{P}_n$. In the subscript we will not explicitly denote the tensor product, e.g. p_{XY} denotes the probability that the system is described by $X\otimes Y$. The initial state of the protocol consisting of n entangled two-qubit states can thus be fully described by the set of probabilities $\mathscr{Q} = \{p_{P_1P_2\ldots P_n} : P_i\in\{I,X,Y,Z\}\}$. We refer to such a system as an *n-qubit bipartite system.*

5.3.1. BILOCAL CLIFFORD CIRCUITS

The first step of the protocol is to apply the bilocal Clifford operations. That is, if Alice applies a Clifford operation $\tilde{C}\in\mathscr{C}_n$ to her qubits, then Bob applies $\tilde{C}^*$, the entry-wise complex conjugate of $\tilde{C}$, to his qubits (see Fig. 5.1). This leads to a permutation of the set $\mathscr{Q}$. In particular, each element p_P of $\mathscr{Q}$ is mapped [175] to $p_{\tilde{C}^\mathsf{T}P\tilde{C}^*}$, or equivalently, $p_{CPC^\dagger}$, where we defined $C=\tilde{C}^\mathsf{T}\in\mathscr{C}_n$. We denote the probabilities that describe the permuted state by $\tilde{p}_{P_1P_2\ldots P_n}$.

We note here that the most general permutation on $\mathscr{Q}$ by local unitaries consists of applying bilocal Cliffords followed by a Pauli string applied to either Alice or Bob's side [38]. These Pauli strings can be used to reorder locally the coefficients of the states.

Since (bilocal) Clifford operations form a group, the Clifford group has a group action on $\mathscr{Q}$. The (normal) subgroup of the Clifford group that fixes $\mathscr{Q}$ point-wise does not change any of the statistics, and is thus not of interest to us. This subgroup consists of all Pauli strings, and quotienting out the Cliffords by this subgroup leads to the symplectic group (over $\mathbb{F}_2$), $\mathrm{Sp}(2n,\mathbb{F}_2)$ [37, 38]. We can thus describe a bilocal Clifford operation by an element $M\in\mathrm{Sp}(2n,\mathbb{F}_2)$. To simplify notation, we sometimes slightly abuse the notation and denote by $C\in\mathrm{Sp}(2n,\mathbb{F}_2)$ the symplectic matrix corresponding to conjugation by $C\in\mathscr{C}_n$, but it should be kept in mind that always the symplectic matrix M is meant.

5.3.2. MEASUREMENTS AND POSTSELECTION

In the second step, Alice and Bob perform measurements in the computational basis on $n-1$ of their qubits. Alice and Bob report their results to each other using classical communication. If the outcomes are equal, they keep the state that was not measured. In this case, the protocol is

called *successful*. If the outcomes are not equal, they discard all states, and the protocol is not successful. The prototypical example of such a protocol is the *DEJMPS protocol* [47]. In the DEJMPS protocol, two bilocal single-qubit Clifford rotations are performed on two Bell-diagonal states, after which a bilocal CNOT is performed between the two pairs. Then, both qubits of one of the pairs is measured in the computational basis, after which they are post-selected on classically correlated outcomes. The probability that a protocol is successful is equal to the probability that all measured states are either in the $|\Phi^+\rangle$ or in the $|\Phi^-\rangle$ state, which correspond to the I and Z Pauli matrix, respectively. This can be seen by the fact that these two states yield correlated outcomes when both sides are measured in the computational basis. The success probability of the protocol is thus equal to

$$p_{\text{suc}} = \sum_{\substack{P_1\in\{I,X,Y,Z\},\\ Q_j\in\{I,Z\}}} \tilde{p}_{P_1 Q_2 \ldots Q_n}, \tag{5.15}$$

where we used the convention that the first two-qubit state is not measured. Moreover, the fidelity of the remaining state and the $|\Phi^+\rangle$ state is equal to

$$F_{\text{out}} = \frac{\sum_{Q_j\in\{I,Z\}} \tilde{p}_{I_1 Q_2 \ldots Q_n}}{p_{\text{suc}}}. \tag{5.16}$$

To simplify notation in the rest of this chapter, we introduce the following two definitions.

Definition 5.3.1. The *base* of an n-qubit bipartite quantum system is given by

$$\mathscr{B} = \{I_1 \otimes Q_2 \otimes \ldots \otimes Q_n \in \mathscr{P}_n : Q_j \in \{I, Z\} \\ \forall\, j \in \{2, \ldots, n\}\}.$$

In the binary representation, the elements of the base are the vectors b satisfying equation (5.17).

$$b = \begin{bmatrix} v \\ w \end{bmatrix}, \quad v, w \in \mathbb{F}_2^n,\ v = 0, w_1 = 0. \tag{5.17}$$

Definition 5.3.2. The *pillars* of an n-qubit bipartite quantum system are given by

$$\mathscr{P} = \{P_1 \otimes Q_2 \otimes \ldots \otimes Q_n \in \mathscr{P}_n : P_1 \in \{I, X, Y, Z\}, \\ Q_j \in \{I, Z\}\ \forall\, j \in \{2, \ldots, n\}\}.$$

The elements of the pillars are represented by the binary vectors p satisfying equation (5.18).

$$p = \begin{bmatrix} v \\ w \end{bmatrix}, \quad v, w \in \mathbb{F}_2^n,\ v_i = 0\ \forall\, i \in \{2, \ldots, n\}. \tag{5.18}$$

The naming of the base and pillars is made clear when the probabilities p_P are ordered in an n-dimensional hypercube, where each dimension corresponds to a qubit pair, see Fig. 5.2.

Using these definitions, equation (5.15) can be rewritten as

$$p_{\text{suc}} = \sum_{P\in\mathscr{P}} \tilde{p}_P, \tag{5.19}$$

and equation (5.16) as

$$F = \frac{\sum_{P\in\mathscr{B}} \tilde{p}_P}{\sum_{P\in\mathscr{P}} \tilde{p}_P}. \tag{5.20}$$

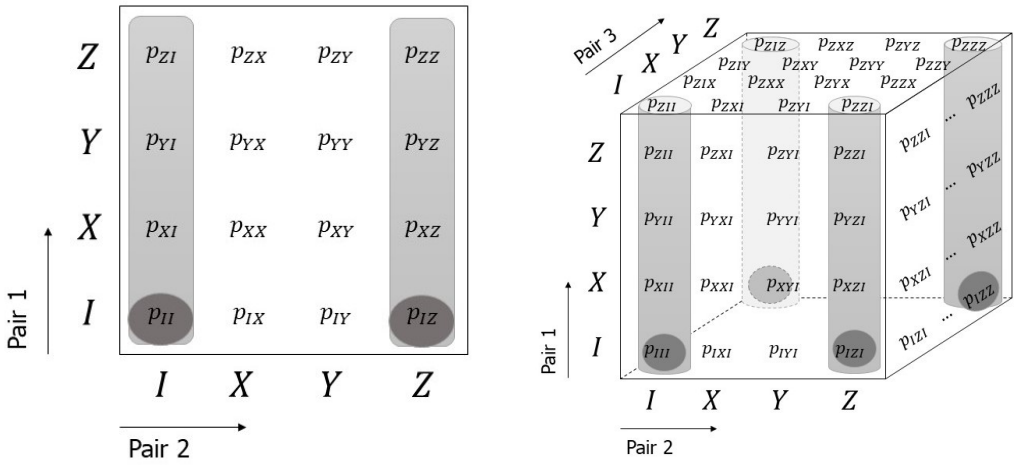

Figure 5.2: Probabilities that describe the state of a 2-pair system (left) and a 3-pair system (right). The light grey rectangles/cylinders highlight the probabilities that correspond to the pillars. The darker circles highlight the probabilities that correspond to the base. For the 3-pair system we have labelled here only the coefficients that are on the front, right and top face.

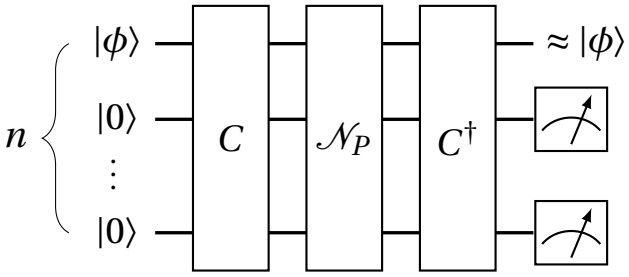

Figure 5.3: Equivalence between bilocal Clifford protocols and a subset of error detection schemes. This circuit detects as errors the set of Pauli strings that do not end up in the pillars after applying the Clifford circuit C.

The fidelity and success probability are referred to as the *distillation statistics*.

In the binary picture, the distillation statistics can be calculated using the inverse of the symplectic matrix, which can be efficiently calculated using 5.12. Let M be the symplectic matrix corresponding to a permutation $P \mapsto CPC^\dagger$, $C \in \mathscr{C}_n$. We wish to determine which binary vectors are mapped to the vectors corresponding to the base and the pillars by M. Since M permutes the binary vectors, this is equivalent to determining where the base and pillar vectors are mapped to by M^{-1}.

Finally, there is a direct analogy between our optimisation over bilocal Clifford protocols, and quantum error detection schemes of the form shown in Fig. 5.3. Such schemes will detect as errors the set of Pauli strings that do not end up in the pillars after applying the Clifford circuit C. We will not pursue this further in this thesis, however.

5.4. PRESERVATION OF DISTILLATION STATISTICS

In many relevant cases, the distillation statistics from equations (5.19) and (5.20) are the relevant parameters for quantifying an entanglement distillation protocol. Furthermore, there exist non-identical bilocal Clifford circuits which result in the same distillation statistics. To find all bilocal Clifford protocols, it is thus sufficient to find a representative bilocal Clifford protocol for each unique tuple of distillation statistics. In this section we characterise these representatives for general input states.

First, we specify the set of Clifford operations that preserve the distillation statistics. We denote this set by $\mathscr{D}_n$. Now observe that $\mathscr{D}_n$ is a subgroup of $\mathrm{Sp}(2n,\mathbb{F}_2)$. Moreover, let $M \in \mathrm{Sp}(2n,\mathbb{F}_2)$ and consider the corresponding distillation protocol. We can freely add or remove elements from $\mathscr{D}_n$ at the end of this protocol, without changing the fidelity and the success probability. That is, all elements in the right coset $\mathscr{D}_n M = \{DM : D \in \mathscr{D}_n\}$ yield the same distillation statistics. Instead of optimising over all possible Clifford circuits it thus suffices to optimise over the right cosets of $\mathscr{D}_n$ in $\mathrm{Sp}(2n,\mathbb{F}_2)$.

5.4.1. RELATING THE PRESERVATION OF THE BASE AND PILLARS

In this section we explain the relation between the base and the pillars, which were introduced in definitions 5.3.1 and 5.3.2, respectively, and the distillation subgroup $\mathscr{D}_n$.

From equations (5.19) and (5.20) it can be observed that for a general initial state, the operations that preserve the distillation statistics are precisely those operations that leave simultaneously both the base and pillars invariant. In the following lemma it is proven that invariance of the base implies invariance of the pillars, and vice versa.

Lemma 5.4.1. *Let $\mathscr{Q}$ be an n-qubit bipartite quantum system with base $\mathscr{B} \subseteq \mathbb{F}_2^{2n}$ and pillars $\mathscr{P} \subseteq \mathbb{F}_2^{2n}$. Let $\pi : \mathbb{F}_2^{2n} \to \mathbb{F}_2^{2n}$, $\pi(a) = Ma$, with $M \in \mathrm{Sp}(2n,\mathbb{F}_2)$. Then $\pi[\mathscr{B}] = \mathscr{B} \iff \pi[\mathscr{P}] = \mathscr{P}$.*

Proof. We first prove $\pi[\mathscr{B}] = \mathscr{B} \implies \pi[\mathscr{P}] = \mathscr{P}$. For this, we first show that the pillars form the symplectic complement of the base, i.e. $\mathscr{B}^\perp = \mathscr{P}$ (see equation (5.9)). Recall from equation (5.17) that $b \in \mathscr{B}$ can be written as $b = \begin{bmatrix} v^b & w^b \end{bmatrix}^\mathsf{T}$ with $v^b = 0$ and $w_1^b = 0$. Note that $\mathscr{B}$ is a subspace of $\mathbb{F}_2^{2n}$. The symplectic inner product between b and $a \in \mathbb{F}_2^{2n}$, is equal to $v^b \cdot w^a + v^a \cdot w^b$, where $\cdot$ indicates the standard vector dot product. This is equal to zero for all $b \in \mathscr{B}$ if and only if $v_i^a = 0$ for all $i \in \{2,\ldots,n\}$, so iff $v \in \mathscr{P}$.

Let $b \in \mathscr{B}$ and $p \in \mathscr{P}$. Then $\omega(b,p) = 0$, and since $M \in \mathrm{Sp}(2n,\mathbb{F}_2^{2n})$, we have that $\omega\big(\pi(b),\pi(p)\big) = 0$ as well. Since by assumption $\pi(b) \in \mathscr{B}$, it follows that $\pi(p) \in \mathscr{P}$. Finally, since π is an automorphism, we know that it is bijective and thus $\pi[\mathscr{P}] = \mathscr{P}$.

For the other direction, we use the fact that $\mathscr{P}^\perp = \mathscr{B}$, see equation (5.10). Then, the above argument can be repeated with $\mathscr{B}$ and $\mathscr{P}$ interchanged to conclude that $\pi[\mathscr{B}] = \mathscr{B} \iff \pi[\mathscr{P}] = \mathscr{P}$. □

From Lemma 5.4.1 we conclude that the operations that preserve the distillation statistics for arbitrary input states are precisely the operations that leave the base invariant. We use this observation to characterise the subgroup $\mathscr{D}_n$ that preserves the distillation statistics. In the trivial case that $n = 1$, we have $\mathscr{D}_1 = \mathrm{Sp}(2,\mathbb{F}_2)$. In this case, the only base element is the identity I, which is always mapped to itself under an automorphism. For all $n > 1$, however, $\mathscr{D}_n$ is a proper subgroup of $\mathrm{Sp}(2n,\mathbb{F}_2)$. Consider for instance the Hadamard gate on the second qubit, which is an element of $\mathrm{Sp}(2n,\mathbb{F}_2)$. This gate induces the swap of X_2 and Z_2 and hereby changes the base.

5.4.2. GENERATORS OF SUBGROUP PRESERVING DISTILLATION STATISTICS

The goal of this section is to characterise the distillation subgroup $\mathscr{D}_n$. In particular, we find the distillation subgroup in terms of a generating set T_n.

Lemma 5.4.2. *The distillation subgroup is generated by the set T_n, i.e. $\mathscr{D}_n = \langle T_n \rangle$, where*

$$T_n = \{H_1, S_1, \ldots, S_n\} \cup \left\{\text{CNOT}_{ij} \mid 1 \leq j < i \leq n\right\} \cup \left\{\text{CNOT}_{ij} \mid 2 \leq i < j \leq n\right\}. \tag{5.21}$$

Proof. By inspection, each element of T_n preserves the base, so by Lemma 5.4.1 we have that $\langle T_n \rangle \subseteq \mathscr{D}_n$. For the other inclusion, let $M = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \in \mathscr{D}_n$. We show that we can reduce such an arbitrary M to the identity matrix by left-multiplication by elements in $\langle T_n \rangle$, see Lemma 5.2.1. First, note that if $M \in \mathscr{D}_n$, then by definition $M[\mathscr{B}] \subseteq \mathscr{B}$ and $M[\mathscr{P}] \subseteq \mathscr{P}$. In the binary picture this implies that

$$B_{ij} = 0 \text{ if } (i,j) \neq (1,1), \qquad D_{12} = \cdots = D_{1n} = 0.$$

Since M has full rank, we cannot have $B_{11} = D_{11} = 0$. Hence, by multiplying M from the left by I, H_1 or $H_1 S_1$, we may assume that $B_{11} = 0$ (such that $B = 0$) and $D_{11} = 1$.

That M has full rank implies that the last n columns of M are linearly independent. By using CNOT gates from $\left\{\text{CNOT}_{ij} \mid 1 \leq j < i \leq n\right\} \subseteq T_n$ and $\left\{\text{CNOT}_{ij} \mid 2 \leq i < j \leq n\right\} \subseteq T_n$, the D submatrix can be reduced to the identity matrix.

Since $D = I$ and $B = 0$, it follows from 5.11 that $A = I$ and C is symmetric. For $1 \leq j < i$ denote $S_{ij} := (S_j \text{CNOT}_{ij})^2 \in \langle T_n \rangle$. Left-multiplication by S_{ij} corresponds to adding row i to row $n+j$ and adding rows i and j to row $n+i$. Note that this preserves the fact that $A = D = I$ and $B = 0$.

For $j = 1, \ldots, n-1$ (in this order), we can multiply M from the left by elements from $\{S_j\} \cup \{S_{ij} \mid i > j\} \subseteq \langle T_n \rangle$ to ensure that $C_{ji} = 0$ for all $1 \leq j \leq i \leq n$. This implies that C is strictly lower triangular. But if C is strictly lower triangular and symmetric, $C = 0$. This implies that $M = I$. □

5.4.3. ORDER OF THE DISTILLATION SUBGROUP

As noted before, for general input states it is sufficient to only consider the right cosets of $\mathscr{D}_n$ in $\text{Sp}(2n, \mathbb{F}_2)$. To see how much looking at cosets of $\mathscr{D}_n$ in $\text{Sp}(2n, \mathbb{F}_2)$ limits the search space of protocols, in this section a formula for the order of $\mathscr{D}_n$ is presented and proved. As mentioned earlier, in the trivial case that $n = 1$ we have $\mathscr{D}_1 = \text{Sp}(2, \mathbb{F}_2)$, and thus $|\mathscr{D}_1| = |\text{Sp}(2, \mathbb{F}_2)| = 6$. For $n \geq 2$ the order of $\mathscr{D}_n$ is given in Theorem 5.4.3.

Theorem 5.4.3. *For an n-to-1 distillation protocol, with $n > 1$, the order of $\mathscr{D}_n$ is given by*

$$|\mathscr{D}_n| = 6 \cdot 2^{n^2-1} \prod_{j=1}^{n-1} (2^j - 1).$$

Proof. First note that $D \in \mathscr{D}_n$ is fully determined by how it maps each of the standard basis vectors $\{e_i : i \in \{1, \ldots, 2n\}\}$ of $\mathbb{F}_2^{2n}$. We count how many transformations of the standard basis vectors are possible.

Let us start by looking at e_{2n}. This is a base element, thus it must again be transformed to a base element, because D preserves the distillation statistics. There are 2^{n-1} base elements, but the identity element, the zero vector, is always mapped to itself by D. Thus there are $2^{n-1} - 1$ possibilities for the transformation of e_{2n}. That all transformations are indeed possible, is proved by giving a construction. Suppose that e_{2n} is mapped to a base element $b \equiv De_{2n} \in \mathscr{B}$. We show that b can be transformed to e_{2n} through left multiplication by elements of $\mathscr{D}_n$. The transformation

from e_{2n} to b can then be obtained by taking the product of the inverses of these generators in reverse order. For this, recall that the action of left multiplication by Hadamard, phase and CNOT gates was given in Lemma 5.2.1.

Note that $b_1, ..., b_{n+1} = 0$. The vector b can be transformed to e_{2n} by taking the following steps.

1. If $b_{2n} = 0$, apply a $\text{CNOT}_{ni} \in \mathscr{D}_n$ gate with i chosen such that $b_{n+i} = 1$. Note that there always is a i such that this is possible, because otherwise b is the zero vector, which corresponds to the identity element $I^{\otimes n}$.
2. For all $i \in \{2, ..., n\}$ with $b_{n+i} = 1$, apply a $\text{CNOT}_{in} \in \mathscr{D}_n$ gate.

Steps 1 and 2 are visually summarised below.

$$b = \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ \hline 0 \\ \cdot \\ \cdot \end{array}\right] \xrightarrow{1} \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ \hline 0 \\ \cdot \\ 1 \end{array}\right] \xrightarrow{2} \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ \hline 0 \\ 0 \\ 1 \end{array}\right] = e_{2n}$$

Given the transformation of e_{2n} by D, we now wish to determine the number of possible transformations for e_n. We know that left multiplication by D preserves the symplectic inner product. Hence, since $\omega(e_n, e_{2n}) = 1$, it must hold that $\omega(De_n, De_{2n}) = 1$. Observe that for every non-zero element $a \in \mathbb{F}_2^{2n}$, exactly for half of the elements of $\mathbb{F}_2^{2n}$ the symplectic inner product with a is equal to one[1]. Thus there are $\frac{|\mathbb{F}_2^{2n}|}{2} = \frac{4^n}{2} = 2^{2n-1}$ possibilities for the transformation of e_n.

We show that each of those transformations can indeed be achieved. Suppose that D has mapped e_n to a vector $c \equiv De_n \in \mathbb{F}_2^{2n}$. Because $\omega(De_n, De_{2n}) = 1$ and De_{2n} is a base vector, we know that there is at least one $i \in \{2, \ldots, n\}$ such that $c_i = 1$. Since we can always apply a $\text{CNOT}_{in} \in \mathscr{D}_n$ gate, which does not affect the vector e_{2n}, we can assume without loss of generality that $c_n = 1$. Now c can be transformed to e_n without affecting e_{2n} by taking the following steps.

3. For all i with $c_i = 1$ apply a $\text{CNOT}_{ni} \in \mathscr{D}_n$ gate.
4. For all $i \neq n$ with $c_{n+i} = 1$, apply a $\text{CZ}_{in} \in \mathscr{D}_n$ gate. This operation results in the addition of row i to row $2n$ and the addition of row n to row $n+i$. Note that this operation leaves the base invariant, so indeed $\text{CZ}_{in} \in \mathscr{D}_n$.
5. If $c_{2n} = 1$, apply the gate $S_n \in \mathscr{D}_n$ on qubit n.

Steps 3 to 5 are visually summarised below.

$$c = \left[\begin{array}{c} \cdot \\ \cdot \\ 1 \\ \hline \cdot \\ \cdot \\ \cdot \end{array}\right] \xrightarrow{3} \left[\begin{array}{c} 0 \\ 0 \\ 1 \\ \hline \cdot \\ \cdot \\ \cdot \end{array}\right] \xrightarrow{4} \left[\begin{array}{c} 0 \\ 0 \\ 1 \\ \hline 0 \\ 0 \\ \cdot \end{array}\right] \xrightarrow{5} \left[\begin{array}{c} 0 \\ 0 \\ 1 \\ \hline 0 \\ 0 \\ 0 \end{array}\right] = e_n$$

Thus indeed, given the transformation of e_{2n}, there are 2^{2n-1} possible transformations of e_n. Combining this with the number of transformations of e_{2n}, we find that there are $2^{2n-1}(2^{n-1} - 1)$ possible transformations for e_n and e_{2n} together.

[1] Let $a \in \mathbb{F}_2^{2n}$, such that $a_k = 1$. Then the vectors $a' \in \mathbb{F}_2^{2n}$ satisfying $\omega(a, a') = 1$ can be constructed by choosing $a'_j \in \{0, 1\}$ randomly for $j \in \{1, \ldots, 2n\} \setminus \{n + k\}$ and then choosing $a'_{n+k} \in \{0, 1\}$ such that $\omega(a, a') = 1$.

The elements of $\mathscr{D}_n$ that leave e_n and e_{2n} invariant form a subgroup that is isomorphic to $\mathscr{D}_{n-1}$, with the number of cosets in $\mathscr{D}_n$ equal to $2^{2n-1}(2^{n-1}-1)$. Thus

$$|\mathscr{D}_n| = 2^{2n-1}(2^{n-1}-1)|\mathscr{D}_{n-1}|.$$

By induction on n it follows that

$$\begin{aligned}|\mathscr{D}_n| &= |\mathscr{D}_1| \prod_{j=2}^{n} 2^{2j-1}(2^{j-1}-1) \\ &= 6\cdot 2^{\sum_{j=2}^{n}(2j-1)} \prod_{j=2}^{n}(2^{j-1}-1) \\ &= 6\cdot 2^{n^2-1} \prod_{j=1}^{n-1}(2^j-1).\end{aligned}$$

□

The following corollary is a direct consequence of Theorem 5.4.3.

Corollary 5.4.4. *The index of* $\mathscr{D}_n$ *in* $\mathrm{Sp}(2n,\mathbb{F}_2)$ *is given by*

$$\left[\mathrm{Sp}(2n,\mathbb{F}_2):\mathscr{D}_n\right] = \frac{1}{3}(2^n-1)\prod_{j=1}^{n}(2^j+1).$$

Proof. Recall that $|\mathrm{Sp}(2n,\mathbb{F}_2)| = 2^{n^2}\prod_{j=1}^{n}(4^j-1)$. As a result,

$$\begin{aligned}\left[\mathrm{Sp}(2n,\mathbb{F}_2):\mathscr{D}_n\right] &= \frac{|\mathrm{Sp}(2n,\mathbb{F}_2)|}{|\mathscr{D}_n|} \\ &= \frac{2^{n^2}\prod_{j=1}^{n}(4^j-1)}{6\cdot 2^{n^2-1}\prod_{j=1}^{n-1}(2^j-1)} \\ &= \frac{\prod_{j=1}^{n}(2^j-1)(2^j+1)}{3\prod_{j=1}^{n-1}(2^j-1)} \\ &= \frac{1}{3}(2^n-1)\prod_{j=1}^{n}(2^j+1).\end{aligned}$$

□

For comparison, we list the values of $|\mathrm{Sp}(2n,\mathbb{F}_2)|$ and $[\mathrm{Sp}(2n,\mathbb{F}_2):\mathscr{D}_n]$ in Table 5.1 for $n = 2,3,4,5$.

	2	3	4	5
$\lvert\mathrm{Sp}(2n,\mathbb{F}_2)\rvert$	720	1451520	47377612800	24815256521932800
$[\mathrm{Sp}(2n,\mathbb{F}_2):\mathscr{D}_n]$	15	315	11475	782595

Table 5.1: Values of $|\mathrm{Sp}(2n,\mathbb{F}_2)|$ and $[\mathrm{Sp}(2n,\mathbb{F}_2):\mathscr{D}_n]$ for $n = 2,3,4,5$.

5.4.4. FINDING A TRANSVERSAL

In this section, we briefly describe a way to find a transversal for the cosets of $\mathscr{D}_n$ in $\mathrm{Sp}(2n,\mathbb{F}_2)$. A transversal is a set that contains exactly one element for each of the cosets. Once this transversal is found, it can be applied to an arbitrary n-qubit input state to calculate all possible distillation statistics that can be achieved using bilocal Clifford circuits. From this set of distillation statistics, the optimal protocol based on any optimality criterion can be selected.

In order to find a transversal, random elements from the symplectic group $\mathrm{Sp}(2n,\mathbb{F}_2)$ are sampled. A sampled element is added to the set of representatives if the corresponding coset is not yet represented in this set. Recall that two elements belong to the same coset if they result in the same distillation statistics (for a general input state). This is the case if and only if the same Pauli strings are mapped to the base. More formally, consider an n-qubit pairs bipartite system with base $\mathscr{B}$ in the binary picture. Let $M_1, M_2 \in \mathrm{Sp}(2n,\mathbb{F}_2)$. Let $\mathscr{V}$ denote the set of binary vectors that are mapped to the base by M_1 and let $\mathscr{W}$ denote the set of binary vectors that are mapped to the base by M_2. Then M_1 and M_2 belong to the same coset if and only if $\mathscr{V} = \mathscr{W}$. Because M_1 and M_2 permute the binary Pauli vectors, this is equivalent to $M_1^{-1}[\mathscr{B}] = M_2^{-1}[\mathscr{B}]$.

The sampling is continued until the set of representatives has size $[\mathrm{Sp}(2n,\mathbb{F}_2) : \mathscr{D}_n]$. Note that finding a transversal in the way described in this section is equivalent to the coupon collector's problem. Hence, it has expected running time $\mathscr{O}([\mathrm{Sp}(2n,\mathbb{F}_2) : \mathscr{D}_n] \log[\mathrm{Sp}(2n,\mathbb{F}_2) : \mathscr{D}_n])$.

5

5.5. REDUCTION FOR n COPIES OF A WERNER STATE

Here we describe our reduction of the search space when the input state is an n-fold tensor product of Werner states. A Werner state has coefficients $p_I = F_{\text{in}}$, $p_X = p_Y = p_Z = \frac{1-F_{\text{in}}}{3}$, and its n-fold tensor product is highly symmetric — it is left invariant under any element of $\mathscr{K}_n$, where

$$\mathscr{K}_n = \langle \{\mathrm{SWAP}_{ij}\}_{1\le i<j\le n} \cup \{H_i\}_{i=1}^n \cup \{S_i\}_{i=1}^n \rangle.$$

We leverage this symmetry by noting that the distinct distillation protocols correspond to the double cosets $\mathscr{D}_n \backslash \mathrm{Sp}(n,\mathbb{F}_2)/\mathscr{K}_n$, similar to our argument before for right cosets for general input states. In this section, we describe how one can rewrite an arbitrary symplectic matrix M to another symplectic matrix M' of a specific form, which is in the same double coset as M. Such a representative M' of the double coset has a smaller number of free parameters, reducing the search space significantly.

Lemma 5.5.1. *Let $M \in \mathrm{Sp}(2n,\mathbb{F}_2)$. There exists M' in the double coset $\mathscr{D}_n M \mathscr{K}_n$ that is of the form*

$$M' = \begin{bmatrix} A & B \\ 0 & A^{\mathsf{T}} \end{bmatrix}, \; A = \left[\begin{array}{c|c} 1 & 0 \\ \hline a & I_{n-1} \end{array}\right], \; B = \left[\begin{array}{c|c} 0 & b^{\mathsf{T}} \\ \hline b & E + ba^{\mathsf{T}} \end{array}\right],$$

where $a, b \in \mathbb{F}_2^{n-1}$ and $E \in \mathbb{F}_2^{(n-1)\times(n-1)}$ is symmetric with zeroes on the diagonal.

Proof. Let $M' \in \mathscr{D}_n M \mathscr{K}_n$ be such that

$$M'_{ij} = \delta_{ij} \quad \text{for } (i,j) \in [n] \times \{2,\dots,k\} \tag{5.22}$$

with $1 \le k \le n$ as large as possible. Note that for $k = 1$ the condition is trivially fulfilled.

Claim: $k = n$.

Proof of claim. Suppose that $k < n$. Then $M'_{k+1,k+1} = 0$, otherwise we can use row operations on M' (left-multiplication by matrices $\mathrm{CNOT}_{k+1,i} \in \mathscr{D}_n$) to obtain $M'_{i,k+1} = \delta_{i,k+1}$ for all $i \in [n]$ while keeping (5.22), contradicting the maximality of k.

Note that the above condition $M'_{k+1,k+1} = 0$ needs to hold after applying operations to M' that preserve the form in equation (5.22). Thus, by permuting rows in $\{k+1,\dots,n\}$ (left-multiplication by matrices $\mathrm{SWAP}_{ij} \in \mathscr{D}_n$) or permuting columns in $\{1,k+1,\dots,n\}$ (right-multiplication by matrices $\mathrm{SWAP}_{ij} \in \mathscr{K}_n$) we deduce that $M'_{ij} = 0$ for $(i,j) \in \{k+1,\dots,n\} \times \{1,k+1,\dots,n\}$. Since we can swap column i and $n+i$ by multiplying from the right with $H_i \in \mathscr{K}_n$, we also have $M'_{ij} = 0$ for $(i,j) \in \{k+1,\dots,n\} \times \{n+1,n+k+1,\dots,2n\}$. To summarise, we have

$$\begin{aligned} M'_{ij} &= 0 \quad \text{for } i \in \{k+1,\dots,n\} \\ &\qquad \text{and } j \in [2n] \setminus \{n+2,\dots,n+k\} \\ M'_{ij} &= \delta_{ij} \quad \text{for } i \in \{2,\dots,k\} \text{ and } j \in \{2,\dots,k\}. \end{aligned}$$

Since rows $k+1,\dots,n$ must have zero symplectic inner product with rows $2,\dots,k$, it follows that rows $k+1,\dots,n$ must in fact be equal to zero. Since M' has full rank, this implies that $k = n$. ■

Consider the first row of M'. We have $M'_{1,j} = 0$ for $j = 2,\dots,n$. If $M'_{1,1} = M'_{1,n+1} = 0$, then the fact that this row has zero symplectic inner product with rows $2,\dots,n$ implies that the first row is equal to zero, which is not possible as M' has full rank. So by multiplying from the right by I, H_1, or S_1H_1 which are in $\mathscr{K}_n$, we may assume that $M'_{1,1} = 1$ and $M'_{1,n+1} = 0$.

Writing $M' = \left[\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right]$, we see that A and B have the following form:

$$A = \left[\begin{array}{c|c} 1 & 0 \\ \hline a & I_{n-1} \end{array}\right], \quad B = \left[\begin{array}{c|c} 0 & d^{\mathsf{T}} \\ \hline b & E' \end{array}\right].$$

Since row 1 has zero symplectic inner product with rows $2,\dots,n$, it follows that $d = b$. Note that for $1 \le i < j \le n-1$ the symplectic inner product of rows $i+1$ and $j+1$ is equal to $a_ib_j + E'_{ji} + a_jb_i + E'_{ij}$. Since this inner product is zero, the matrix $E := E' + ba^{\mathsf{T}}$ is symmetric. By multiplying from the right by $H_iS_iH_i \in \mathscr{K}_n$ $(i = 2,\dots,n)$ if necessary, we may set the diagonal elements of E' such that the diagonal elements of E are zero.

Recall from Lemma 5.4.2 that $S_{ij} := (S_j\mathrm{CNOT}_{ij})^2 \in \mathscr{D}_n$ for $1 \le j < i$. Recall furthermore that left-multiplication of M' by S_{ij} amounts to adding row i to row $n+j$ and adding rows i and j to row $n+i$. By left-multiplication by elements $S_{ij} \in \mathscr{D}_n$ and $S_i \in \mathscr{D}_n$ we may (without changing the first n rows of M') assume that C is a strictly upper triangular matrix. Since the first n columns of M' must have pairwise zero symplectic inner product, this implies that in fact $C = 0$. Since $A^{\mathsf{T}}D + C^{\mathsf{T}}B = I_n$, it follows that $D = (A^{\mathsf{T}})^{-1} = A^{\mathsf{T}}$, where we have used that A is self-inverse.

□

Note that for any permutation $\pi \in S_{n-1}$, we can replace a,b,E by $\pi(a),\pi(b),\pi(E)$ (permuting both rows and columns) by multiplying M' simultaneously from the left and the right by elements SWAP_{ij}, since SWAP_{ij} is an element of $\mathscr{D}_n$ and $\mathscr{K}_n$ for $2 \le i < j \ne n$. Also, we can replace (a,b) by (b,a) or $(a,a+b)$ by multiplication from the left and right by elements from $\{S_1,H_1\}$. Hence, to cover all cases, it suffices to enumerate over the triples (a,b,E) where $a \le b \le a+b$ and E runs over the adjacency matrices of graphs on $n-1$ nodes (up to isomorphism).

5.6. OPTIMISATION RESULTS

In the previous sections we have outlined our methods for finding all possible bilocal Clifford protocols, which were described in Section 5.3. In the following we report our findings, first for up to $n = 5$ general Bell-diagonal input states, second for up to $n = 8$ identical Werner states.

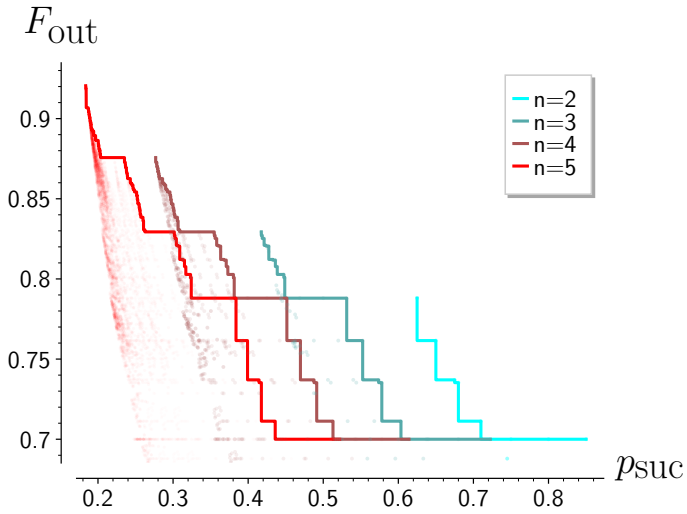

Figure 5.4: Achievable $(p_{\text{suc}}, F_{\text{out}})$ pairs for $n = 2$ to 5 copies of a state with coefficients $p_I = 0.7$, $p_X = 0.15$, $p_Y = 0.10$, $p_Z = 0.05$. The highest achievable pairs are indicated by a solid line for each number of copies. Not included in the plot are those distillation protocols with fidelity smaller than $F = 0.68$.

5

5.6.1. ACHIEVED DISTILLATION STATISTICS FOR GENERAL INPUT STATES

In Fig. 5.4 we show the achievable $(p_{\text{suc}}, F_{\text{out}})$ pairs for $n = 2,3,4,5$ copies of a state with coefficients $p_I = 0.7$, $p_X = 0.15$, $p_Y = 0.10$, $p_Z = 0.05$. We also plot the envelope, indicating the best performing schemes. Moreover, our results for $n = 5$ clearly show that picking an arbitrary coset does not give a good protocol in general.

Furthermore, we consider the $n = 2$ scenario where the two input states are equal, i.e. both have equal values for the p_I, p_X, p_Y, p_Z parameters. By comparing all analytic expressions of the output fidelity as a function of p_I, p_X, p_Y, p_Z, we find that the DEJMPS protocol achieves the highest fidelity out of all bilocal Clifford protocols (see [76] for the details).

While we do not explore this direction, the results can also be applied to less symmetric cases, i.e. when the n pairs are not the same. This situation is, for example, relevant when states arrive at different times, and thus experience different amounts of decoherence.

5.6.2. ACHIEVED DISTILLATION STATISTICS FOR n-FOLD WERNER STATES

Here we show our results for the case of an n-fold tensor product of a Werner state. First, we list the number of cases to check (i.e. the number of triples (a, b, E), see Section 5.5) and the number of distinct distillation protocols for this scenario in Table 5.2.

	2	3	4	5	6	7	8
Cases	2	10	60	561	6358	111540	2917980
Distinct protocols	2	4	12	31	86	303	1131

Table 5.2: Number of cases to check (i.e. the number of triples (a, b, E) and number of distinct distillation protocols for n identical Werner states.

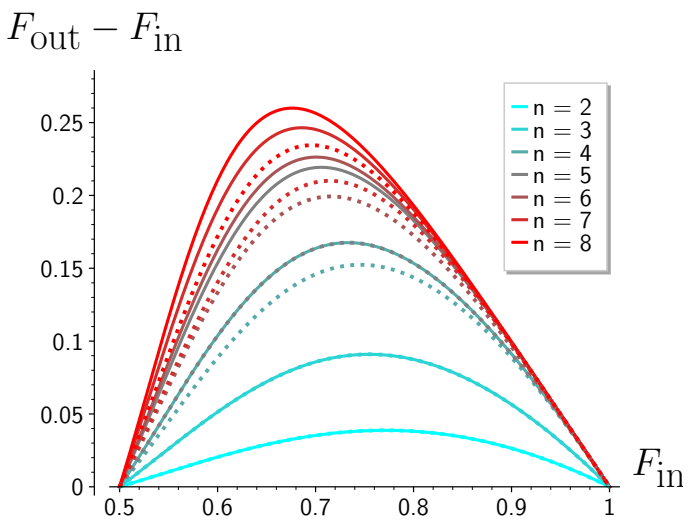

Figure 5.5: Comparison between the increase in fidelity $F_{\text{out}} - F_{\text{in}}$ with our optimisation (solid) and concatenated DEJMPS protocol (dotted), for $n = 2$ to 8 identical Werner states with fidelity F_{in}. Note how the $n = 5$ concatenated DEJMPS protocol overlaps with an optimised $n = 4$ protocol.

The number of cases and distinct protocols still rapidly grow with n, but our reduction allowed to consider all possible distillation protocols for $n = 8$ in about a day of computer run-time. This should be compared with a naive optimisation over all elements of the symplectic group — for $n = 8$ the ratio between the order of the symplectic group and the number of cases to check is approximately $2 \cdot 10^{34}$.

In order to gauge the advantage of the optimal protocols that we find for Werner states, we compare them with the class of protocols we call *concatenated DEJMPS protocols*. These are bilocal Clifford protocols that are built from multiple iterations of the DEJMPS protocol, see Appendix 5.8.2 for more information. The concatenated DEJMPS protocols form a natural generalisation of the (nested) entanglement pumping protocols [47].

We first investigate the increase in fidelity $F_{\text{out}} - F_{\text{in}}$ conditioned on the success of the distillation protocol. We plot the increase in fidelity as a function of the input fidelity F_{in} for $n = 2, 3, \ldots, 8$ in Fig. 5.5. The dotted lines correspond to the concatenated DEJMPS protocols, the solid lines correspond to the single protocol that achieves the highest output fidelity found with our optimisation. For completeness, we show the success probabilities and fidelities for the optimised protocols for $n = 2, 3, \ldots, 8$ in Tables 5.3, 5.4, 5.5 and 5.6 in Appendix 5.8.5.

Let us now discuss Fig. 5.5. First we observe that for $n = 2, 3$, the optimal protocols correspond to the original DEJMPS [42] and double-selection [54] protocols. However, for $n > 3$, we find distillation protocols that outperform the concatenated DEJMPS protocols.

We find that the optimal protocol for $n = 4$ achieves the same fidelity as the concatenated DEJMPS protocol for $n = 5$, and can be executed with a circuit of the same depth as the concatenated DEJMPS protocol. This protocol achieves the same distillation statistics as the protocol found with different means in the recent work from [181].

For $n = 5$ there is a large gap between the optimised protocols and the concatenated DEJMPS protocol. We make this now more quantitative by expanding the F_{out} for high input fidelity $F_{\text{in}} \approx 1$. For $n = 5$, the concatenated DEJMPS protocol has quadratic scaling in the infidelity,

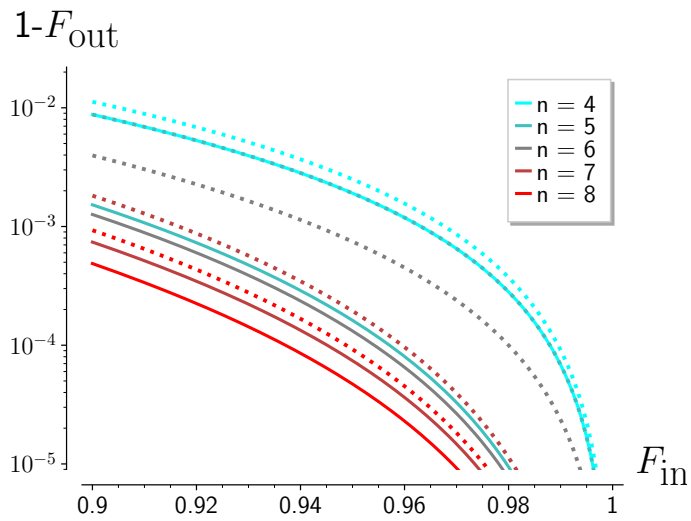

Figure 5.6: Comparison between the achieved infidelities $1 - F_{\text{out}}$ with our optimisation (solid) and concatenated DEJMPS protocol (dotted), for $n = 4$ to 8 identical Werner states with fidelity F_{in}.

$$1 - \frac{2}{3}(1-F)^2 + \mathcal{O}\left((1-F)^3\right), \tag{5.23}$$

while the optimised protocol has a cubic scaling in the infidelity

$$1 - \frac{10}{9}(1-F)^3 + \mathcal{O}\left((1-F)^4\right). \tag{5.24}$$

This is particularly surprising, since previous protocols with five or less pairs [87] have a scaling that is at most quadratic in the infidelity. We list the scaling of the found protocols in Table. 5.7.

Next, we investigate the behavior of the protocols for high fidelities $F_{\text{in}} \approx 1$. In Fig. 5.6 we plot the infidelity $1 - F_{\text{out}}$ as a function of the input fidelity F_{in}. We observe that it is possible to reach fidelities of around 0.999 using six copies of Werner states with fidelity $F_{\text{in}} = 0.9$. We do not plot the results for $n = 2, 3$ since we find no improvements with respect to previous protocols.

We have seen that the optimised distillation protocols are capable of achieving a higher fidelity than the concatenated DEJMPS protocols. However, the optimised distillation protocols also have a lower success probability. This motivates us to investigate a metric that combines the success probability and the quality of the resultant state. As a metric we use the distillable entanglement rate which we approximate by combining the distillation protocol together with a *hashing* protocol [11]. That is, given n entangled pairs, we first perform an n-to-1 distillation protocol and then use the output as input for the hashing protocol. The rate r at which this procedure produces maximally entangled state is given by

$$r = \frac{\big(1 - H(p)\big) \cdot p_{\text{suc}}}{n}, \tag{5.25}$$

where $H(p)$ is the entropy in bits of the probability distribution $p = \big(p_{\text{I}}, p_X, p_Y, p_Z\big)$ corresponding to the output state. This metric has been used previously and is sometimes called the hashing yield [87].

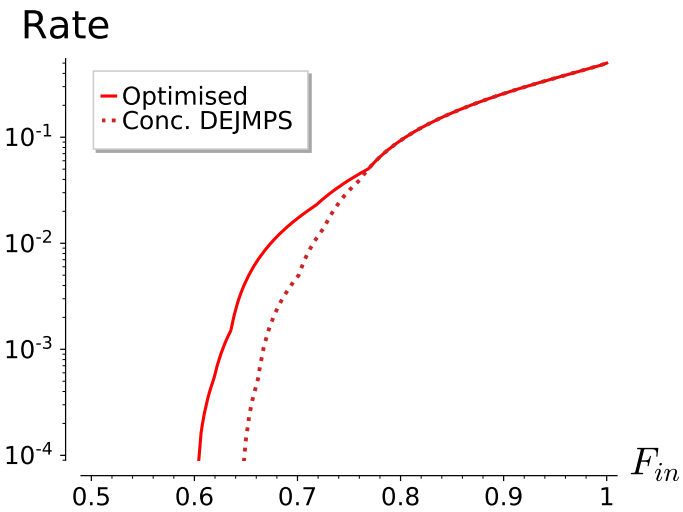

Figure 5.7: Comparison between the achieved rates after distilling and then hashing with our optimisation (solid) and the concatenated DEJMPS (dotted) protocol. For both cases, we take the envelope of all protocols on $n = 2$ to 8 identical Werner states with fidelity F_{in}.

We plot the comparison between the achieved rate of all found distillation protocols and the concatenated DEJMPS protocols in Fig. 5.7. We find that for $n > 2$ and fidelities $\lesssim 0.78$ the optimal bilocal Clifford protocols achieve higher rates than concatenated DEJMPS protocols. Furthermore, these protocols achieve a non-zero rate with an input fidelity of ~ 0.6, as opposed to concatenated DEJMPS protocols which require a fidelity of ~ 0.65. On the other hand, fidelities $\gtrsim 0.78$ would be expected for practical distillation, showing that concatenated DEJMPS protocols perform the best out of all bilocal Clifford distillation protocols for achieving the highest asymptotic rate.

Let us conclude with an investigation of circuits that achieve the highest fidelity for $n = 4$ to 8. Interestingly, these protocols can be implemented with low-depth circuits. We performed a search over circuits of the form described in Appendix 5.8.3, to find circuits that achieve the highest fidelity. We report these circuits in Appendix 5.8.5. For $n = 4$ to 8, we find a total number of two-qubit gates of 4, 7, 8, 11 and 13. Furthermore, the corresponding circuit depths are 3, 5, 6, 6 and 7, respectively. For comparison, the circuit from [181] for $n = 4$ pairs has 4 two-qubit gates and depth 5. This protocol can be converted to our optimal $n = 4$ protocol by left-multiplication with elements in $\mathscr{D}_n$ and right-multiplication with elements in $\mathscr{K}_n$. Therefore, both protocols achieve the same distillation statistics. The protocol from [105] for $n = 5$ pairs, which achieves the same distillation statistics as the concatenated DEJMPS protocol, has 4 two-qubit gates and depth 4.

5.7. CONCLUSIONS AND DISCUSSIONS

Our goal in this chapter was to find good distillation protocols requiring modest resources. For this, we introduced the class of bilocal Clifford protocols which generalises many existing protocols. The protocols in this class require only a single round of communication between the end parties and the implementation of Clifford gates. Within this class, we leveraged group theoretic tools to find all distillation protocols for up to $n = 5$ pairs for general Bell-diagonal states and up to $n = 8$ pairs for the n-fold tensor product of a Werner state.

Some of the protocols that we found strongly improve upon the fidelities and rates of previous

protocols for certain values of the initial fidelity. At the same time, we have shown that concatenated DEJMPS protocols, a generalisation of a previous protocol, achieve the highest rate out of all bilocal Clifford protocols for relevant input fidelities. Moreover, we give explicit circuits for the optimal protocols for the n-fold Werner state case, with $n = 2$ to $n = 8$. These circuits have comparable depth and number of two-qubit gates as previous protocols, indicating the experimental feasibility of the new protocols. If the improved performance holds with noisy operations, then it will translate in improved forecasts for the performance of near-term quantum networks [105, 167] or distributed quantum computation [119]. Finally, since we have enumerated all bilocal Clifford protocols up to $n = 5$, it is possible to pick and choose the protocol that maximises any figure of merit for any particular set of input states. Our software can be found at [76].

In this work we considered distilling one entangled pair out of n pairs. The results here could be extended to n to m distillation protocols by generalising Lemma 5.5.1 and the characterisation of the distillation subgroup to the case of n to m distillation. Such distillation protocols would allow for a more refined trade-off between the fidelity and the success probability/rate. Another interesting avenue would be to generalise the tools to higher dimensional entangled states.

5.8. APPENDIX

5.8.1. BACKGROUND ON BINARY PICTURE

For completeness, we give here more background and derivations on the binary picture used in this work.

Firstly, we give a derivation of equation (5.1). Suppose that we have two elements $\tau_a, \tau_b \in \overline{\mathscr{P}}_n$, with $a = \begin{bmatrix} v^a \\ w^a \end{bmatrix} = [v_1^a \dots v_n^a \; w_1^a \dots w_n^a]^\mathsf{T}$ and $b = \begin{bmatrix} v^b \\ w^b \end{bmatrix} = \left[v_1^b \dots v_n^b \; w_1^b \dots w_n^b \right]^\mathsf{T}$. Then

$$\begin{aligned} \tau_a \tau_b &= (\tau_{v_1^a w_1^a} \otimes \dots \otimes \tau_{v_n^a w_n^a})(\tau_{v_1^b w_1^b} \otimes \dots \otimes \tau_{v_n^b w_n^b}) \\ &= \bigotimes_{k=1}^{n} \tau_{v_k^a w_k^a} \tau_{v_k^b w_k^b}. \end{aligned} \tag{5.26}$$

For all $k \in \{1, \dots, n\}$, we have

$$\begin{aligned} \tau_{v_k^a w_k^a} \tau_{v_k^b w_k^b} &= X^{v_k^a} Z^{w_k^a} X^{v_k^b} Z^{w_k^b} \\ &= X^{v_k^a} (-1)^{v_k^b w_k^a} X^{v_k^b} Z^{w_k^a} Z^{w_k^b} \\ &= (-1)^{v_k^b w_k^a} X^{v_k^a + v_k^b} Z^{w_k^a + w_k^b} \\ &= (-1)^{v_k^b w_k^a} \tau_{v_k^a + v_k^b, w_k^a + w_k^b}. \end{aligned} \tag{5.27}$$

As a result,

$$\begin{aligned} \tau_a \tau_b &= \bigotimes_{k=1}^{n} (-1)^{v_k^b w_k^a} \tau_{v_k^a + v_k^b, w_k^a + w_k^b} \\ &= (-1)^{\sum_{k=1}^n v_k^b w_k^a} \tau_{v^a + v^b, w^a + w^b} \\ &= (-1)^{v^b \cdot w^a} \tau_{a+b}. \end{aligned} \tag{5.28}$$

Here $v^b \cdot w^a$ is the standard vector dot product. We can rewrite this dot product in terms of the vectors a and b:

$$v^b \cdot w^a = b^\mathsf{T} \Xi a, \qquad \Xi = \begin{bmatrix} 0 & I_n \\ 0 & 0 \end{bmatrix}. \tag{5.29}$$

Hence, the product of τ_a and τ_b is given by

$$\tau_a \tau_b = (-1)^{b^\mathsf{T} \Xi a} \tau_{a+b}. \tag{5.30}$$

Combining equation (5.30) for $\tau_a \tau_b$ and $\tau_b \tau_a$, we finally obtain

$$\tau_a \tau_b = (-1)^{b^\mathsf{T} \Xi a + b^\mathsf{T} \Xi^\mathsf{T} a} \tau_b \tau_a = (-1)^{b^\mathsf{T} \Omega a} \tau_b \tau_a, \tag{5.31}$$

$$\text{where } \Omega = \Xi + \Xi^\mathsf{T} = \begin{bmatrix} 0 & I_n \\ I_n & 0 \end{bmatrix}. \tag{5.32}$$

Let $C \in \mathscr{C}_n$ be a Clifford operation and $f : \overline{\mathscr{P}}_n \to \overline{\mathscr{P}}_n$, $f(P) = CPC^\dagger$ be the corresponding automorphism. Let $\pi : \mathbb{F}_2^{2n} \to \mathbb{F}_2^{2n}$ be the representation of f in the binary picture. Let $a, b \in \mathbb{F}_2^{2n}$. Then we know that $C\tau_{a+b}C^\dagger = (-1)^{b^\mathsf{T} \Xi a} C\tau_a \tau_b C^\dagger = (-1)^{b^\mathsf{T} \Xi a} C\tau_a C^\dagger C\tau_b C^\dagger$. In the binary representation, the prefactor $(-1)^{b^\mathsf{T} \Xi a}$ does not make a difference. Thus, $\pi(a+b) = \pi(a) + \pi(b)$, so π is a linear map, and there exists a binary $2n \times 2n$ matrix M such that $\pi(a) = Ma$ for all $a \in \mathbb{F}_2^{2n}$.

5.8.2. CONCATENATED DEJMPS PROTOCOLS

Here we describe the distillation protocols which we compare our results with. These are all based on the so-called DEJMPS protocol [42]. The DEJMPS protocol takes two pairs of Bell-diagonal states, and outputs one state. It performs bilocal single-qubit rotations on both pairs, then a bilocal CNOT, and finally a measurement on one of the pairs where a success is achieved only when correlated outcomes are observed. It is clear that the DEJMPS protocol is an example of a bilocal Clifford protocol. The DEJMPS protocol can be generalised to a number of pairs $n > 2$ by applying the DEJMPS protocol multiple times.

Since the DEJMPS protocol corresponds to 2-1 distillation, the possible ways of combining the different pairs correspond to the number of binary trees on n unlabeled nodes for an n-fold tensor product of input states. Furthermore, for each of the performed DEJMPS protocols (corresponding to each parent of the binary tree), we consider all possible single-qubit rotations. The *concatenated DEJMPS protocol* is then the protocol that has the highest fidelity out of all such protocols, and is found by calculating the fidelity of each possible configuration. Note that this optimisation includes well known variants of DEJMPS such as (nested) entanglement pumping protocols [20, 47, 48] or double selection [54].

5.8.3. DISTILLATION CIRCUITS

In this section we are concerned with finding circuits that achieve the highest fidelity for $n = 4$ to 8 for an n-fold tensor product of a Werner state[2].

We first note that one could use techniques for general Clifford circuit decompositions to decompose the symplectic matrices of the form in 5.5.1. However, we found that this would in general lead to circuits with high depths. Instead, we first find that any distillation protocol has a circuit in a given form. Then, we randomly generate circuits of that form, until we find circuits that achieve the highest fidelity, and have small depth.

5.8.4. REDUCING THE CIRCUIT SEARCH SPACE

We use the Bruhat decomposition from [19, 103], which allows to write any Clifford circuit C in the form $C = FWF'$, with F and, F' elements of the so-called Borel subgroup[3] and W a layer of Hadamard gates followed by a permutation $\sigma \in \mathscr{S}_n$. The Borel subgroup $\mathscr{B}_n$ is generated by X_i, S_i for $1 \leq i \leq n$ and CNOT_{ij} with $1 \leq j \leq i \leq n$. For convenience, we denote such CNOT gates as $\mathrm{CNOT}^{\uparrow}$ gates. Now note that the Borel subgroup $\mathscr{B}_n$ is a subgroup of the distillation subgroup $\mathscr{D}_n$. This implies that the F part of any circuit in the form $C = FWF'$ does not change the distillation statistics. Thus, any distillation protocol has a corresponding circuit of the form WF'. Furthermore, since the distillation subgroup $\mathscr{D}_n$ contains elements that arbitrarily permute qubits 2 to n, we can restrict to permutations that are either the identity, or exchange qubit 1 with j. In practice, we have found that it sufficient to only consider $W = H_2 H_3 \dots H_n$.

By the results from [19], any element F' from $\mathscr{B}_n$ can be written as a layer of $\mathrm{CNOT}^{\uparrow}$ gates, a layer of CZ gates, a layer of phase gates and a layer of Pauli gates. Firstly, the Pauli gates are in the kernel of ϕ, and thus can be left out. Secondly, the layer of S gates can be moved to the beginning. To see this, first note that phase gates commute with CZ gates. Then, since $\mathrm{CNOT}_{ij} S_i = S_i \mathrm{CNOT}_{ij}$ and $S_j \mathrm{CNOT}_{ij} = \mathrm{CNOT}_{ij} \mathrm{CZ}_{ij} S_i S_j$, the layer of S gates can be moved to the beginning. Since Werner states are invariant under S, the layer of S gates can be removed without changing the distillation statistics. In the above process of moving the S gates to the beginning, the layer that

[2] We are not interested in the cases $n = 2$ and $n = 3$, since for those cases the concatenated DEJMPS protocols are optimal.

[3] We use a different convention from [19, 103], where the target index for the CNOT gates in the Borel subgroup is larger than the control index.

only had $\mathrm{CNOT}^{\uparrow}$ gates will now have CZ gates as well. In the binary picture we have the following identities,

$$\mathrm{CNOT}_{ij}\mathrm{CZ}_{kl} = \mathrm{CZ}_{kl}\mathrm{CNOT}_{ij}\,, \tag{5.33}$$

$$\mathrm{CNOT}_{ij}\mathrm{CZ}_{ij} = \mathrm{CZ}_{ij}\mathrm{CNOT}_{ij}\,, \tag{5.34}$$

$$\mathrm{CNOT}_{ik}\mathrm{CZ}_{ij} = \mathrm{CZ}_{ij}\mathrm{CNOT}_{ik}\,, \tag{5.35}$$

$$\mathrm{CNOT}_{ik}\mathrm{CZ}_{jk} = \mathrm{CZ}_{ij}\mathrm{CZ}_{jk}\mathrm{CNOT}_{ik}\,, \tag{5.36}$$

where the i, j, k, l are assumed to be distinct, and can be verified using Lemma 5.2.1. By using the above identities, the CZ gates can be moved through to the original layer of CZ gates.

It is thus sufficient to consider only elements F' that consist of a layer of $\mathrm{CNOT}^{\uparrow}$ gates and a layer of CZ gates. Now, to find circuits we randomly generate circuits consisting of a layer of $\mathrm{CNOT}^{\uparrow}$ gates, a layer of CZ gates, and a Hadamard gate on all qubit except the first. We found several circuits that achieved the largest fidelity, and choose the one with smallest depth. We report the found circuits in Fig. 5.8.

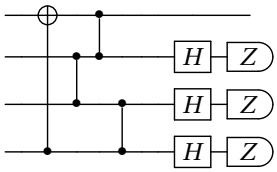

(a) $n = 4$, #2-qubit gates = 4, depth = 3.

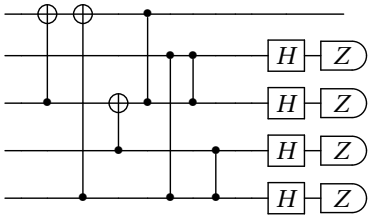

(b) $n = 5$, #2-qubit gates = 7, depth = 5.

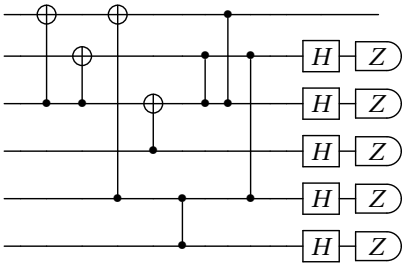

(c) $n = 6$, #2-qubit gates = 8, depth = 6

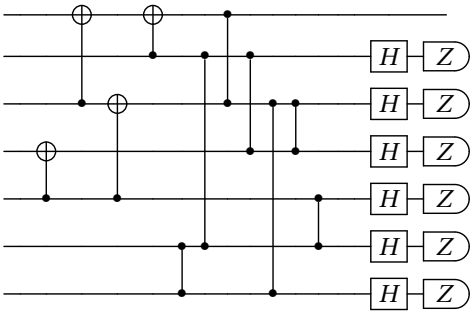

(d) $n = 7$, #2-qubit gates = 11, depth = 6.

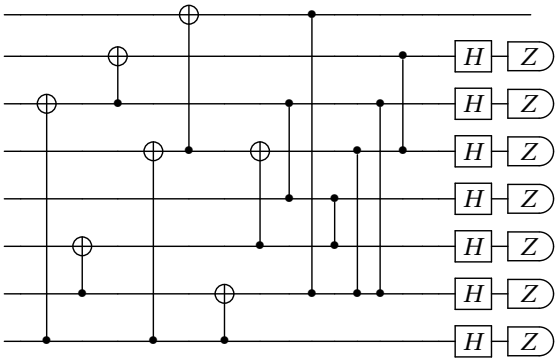

(e) $n = 8$, #2-qubit gates = 13, depth = 7.

Figure 5.8: Circuits that achieve the maximum fidelity for n. These circuits are applied by both Alice and Bob, after which they measure the last $n-1$ qubits, and communicate their outcomes to each other. When the outcomes for all individual qubit pairs are correlated, the distillation protocol was deemed successful.

5.8.5. ANALYTICAL EXPRESSIONS

We report here the analytical expressions of the fidelity and success probability that correspond to the found optimal schemes. The input state is an n-fold tensor product of a Werner state with fidelity F. For completeness, we report here as well the distillation statistics expressed in the infidelity $\epsilon \equiv 1 - F$, and the scaling of the output fidelity as a function of the infidelity.

n	p_{suc}
2	$\frac{8}{9}F^2 - \frac{4}{9}F + \frac{5}{9}$
3	$\frac{32}{27}F^3 - \frac{4}{9}F^2 + \frac{7}{27}$
4	$\frac{32}{27}F^4 - \frac{4}{9}F^2 + \frac{4}{27}F + \frac{1}{9}$
5	$\frac{80}{27}F^4 - \frac{80}{27}F^3 + \frac{10}{9}F^2 - \frac{5}{27}F + \frac{2}{27}$
6	$\frac{128}{243}F^6 + \frac{320}{243}F^5 - \frac{256}{243}F^4 + \frac{16}{243}F^3 + \frac{40}{243}F^2 - \frac{14}{243}F + \frac{1}{27}$
7	$\frac{2048}{2187}F^7 - \frac{128}{2187}F^6 + \frac{320}{729}F^5 - \frac{796}{2187}F^4 - \frac{44}{2187}F^3 + \frac{49}{729}F^2 - \frac{37}{2187}F + \frac{37}{2187}$
8	$\frac{6656}{6561}F^8 - \frac{1024}{6561}F^7 + \frac{1664}{6561}F^6 - \frac{64}{6561}F^5 - \frac{1120}{6561}F^4 + \frac{416}{6561}F^3 - \frac{4}{6561}F^2 - \frac{16}{6561}F + \frac{53}{6561}$

Table 5.3: Success probability for the protocols with the highest output fidelity for $n = 2$ to 8.

n	$p_{\text{suc}} \cdot F_{\text{out}}$
2	$\frac{10}{9}F^2 - \frac{2}{9}F + \frac{1}{9}$
3	$\frac{28}{27}F^3 - \frac{1}{9}F + \frac{2}{27}$
4	$\frac{8}{9}F^4 + \frac{8}{27}F^3 - \frac{2}{9}F^2 + \frac{1}{27}$
5	$\frac{32}{27}F^5 - \frac{20}{27}F^4 + \frac{10}{9}F^3 - \frac{20}{27}F^2 + \frac{5}{27}F$
6	$\frac{32}{27}F^6 - \frac{112}{243}F^5 + \frac{80}{243}F^4 + \frac{8}{243}F^3 - \frac{32}{243}F^2 + \frac{10}{243}F + \frac{1}{243}$
7	$\frac{2368}{2187}F^7 - \frac{592}{2187}F^6 + \frac{196}{729}F^5 - \frac{44}{2187}F^4 - \frac{199}{2187}F^3 + \frac{20}{729}F^2 - \frac{2}{2187}F + \frac{8}{2187}$
8	$\frac{6784}{6561}F^8 - \frac{51}{6561}F^7 - \frac{32}{6561}F^6 + \frac{832}{6561}F^5 - \frac{560}{6561}F^4 - \frac{8}{6561}F^3 + \frac{52}{6561}F^2 - \frac{8}{6561}F + \frac{13}{6561}$

Table 5.4: Product of the success probability and the output fidelity for the protocols with the highest output fidelity for $n = 2$ to 8.

n	p_{suc}
2	$1-\frac{4}{3}\epsilon+\frac{8}{9}\epsilon^2$
3	$1-2\epsilon+\frac{4}{3}\epsilon^2$
4	$1-2\epsilon+\frac{4}{3}\epsilon^2-\frac{8}{27}\epsilon^3$
5	$1-\frac{14}{3}\epsilon+\frac{28}{3}\epsilon^2-\frac{256}{27}\epsilon^3+\frac{400}{81}\epsilon^4-\frac{256}{243}\epsilon^5$
6	$1-5\epsilon+\frac{32}{3}\epsilon^2-12\epsilon^3+\frac{608}{81}\epsilon^4-\frac{608}{243}\epsilon^5+\frac{256}{729}\epsilon^6$
7	$1-7\epsilon+\frac{190}{9}\epsilon^2-\frac{944}{27}\epsilon^3+\frac{928}{27}\epsilon^4-\frac{544}{27}\epsilon^5+\frac{1600}{243}\epsilon^6-\frac{2048}{2187}\epsilon^7$
8	$1-\frac{23}{3}\epsilon+\frac{244}{9}\epsilon^2-\frac{1540}{27}\epsilon^3+\frac{6280}{81}\epsilon^4-\frac{16832}{243}\epsilon^5+\frac{28768}{729}\epsilon^6-\frac{9472}{729}\epsilon^7+\frac{4096}{2187}\epsilon^8$

Table 5.5: Success probability for the protocols with the highest output fidelity for $n = 2$ to 8, expressed in the infidelity $\epsilon \equiv 1-F$.

n	$p_{\text{suc}} \cdot F_{\text{out}}$	F_{out}
2	$1-2\epsilon+\frac{10}{9}\epsilon^2$	$1-\frac{2}{3}\epsilon-\mathcal{O}\left(\epsilon^2\right)$
3	$1-3\epsilon+\frac{10}{3}\epsilon^2-\frac{4}{3}\epsilon^3$	$1-\frac{1}{3}\epsilon-\mathcal{O}\left(\epsilon^2\right)$
4	$1-3\epsilon+\frac{10}{3}\epsilon^2-\frac{44}{27}\epsilon^3+\frac{8}{27}^4$	$1-\frac{2}{3}\epsilon^2-\mathcal{O}\left(\epsilon^3\right)$
5	$1-5\epsilon+\frac{92}{9}\epsilon^2-\frac{284}{27}\epsilon^3+\frac{440}{81}\epsilon^4-\frac{272}{243}\epsilon^5$	$1-\frac{10}{9}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$
6	$1-\frac{17}{3}\epsilon+\frac{122}{9}\epsilon^2-\frac{466}{27}\epsilon^3+\frac{992}{81}\epsilon^4-\frac{1112}{243}\epsilon^5+\frac{512}{729}\epsilon^6$	$1-\frac{8}{9}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$
7	$1-7\epsilon+\frac{190}{9}\epsilon^2-\frac{320}{9}\epsilon^3+\frac{2936}{81}\epsilon^4-\frac{1816}{81}\epsilon^5+\frac{5680}{729}\epsilon^6-\frac{2560}{2187}\epsilon^7$	$1-\frac{13}{27}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$
8	$1-8\epsilon+\frac{259}{9}\epsilon^2-\frac{544}{9}\epsilon^3+\frac{2180}{27}\epsilon^4-\frac{17000}{243}\epsilon^5+\frac{27872}{729}\epsilon^6-\frac{2912}{243}\epsilon^7+\frac{3584}{2187}\epsilon^8$	$1-\frac{8}{27}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$

Table 5.6: Product of the success probability and output fidelity (first column) and the scaling of the output fidelity around $\epsilon \approx 0$ (second column) for the protocols with the highest output fidelity for $n = 2$ to 8. Here $\epsilon \equiv 1-F$.

n	F_{out}
2	$1-\frac{2}{3}\epsilon-\mathcal{O}\left(\epsilon^2\right)$
3	$1-\frac{1}{3}\epsilon-\mathcal{O}\left(\epsilon^2\right)$
4	$1-\frac{2}{3}\epsilon^2-\mathcal{O}\left(\epsilon^3\right)$
5	$1-\frac{10}{9}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$
6	$1-\frac{8}{9}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$
7	$1-\frac{13}{27}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$
8	$1-\frac{8}{27}\epsilon^3-\mathcal{O}\left(\epsilon^4\right)$

Table 5.7: Scaling of the output fidelity around $\epsilon \approx 0$ for the protocols with the highest output fidelity for $n = 2$ to 8, where $\epsilon \equiv 1 - F$.

6

CONCLUSION

The research presented in this thesis focused on the problem of entanglement distribution. Simply put, the two main problems facing (practical) implementation of entanglement distribution over quantum networks are *loss* and *noise*. Quantum repeaters are meant to overcome the effects of loss, but in practice their implementation always comes at the cost of more incurred noise. This additional noise can be overcome by the use of entanglement distillation.

In the first two chapters, we focused on the assessment of a basic building block for quantum networks, a single quantum repeater. We then considered finding schemes for the concatenation of multiple such quantum repeaters, along with the inclusion of basic distillation protocols. Finally, we considered a systematic way of optimising over a relevant class of (more complex) distillation protocols.

6.1. SUMMARY OF RESULTS

We summarise here the main contributions of this thesis.

- The cut-off is a tool that allows for the trade-off between the rate and the fidelity of the received qubits. The cut-off is easy to implement experimentally, yet allows for an increase in the secret-key generated in comparison to the no cut-off scenario, especially as the distance increases. In fact, for the case of a single repeater node as presented in chapters 2 and 3, we observe that it allows for an approximately one-third increase in the total distance over which non-zero secret-key can be generated. While the cut-off makes the theoretical analysis of quantum repeaters more complicated, this is outweighed by the practical benefits it provides.

- The number of (near-deterministic) entanglement distribution schemes over linear repeater chains grows super-exponentially in the number of elementary links. In practice, one cannot hope to exploit symmetry of the setup to simplify the optimisation. We have proposed a heuristic optimisation over repeater schemes which runs in polynomial time. The optimisation is general, and can be applied to a wide range of different experimental platforms. We have found significant improvements (in the sense of higher fidelity and generation rates) for asymmetric repeater chains when compared with a naive optimisation. Furthermore,

we have found heuristics on what makes a 'good' near-deterministic scheme. In particular, the banded swapping heuristic relates the fidelities and the distances over which entanglement has been generated for good schemes.

- Entanglement distillation on a small number of copies ($\sim 4-8$) can achieve higher fidelities than thought previously, improving on known protocols and naive generalisations thereof. At the same time, these new protocols can be implemented using bilocal Clifford circuits (with modest depth and number of two-qubit gates) and a single round of communication only. Furthermore, all such protocols can be classified for general Bell-diagonal states and the n-fold tensor product of a Werner state for up to 5 and 8 pairs.

6.2. OUTLOOK

Here we will discuss future research directions relevant to this thesis and, more importantly, the creation of future quantum networks. We discuss first single quantum repeater setups, building up to linear repeater chains, and finally concluding with a discussion on (bilocal Clifford) entanglement distillation protocols.

6.2.1. SINGLE QUANTUM REPEATERS

Quantum repeaters form an essential building block for quantum networks. The most basic setup is that of a single quantum repeater node situated between two parties. Demonstrating a proof-of-principle implementation of a single quantum repeater node forms a vital step towards constructing quantum networks.

There have been several proposals for such a single quantum repeater node. We compared four different proposals to find which would work best for a proof-of-principle implementation of a single repeater node. We found that the single-photon scheme was the most suitable candidate. One of the benefits of this scheme was that it did not require storage of quantum states over time. Interestingly, one of the first experimental implementations that has surpassed the secret-key capacity [108] was essentially the same as the single-photon scheme, and in particular did not require the storage of quantum states. However, extending such proof-of-principle quantum repeater schemes to an additional node, let alone multiple, will necessarily require the storage of states. In this sense, the assessment of quantum repeater schemes/implementations that are more amenable for scaling up to repeater chains is of importance. At the same time, a detailed analysis of multi-node repeater chains that involve the storage of states is complicated by the stochastic nature of the operations involved. We expect that a proper assessment cannot be treated analytically, and will require the simulation of such schemes.

From a scientific point of view, proof-of-principle experiments for quantum repeaters are valuable - surpassing the secret-key capacity allows one to make a clear statement on the performed experiment. Namely, the implemented setup achieved something that would have been impossible without a quantum repeater. It is thus not surprising that comparing the performance of quantum repeaters with the secret-key capacity has been a staple in the quantum information literature so far. However, after sufficient time proof-of-principle experiments will have become hopefully more commonplace. At this point in time, quantum repeater setups will evolve to become more of a means to an end. Both theoretical and experimental efforts will then not focus on surpassing the maximum achievable secret key (or any other practically relevant quantity) per channel use for direct transmission, but on increasing the secret key (or any other practically relevant quantity) per unit time.

6.2.2. LINEAR REPEATER CHAINS

For large enough distances, a single repeater will not suffice to meet future entanglement rate requirements. This necessitates the introduction of multiple repeater nodes, leading in the simplest case to linear repeater chains. In this thesis we considered near-deterministic schemes over linear repeater chains. Such schemes have the benefit of delivering an entangled state with high probability at a fixed time. This comes at the cost of increased requirements on the quantum memories, since states are in general stored longer. This trade-off becomes especially important at the level of quantum networks, where the failure of generating entanglement within a certain time-frame will impact other processes in the network. At the same time, near-deterministic schemes will in general occupy components of the network for a longer time. One could imagine that near-deterministic schemes would be useful for creating 'backbones' in quantum networks - sections of a quantum network dedicated to generating entanglement consistently at a fixed rate and fidelity. In such a backbone, the consistent generation of entanglement between fixed nodes is one of the most important features.

Furthermore, the relation between schemes being near-deterministic and the practicality of entanglement distillation is worthwhile to investigate. While entanglement distillation from multiple states to one state can increase fidelity, all states need to be present. Due to the stochastic nature of the operations in a network, the states will arrive at different times. The additional experienced decoherence due to this can quickly make distillation fruitless. This time-delay can be decreased by using near-deterministic scheme, since the success of such schemes has a small variance by construction. It is not clear in which cases the trade-off between a smaller variance against a total larger overhead in waiting time will provide a benefit.

Our optimisation results have been mostly used in a quantitative fashion - that is, results on the achievable pairs of fidelities and entanglement generation rates for the given experimental parameters. However, we have also implemented heuristics that restrict the types of considered schemes, which were inspired by optimisation results without the heuristics. A concrete example is given by the banded swapping heuristic, which relates the allowed fidelities and lengths of the (multi-hop) links. A further investigation of such qualitative statements can inspire design choices for quantum networks. In particular for (strongly) asymmetric repeater chains it is not clear whether certain schemes are bad or good *a priori*. For example, given two (multi-hop) links of the same length, but different parameters, can statements be made on which link should be preferred for entanglement distillation? If such statements do not depend in a complicated fashion on the underlying physical parameters, they could potentially lead to relatively simple guidelines for designing quantum networks and the operations running on them.

6.2.3. ENTANGLEMENT DISTILLATION PROTOCOLS

Entanglement distillation protocols allow for overcoming the noise incurred in any realistic implementation of a quantum network. For this reason, they form an important part of any future quantum internet. This motivated the search for new experimentally feasible distillation protocols. We introduced the class of bilocal Clifford protocols, and enumerated/optimised over all such protocols for up to $n = 8$ pairs. For more than $n = 3$ pairs, we found explicit protocols that achieve higher fidelities than previously known protocols. Importantly, these protocols require only a single round of communication, and circuits that have both low depth and a small number of two-qubit gates.

The results found in this thesis can be generalised by distilling to a number of copies $m \geq 1$, instead of only distilling to one pair. Furthermore, the results herein could be generalised to the case where one relaxes the post-selection condition on the measurements. That is, certain observed patterns of (anti-)correlated outcomes can be counted as a success. This leads to a trade-off between the success probability and the fidelity. An example of a protocol that falls into this more

general class is the hashing protocol, which is known to be able to distill certain Bell-diagonal states to maximally entangled states at a non-zero rate. This protocol has been used in the literature to assess how well certain tasks can be performed. Exploring a class of protocols which the hashing protocol belongs to, in both the finite and asymptotic case, would thus be of value.

The techniques employed rely strongly on the input states being Bell-diagonal and the allowed operations being bilocal Clifford gates. Certain techniques could carry-over to the case of arbitrary bilocal unitaries on n-fold tensor products of Werner states. In that case, the corresponding distillation subgroup and the subgroup leaving the states invariant would become infinite. The first can be seen by noting that the bilocal application of $\exp(i \cdot tZ)$ on any of the measured out registers commutes with the measurements. The second can be seen by the fact that the n-fold tensor product of a Werner state is invariant under (at least) the subgroup generated by all possible SWAP operations and local unitaries. Ideally, the corresponding double cosets would have finite index in the unitary group on n qubits. This would reduce the optimisation over a continuous set to one over a finite (but potentially very large) set. Even for infinite index, a characterisation of the double cosets might still be possible.

7

ACKNOWLEDGEMENTS

Here I would like to get informal and show my gratitude to those who were invaluable during my time as a PhD candidate. These will be presented mostly (pseudo-)randomly.

First off, I would like to thank **Stephanie** Wehner for providing me with support and opportunities over the years, of which especially the mentoring during my Master's project I look back upon fondly. It has been truly amazing to see how you have pushed the development of QuTech and its quantum internet efforts over the time I've spent here, and I am curious to see what the future will bring under your leadership. Querido **David**, gracias por tu apoyo, calma y supervisión durante mis estudios de doctorado y máster, en los que me hiciste ganar interés en aspectos más teóricos de información cuántica, concretamente el estudio de canales cuánticos y medidas de entrelazamiento. Quizás algún día seremos capaces de resolver la NPPT-conjetura de destilabilidad (: Afortunadamente, y aunque nuestras dinámicas han ido cambiando durante las diferentes etapas académicas, siempre he podido contar con tu amabilidad, honestidad y los ocasionales (o no) juegos de palabras que han hecho estos últimos cuatro años tan formativos. ¡Gracias!

One of my favourite aspects of my time spent here in Delft were my lovely (ex-)colleagues. I am still amazed and grateful that a collection of such talented people can be so kind and social. I would like to highlight a few of these people, in no particular order. **Filipku**, we got to work on two out of my four papers herein — I'm not sure if I would have managed without you and your inspiration. I'm glad we are still in contact, both as friends and as colleagues. Unfortunately I don't think we will get to finish our project on unspeakable information, for some reason I find it hard to discuss :(**Axel**, thank you for simultaneously endangering my life and saving it on that mountain in Austria. For a (climbing/work) monster you manage to be very kind, patiently obliging me with explanations on all things graph related. With my thesis done, I hope to get to spend some more time on Vim :) **Jérémy**, your silent passion is something that I admire — during our time in Washington D.C. I learnt more about algorithmic complexity theory than I did about quantum information theory, and I'm happy I did! Kaushik — **KC** — thank you for the cooking, humour, and time spent dancing! Our occasional walk & talks were much needed and appreciated. Let us finish the ritual of Chüd someday soon, KC. **Jonas**, I am not sure how you managed to import the whole of Wikipedia (and more) into your brain, including all things esoteric and obscure — I enjoyed very much picking your brain during your time here. **Jed**, thank you for the conversations, especially the recent ones. They somehow manage to oscillate between the lamest of 'quantum

"puns" ', to the future of quantum information and our responsibility as researchers in it (and back again to the "puns"). **Corsin**, thank you for providing one of the first impetuses to dive into math proper. **Sébastian**, I am not sure how you have survived with me for this long, with our differences in music tastes and all. You made for an excellent neighbour in the office, and also for an excellent colleague virtually. I will miss the fact that you've accepted my antics (and may have even started to appreciate them). **Mark**, I really enjoyed our time together, getting to know how to push your buttons (physically or otherwise) — thank you for not reporting all of my misconduct in the office. I'm glad there was also someone else who could enjoy my sense of 'humour', from the dreaded Helvetica Scenario to T&E Awesome Show! (great job!) **Matt**, it's great to have you as a friend — so many niche shared interests! Thank you in particular for not only sharing my interests in quadruple parallel universes, but also experiencing them together with me. Also, a big thank you for the emergency thesis recompiling. Hopefully see you sometime soon again! **Ben**, you're an amazing cook, gym-buddy, skeleton-mover and, occasionally, a surprisingly good improv rapper. Your passion(s) and ideals are really under-appreciated. **Gláucia**, if more people would be as bad as you with hiding their thoughts/feelings, the world would not only be a better place, it would also be a lot more funny. I'm looking forward to keep following your future career, both academically and musically. **Leon**, it's nice we got to reconnect a bit recently, thanks for the conversations and for destroying me in chess! **Guus**, it was nice to have someone to talk to about both general relativity and algebraic topology, it would be nice if we could delve such topics together sometime soon :) **Carlo**, thank you for the vegan cooking tips, and I will fondly remember your 'joke' on skull-shaped pirates on a skull-shaped ship sailing to a skull-shaped island to find a sku- (even though it would've been even better if I was there for the first part). It's nice to know another Lynch enthusiast in theoryland :) **Lennart**, your sense of humour and knowledge about (world) history is very impressive. I look fondly back at our adventures together as a (half-)wizarding duo during the evenings hosted by the amazing dungeon master and theorist extraordinaire **Joel**. **Francisco**, I'm glad to experience some of your warm Portuguese demeanour, especially in this cold Dutch weather. Ik ben blij dat je niet hebt geluisterd naar mijn advies om geen Nederlands te leren — wat je hebt geleerd tot dusver is zeer indrukwekkend! **Gayane**, we both make for amazing actors/models, and I appreciate our mock conversations about the 'plus' symbol (and the actual scientific conversations about distillation). Thank you for the recent inspiring talks on the academic lifestyle! **Yves**, not only does your drive to dig deeper make you a great student, but you're also great to be around with socially, please do keep that spark. **Thinh**, I always liked our discussions, in particular those about working out, (classical) music and music theory, all very welcome distractions! It's great to see that you have surpassed me in my piano skills as fast and decidedly as you did! As an aside, I'm still using your topology book (: **David**, thank you for the nice times, both at uni and at Klokssteeg 17 — your presence and laughter lights up the room :) **Siddhant**, thank you for the chess/humour/chess humour and for working through all of the openings (maybe sometime in the future for real? ;)) From both your social nature and academic motivation, I'm sure you'll get far in your career, academic or otherwise. (New) **Matt**, don't fret — I'm happy to remain your piano practice fairy! Hopefully we'll get to enjoy the concert soon. **Wojciech**, it is good that there is at least one other person to talk with about both TRVE KVLT stuff and squats during lunch breaks! Looking forward to our concert together, fingers crossed. **Hana**, thank you for the very open conversations, both online and during the walks, especially during the (first?) height of the pandemic. Qīn'ài de **Gao**, I hope you're enjoying your times here in the cold Netherlands. Though sometimes there is a bit of a language barrier between us, at least there is one universal language that we both share ;) *Nunc fluens* — plenty of new people have joined the Wehner group over time, all of them which I feel I did not get to know properly yet. **Álvaro**, **Aram**, **Bart**, **Bethany**, **Ravi**, **Scarlett** — good luck in your endeavours, both here at QuTech and in the future. Of course, a warm thank you for **Helena** — without your help I would mostly likely still be sending out committee invites by now.

Of course, there's my two lovely paranymphs, Daniël and Tim. **Daniël**, your focus and commitment is truly impressive. Why you had the feeling I wasn't a lost cause with keeping track of my calendar, I do not know, but me and my future employer are happy you did. I still think about klaprozen, and I seriously think that if more people did, there would be a lot less sadness in this world. It's pretty crazy to think that we know each other for about 11 years or so already! Here's to many more :) **Tim**, you're way too humble! You're a kind and well-spoken theorist, a rare combination indeed. I have always enjoyed our discussions, from math/physics to music, but also topics where I was completely out of my depth, such as culture or language (Dutch or otherwise). It is unfortunate we did not get to collaborate (musically/academically) as much as we would've liked, but luckily there'll always be time for that! Never lose your curiosity for useless things :)

Besides my time here as a PhD student, I got to enjoy six months at QuTech as a software engineer, or at least that was the official job title on the contract ;) Thanks to **Thomas** and **Remon**, for showing your immense skill at juggling difficulties both scientific and practical — truly impressive. **Bruno**, thanks for your patient and incredibly clear explanations. I hope to reach the same enjoyment in work/'work' as you do. **Josh**, thank you for your trust in me, and for being a great leader. You've found that delicate balance between treating people as, well, *people*, but also expecting them to push themselves beyond what they initially thought they were capable of, and I mean that in the best way possible. I look forward to seeing what you and your team will be able to pull off in the future.

I would like to thank the rest of my bandmates, **Tim**, **Sebas**, **Bas** — having something to share amongst yourself and that other people resonate with is a rare treasure indeed. Some of my favourite moments so far have been with performing with you guys, both practicing, recording and performing live. Meeting and hanging out with you guys naturally led to seeing the huge friend group from Rotterdam more often (you know who you are). From my first festival, to great holidays in Spain and Belgium, to great Sinterklaasavonden, to memorable (and sometimes not so memorable anymore..) parties... it's great to have such a large welcoming group of people.

Halfway through my PhD I picked up dancing Zouk, which has led me to meeting a host of people that I now hold dear. With me from the beginning was **Corien**. Though we're still both Zouk teenagers, I still feel proud how much you've improved as a dancer! Thank you for making me feel welcome in your home and family! **Laura**, I think everyone, myself definitely included, could learn from your open-mindedness and willingness to jump in at the deep end. I hope I will have plenty of more years where I can patiently learn from you. **Eric**, thank you for your views on the world, relationships, dancing, and for our shared interest in (applied) knot theory. **Willemijn**, thank you for the walks + talks about life/love, and your openness and honesty as a person — I don't think I will ever not be smiling while dancing with you! **Alma**, I wouldn't be half as good at dancing as I am now without our practice sessions (and lovely car rides) together, I'm looking forward to many more. **Laura B.**, your enthusiasm for both science, education, dance and exploration is something I got to appreciate a lot since we got to know each other proper, i.e. after Emmaus. Hopefully I'll have more time for Quantum4kids soon! Of course, I cannot forget about **Victoria**, motivating me to keep on practicing, and for happily obliging me every time I wanted to try out a new dance move (which we *never* did during working hours, obviously).

Ana, thank you so much for the conversations and the time we got to spend so far (get it?). It's good to know that the intersection of people that are into Death Grips, math, experimental movies is non-empty. **Joost**, thanks not only for the cover, but also for the rest of the (math/visual/visual math) inspiration. Your natural curiosity and tenacity are sure to get you far in life. **Dave C**, I'm glad you reached out to me, and for being in contact with me as you've been doing — and that from the other side of the world!

There are some friends I've had the luck of having with me already for a very long time. **Stefan**, Steef, Stafna, Nefats, ~~Stoeph..~~, I've been thinking about how this all started from the 'yoyo-incident', progressing all the way from the absolutely crazy times in high school (WoW marathons/ouzo/my rapping attempts/'gkick'/Faustas/crashing Jerry's party/KC and the Sunshine Band to just name a random few), to us both starting to resemble something you could call adults. Naturally, this also led me to reflect on how my life would've been had I not met you. Probably a lot more boring, though my mother would've had a lot less to worry about for sure — I thank you for the first part. Furthermore, I know you've been working very hard on the cure for my longface-syndrome, I hope to see the day where you'll rightfully claim your Nobel prize. **Bender**, from being that 'friend' who would cut me off on the bike, to daring with me the sickest of free-running moves, to growing up to being such a gentle person — who would've known? I'm glad I stuck around to find out. **Stallion**, is there anything we haven't accomplished in the virtual world, barring legendary co-op Halo 2? Maybe someday :) Collectively I would like to thank the Stallion Squad for our time together, whether it was playing DnD or videogames. In particular our matches playing Overwatch was one of the main things to keep me sane during my quarantine. **Patrick**, thanks for your views on life and music and for creating that space in which we both can be our unabashed selves. Few people are so dedicated to bettering themselves or the world, most likely because they haven't realised there isn't that much of a distinction. Of course I cannot skip mentioning your guidance during my pilot training.

Núria, mona, guapa, mamacita — it has been very nice to slowly get to know you better, and thank you for taking me as I am (and for doing so better than I do myself). Gràcies for being optimistic when I wasn't, mercès for picking me up when I couldn't myself, and merci for joining me in this wild ride that is life.

I cannot forget thanking my family, **Moeke**, **Sjakie** and **Am**. The amount of unconditional(!) support over the years, especially from my mother, have been truly invaluable. I am not sure how to repay this 'debt', but I hope you being able to read through a ~150 page thesis on quantum communication is a good start.

This thesis is dedicated to **Sean Malone**, a man of singular artistry, dedication and kindness. A man who saw further than most, and suffered because of that. Thank you for the inspiration you've provided all the hours, days, and years so far, and especially for those moments that are yet to come. I wish I could have shared this work with you — I like to think it would've made both of us proud. Rest in peace Sean, we were all lucky to have been alive while you were making music.

BIBLIOGRAPHY

[1] Mohamed H. Abobeih, Julia Cramer, Michiel A. Bakker, Norbert. Kalb, Matthew. Markham, Daniel J. Twitchen, and Tim Hugo. Taminiau. One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment. *Nature Communications*, 9(1):2552, 2018.

[2] Mikael Afzelius, Christoph Simon, Hugues De Riedmatten, and Nicolas Gisin. Multimode quantum memory based on atomic frequency combs. *Physical Review A*, 79(5):052329, 2009.

[3] Emil Artin. *Geometric algebra*, chapter 6, pages 143–147. Interscience Publishers New York, 1 edition, 1957.

[4] Koji Azuma, Kiyoshi Tamaki, and Hoi-Kwong Lo. All-photonic quantum repeaters. *Nature communications*, 6, 2015.

[5] Manuel A Ballester, Stephanie Wehner, and Andreas Winter. State discrimination with post-measurement information. *IEEE Transactions on Information Theory*, 54(9):4183–4198, 2008.

[6] Bhaskar Roy Bardhan and Mark M Wilde. Strong converse rates for classical communication over thermal and additive noise bosonic channels. *Physical Review A*, 89(2):022302, 2014.

[7] Sean D. Barrett and Pieter Kok. Efficient high-fidelity quantum computation using matter qubits and linear optics. *Physical Review A*, 71(6):060310, 2005.

[8] Normand J Beaudry, Tobias Moroder, and Norbert Lütkenhaus. Squashing models for optical measurements in quantum communication. *Physical review letters*, 101(9):093601, 2008.

[9] Charles H Bennett, Herbert J Bernstein, Sandu Popescu, and Benjamin Schumacher. Concentrating partial entanglement by local operations. *Physical Review A*, 53(4):2046, 1996.

[10] Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *International Conference on Computer System and Signal Processing, IEEE, 1984*, pages 175–179, 1984.

[11] Charles H Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A Smolin, and William K Wootters. Purification of noisy entanglement and faithful teleportation via noisy channels. *Physical review letters*, 76(5):722, 1996.

[12] Charles H. Bennett, David P. Divincenzo, John A. Smolin, and William K. Wootters. Mixed-state entanglement and quantum error correction. *Physical Review A*, 54(5):3824–3851, Jan 1996.

[13] Hannes Bernien, Bas Hensen, Wolfgang Pfaff, Gerwin Koolstra, MS Blok, Lucio Robledo, TH Taminiau, Matthew Markham, DJ Twitchen, Lilian Childress, et al. Heralded entanglement between solid-state qubits separated by three metres. *Nature*, 497(7447):86, 2013.

[14] Machiel Blok, Norbert Kalb, Andreas Reiserer, Tim Taminiau, and Ronald Hanson. Towards quantum networks of single spins: analysis of a quantum memory with an optical interface in diamond. *Faraday discussions*, 184:173–182, 2015.

[15] Matthias Bock, Pascal Eich, Stephan Kucera, Matthias Kreis, Andreas Lenhard, Christoph Becher, and Jürgen Eschner. High-fidelity entanglement between a trapped ion and a telecom photon via quantum frequency conversion. *Nature communications*, 9(1):1–7, 2018.

[16] Stefan Bogdanovic, Suzanne B van Dam, Cristian Bonato, Lisanne C Coenen, AJ Zwerver, Bas Hensen, Madelaine SZ Liddy, Thomas Fink, Andreas Reiserer, Marko Loncar, and Ronald Hanson. Design and low-temperature characterization of a tunable microcavity for diamond-based quantum networks. *Applied Physics Letters*, 110(17):171103, 2017.

[17] Sebastiaan Brand, Tim Coopmans, and David Elkouss. Efficient computation of the waiting time and fidelity in quantum repeater chains. *IEEE Journal on Selected Areas in Communications (2020).*, 2020.

[18] Gilles Brassard, Anne Broadbent, Joseph Fitzsimons, Sébastien Gambs, and Alain Tapp. Anonymous quantum communication. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 460–473. Springer, 2007.

[19] Sergey Bravyi and Dmitri Maslov. Hadamard-free circuits expose the structure of the Clifford group. *arXiv preprint arXiv:2003.09412v1 [quant-ph]*, March 2020.

[20] Hans J. Briegel, Wolfgang Dür, Juan I Cirac, and Peter Zoller. Quantum repeaters: the role of imperfect local operations in quantum communication. *Physical Review Letters*, 81(26):5932, 1998.

[21] Dagmar Bruß. Optimal eavesdropping in quantum cryptography with six states. *Physical Review Letters*, 81(14):3018, 1998.

[22] Guido Burkard, Hans-Andreas Engel, and Daniel Loss. Spintronics and quantum dots for quantum computing and quantum communication. *Fortschritte der Physik: Progress of Physics*, 48(9-11):965–986, 2000.

[23] Carlos Cabrillo, Juan Ignacio Cirac, Priscila Garcia-Fernandez, and Peter Zoller. Creation of entangled states of distant atoms by interference. *Physical Review A*, 59(2):1025, 1999.

[24] Earl T Campbell and Simon C Benjamin. Measurement-based entanglement under conditions of extreme photon loss. *Physical review letters*, 101(13):130502, 2008.

[25] Jiu-Peng Chen, Chi Zhang, Yang Liu, Cong Jiang, Wei-Jun Zhang, Zhi-Yong Han, Shi-Zhao Ma, Xiao-Long Hu, Yu-Huai Li, Hui Liu, et al. Twin-field quantum key distribution over a 511 km optical fibre linking two distant metropolitan areas. *Nature Photonics*, pages 1–6, 2021.

[26] Jiu-Peng Chen, Chi Zhang, Yang Liu, Cong Jiang, Weijun Zhang, Xiao-Long Hu, Jian-Yu Guan, Zong-Wen Yu, Hai Xu, Jin Lin, et al. Sending-or-not-sending with independent lasers: Secure twin-field quantum key distribution over 509 km. *Physical review letters*, 124(7):070501, 2020.

[27] Yan Chen, Michael Zopf, Robert Keil, Fei Ding, and Oliver G Schmidt. Highly-efficient extraction of entangled photons from quantum dots using a broadband optical antenna. *Nature communications*, 9(1):1–7, 2018.

[28] Lilian Childress and Ronald Hanson. Diamond NV centers for quantum computing and quantum networks. *MRS bulletin*, 38(2):134–138, 2013.

[29] Lilian Childress, Jacob Taylor, Anders Søndberg Sørensen, and Mikhail Lukin. Fault-tolerant quantum communication based on solid-state photon emitters. *Physical review letters*, 96(7):070504, 2006.

[30] Lilian Childress, Jacob Taylor, Anders Søndberg Sørensen, and Mikhail D Lukin. Fault-tolerant quantum repeaters with minimal physical resources and implementations based on single-photon emitters. *Physical Review A*, 72(5):052330, 2005.

[31] Chin-Wen Chou, Hugues De Riedmatten, Daniel Felinto, Sergey V Polyakov, Steven J Van Enk, and Jeff Kimble. Measurement-induced entanglement for excitation stored in remote atomic ensembles. *Nature*, 438(7069):828, 2005.

[32] Matthias Christandl and Alexander Müller-Hermes. Relative entropy bounds on quantum, private and repeater capacities. *Communications in Mathematical Physics*, 353(2):821–852, 2017.

[33] Matthias Christandl and Stephanie Wehner. Quantum anonymous transmissions. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 217–235. Springer, 2005.

[34] Isaac L Chuang, Debbie W Leung, and Yoshihisa Yamamoto. Bosonic quantum codes for amplitude damping. *Physical Review A*, 56(2):1114, 1997.

[35] Julia Cramer, Norbert Kalb, M Adriaan Rol, Bas Hensen, Machiel S Blok, Matthew Markham, Daniel J Twitchen, Ronald Hanson, and Tim H Taminiau. Repeated quantum error correction on a continuously encoded qubit by real-time feedback. *Nature communications*, 7:11526, 2016.

[36] Gijs De Lange, Zhi-Hui Wang, Diego Riste, Viatcheslav Dobrovitski, and Ronald Hanson. Universal dynamical decoupling of a single solid-state spin from a spin bath. *Science*, 330(6000):60–63, 2010.

[37] Jeroen Dehaene and Bart De Moor. Clifford group, stabilizer states, and linear and quadratic operations over GF(2). *Physical Review A*, 68(4):042318, 2003.

[38] Jeroen Dehaene, Maarten Van den Nest, Bart De Moor, and Frank Verstraete. Local permutations of products of Bell states and entanglement distillation. *Physical Review A*, 67(2), February 2003.

[39] Jeroen Dehaene, Maarten Van den Nest, Bart De Moor, and Frank Verstraete. Local permutations of products of Bell states and entanglement distillation. *Physical Review A*, 67(2):022310, 2003.

[40] Aymeric Delteil, Zhe Sun, Wei-bo Gao, Emre Togan, Stefan Faelt, and Ataç Imamoğlu. Generation of heralded entanglement between distant hole spins. *Nature Physics*, 12(3):218, 2016.

[41] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Physical review letters*, 77(13):2818, 1996.

[42] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Physical Review Letters*, 77(13):2818–2821, September 1996.

[43] Marcus William Doherty, Neil Manson, Paul Delaney, Fedor Jelezko, Jörg Wrachtrup, and Lloyd Hollenberg. The nitrogen-vacancy colour centre in diamond. *Physics Reports*, 528(1):1–45, 2013.

[44] Luming Duan and Jeff Kimble. Scalable photonic quantum computation through cavity-assisted interactions. *Physical review letters*, 92(12):127902, 2004.

[45] Luming Duan, Mikhail Lukin, Juan Ignacio Cirac, and Peter Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–418, 2001.

[46] Wolfgang Dür and Hans Briegel. Entanglement purification for quantum computation. *Physical review letters*, 90(6):067901, 2003.

[47] Wolfgang Dür and Hans J Briegel. Entanglement purification and quantum error correction. *Reports on Progress in Physics*, 70(8):1381, 2007.

[48] Wolfgang Dür, Hans J. Briegel, Juan Ignacio Cirac, and Peter Zoller. Quantum repeaters based on entanglement purification. *Physical Review A*, 59(1):169, 1999.

[49] Artur K Ekert. Quantum cryptography based on Bell's theorem. *Physical review letters*, 67(6):661, 1991.

[50] Fabian Ewert and Peter van Loock. 3/4-efficient Bell measurement with passive linear optics and unentangled ancillae. *Physical review letters*, 113(14):140403, 2014.

[51] Kun Fang, Xin Wang, Marco Tomamichel, and Runyao Duan. Non-asymptotic entanglement distillation. *IEEE Transactions on Information Theory*, 65(10):6454–6465, 2019.

[52] Xiao-Tian Fang, Pei Zeng, Hui Liu, Mi Zou, Weijie Wu, Yan-Lin Tang, Ying-Jie Sheng, Yao Xiang, Weijun Zhang, Hao Li, et al. Implementation of quantum key distribution surpassing the linear rate-transmittance bound. *Nature Photonics*, 14(7):422–425, 2020.

[53] Mark Fox. *Quantum optics: an introduction*, volume 15. Oxford University Press, 2006.

[54] Keisuke Fujii and Katsuji Yamamoto. Entanglement purification with double selection. *Physical Review A*, 80(4):042308, 2009.

[55] Wei-bo Gao, Atac Imamoglu, Hannes Bernien, and Ronald Hanson. Coherent manipulation, measurement and entanglement of individual solid-state spins using optical fields. *Nature Photonics*, 9(6):363, 2015.

[56] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum-enhanced positioning and clock synchronization. *Nature*, 412(6845):417–419, 2001.

[57] Oleg Gittsovich, Normand J Beaudry, Varun Narasimhachar, R Romero Alvarez, Tobias Moroder, and Norbert Lütkenhaus. Squashing model for detectors and applications to quantum-key-distribution protocols. *Physical Review A*, 89(1):012325, 2014.

[58] Kenneth Goodenough. https://github.com/kdgoodenough/repeaterchainoptimisation.

[59] Kenneth Goodenough, David Elkouss, and Stephanie Wehner. Assessing the performance of quantum repeaters for all phase-insensitive gaussian bosonic channels. *New Journal of Physics*, 18(6):063005, 2016.

[60] Daniel Gottesman. Theory of fault-tolerant quantum computation. *Physical Review A*, 57(1):127, 1998.

[61] Daniel Gottesman and Hoi-Kwong Lo. Proof of security of quantum key distribution with two-way classical communications. *IEEE Transactions on Information Theory*, 49(2):457–475, 2003.

[62] Warren P Grice. Arbitrarily complete Bell-state measurement using only linear optical elements. *Physical Review A*, 84(4):042331, 2011.

[63] Saikat Guha, Hari Krovi, Christopher A Fuchs, Zachary Dutton, Joshua A Slater, Christoph Simon, and Wolfgang Tittel. Rate-loss analysis of an efficient quantum repeater architecture. *Physical Review A*, 92(2):022357, 2015.

[64] Ronald Hanson, Leo P Kouwenhoven, Jason R Petta, Seigo Tarucha, and Lieven MK Vandersypen. Spins in few-electron quantum dots. *Reviews of modern physics*, 79(4):1217, 2007.

[65] Bas Hensen, Hannes Bernien, Anaïs E Dréau, Andreas Reiserer, Norbert Kalb, Machiel S Blok, Just Ruitenberg, Raymond FL Vermeulen, Raymond N Schouten, Carlos Abellán, et al. Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682, 2015.

[66] Bas Hensen, Norbert Kalb, MS Blok, AE Dréau, Andreas Reiserer, RFL Vermeulen, RN Schouten, M Markham, DJ Twitchen, Kenneth Goodenough, et al. Loophole-free Bell test using electron spins in diamond: second experiment and additional analysis. *Scientific reports*, 6:30289, 2016.

[67] SLN Hermans, M Pompili, HKC Beukers, S Baier, J Borregaard, and R Hanson. Qubit teleportation between non-neighboring nodes in a quantum network. *arXiv preprint arXiv:2110.11373*, 2021.

[68] Adrian Holzäpfel, Jean Etesse, Krzysztof T Kaczmarek, Alexey Tiranov, Nicolas Gisin, and Mikael Afzelius. Optical storage on the timescale of a second in a solid-state atomic frequency comb memory using dynamical decoupling. *arXiv preprint arXiv:1910.08009*, 2019.

[69] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of modern physics*, 81(2):865, 2009.

[70] Daniel Huber, Marcus Reindl, Johannes Aberl, Armando Rastelli, and Rinaldo Trotta. Semiconductor quantum dots as an ideal source of polarization-entangled photon pairs on-demand: a review. *Journal of Optics*, 20(7):073002, 2018.

[71] David Hucul, Ismail Volkan Inlek, Grahame Vittorini, Clayton Crocker, Shantanu Debnath, SM Clark, and Christopher Monroe. Modular entanglement of atomic qubits using photons and phonons. *Nature Physics*, 11(1):37–42, 2015.

[72] Peter C Humphreys, Norbert Kalb, Jaco PJ Morits, Raymond N Schouten, Raymond FL Vermeulen, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. Deterministic delivery of remote entanglement on a quantum network. *Nature*, 558(7709):268, 2018.

[73] Ismail Volkan Inlek, Clayton Crocker, Martin Lichtman, Ksenia Sosnova, and Christopher Monroe. Multispecies trapped-ion node for quantum networking. *Physical review letters*, 118(25):250502, 2017.

[74] Solomon Ivan, Krishna Kumar Sabapathy, and Rajiah Simon. Operator-sum representation for bosonic Gaussian channels. *Physical Review A*, 84(4):042311, 2011.

[75] S Rao Jammalamadaka and Ambar Sengupta. *Topics in circular statistics*, volume 5. World Scientific, 2001.

[76] Sarah Jansen. Enumerating distillation protocols Github repository. https://github.com/sarahjansen08/enumerating-distillation-protocols. Accessed: 2021-02-04.

[77] Liang Jiang, Jacob M Taylor, Navin Khaneja, and Mikhail D Lukin. Optimal approach to quantum communication using dynamic programming. *Proceedings of the National Academy of Sciences*, 104(44):17291–17296, 2007.

[78] Liang Jiang, Jacob M Taylor, Kae Nemoto, William J Munro, Rodney Van Meter, and Mikhail D Lukin. Quantum repeater with encoding. *Physical Review A*, 79(3):032325, 2009.

[79] Cody Jones, Danny Kim, Matthew T Rakher, Paul G Kwiat, and Thaddeus D Ladd. Design and analysis of communication protocols for quantum repeater networks. *New Journal of Physics*, 18(8):083015, 2016.

[80] Richard Jozsa, Daniel S Abrams, Jonathan P Dowling, and Colin P Williams. Quantum clock synchronization based on shared prior entanglement. *Physical Review Letters*, 85(9):2010, 2000.

[81] Norbert Kalb, Peter C Humphreys, Jesse Slim, and Ronald Hanson. Dephasing mechanisms of diamond-based nuclear-spin memories for quantum networks. *Physical Review A*, 97(6):062330, 2018.

[82] Norbert Kalb, Andreas Reiserer, Stephan Ritter, and Gerhard Rempe. Heralded storage of a photonic quantum bit in a single atom. *Physical review letters*, 114(22):220501, 2015.

[83] Norbert Kalb, Andreas A Reiserer, Peter C Humphreys, Jacob JW Bakermans, Sten J Kamerling, Naomi H Nickerson, Simon C Benjamin, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. Entanglement distillation between solid-state quantum network nodes. *Science*, 356(6341):928–932, 2017.

[84] Gerd Keiser. *Optical Fiber Communications*. McGraw-Hill, 4th edition, 2011.

[85] Aeysha Khalique and Barry C Sanders. Practical long-distance quantum key distribution through concatenated entanglement swapping with parametric down-conversion sources. *arXiv preprint arXiv:1501.03317*, 2015.

[86] Pieter Kok and Samuel L Braunstein. Postselected versus nonpostselected quantum teleportation using parametric down-conversion. *Physical Review A*, 61(4):042304, 2000.

[87] Stefan Krastanov, Victor V Albert, and Liang Jiang. Optimized entanglement purification. *Quantum*, 3:123, 2019.

[88] Marko Krčo and Prabasaj Paul. Quantum clock synchronization: Multiparty protocol. *Physical Review A*, 66(2):024305, 2002.

[89] Hari Krovi, Saikat Guha, Zachary Dutton, Joshua A Slater, Christoph Simon, et al. Practical quantum repeaters with parametric down-conversion sources. *arXiv preprint arXiv:1505.03470*, 2015.

[90] Hari Krovi, Saikat Guha, Zachary Dutton, Joshua A Slater, Christoph Simon, and Wolfgang Tittel. Practical quantum repeaters with parametric down-conversion sources. *Applied Physics B*, 122(3):52, 2016.

[91] Seung-Woo Lee and Hyunseok Jeong. Near-deterministic quantum teleportation and resource-efficient quantum computation using linear optics and hybrid qubits. *Physical Review A*, 87(2):022326, 2013.

[92] Seung-Woo Lee, Kimin Park, Timothy C Ralph, and Hyunseok Jeong. Nearly deterministic Bell measurement for multiphoton qubits and its application to quantum information processing. *Physical review letters*, 114(11):113603, 2015.

[93] Seung-Woo Lee, Kimin Park, Timothy C Ralph, and Hyunseok Jeong. Nearly deterministic Bell measurement with multiphoton entanglement for efficient quantum-information processing. *Physical Review A*, 92(5):052324, 2015.

[94] Yang Liu, Zong-Wen Yu, Weijun Zhang, Jian-Yu Guan, Jiu-Peng Chen, Chi Zhang, Xiao-Long Hu, Hao Li, Cong Jiang, Jin Lin, et al. Experimental twin-field quantum key distribution through sending or not sending. *Physical Review Letters*, 123(10):100505, 2019.

[95] Hoi-Kwong Lo, Hoi Fung Chau, and Mohammed Ardehali. Efficient quantum key distribution scheme and a proof of its unconditional security. *Journal of Cryptology*, 18(2):133–165, 2005.

[96] Hoi-Kwong Lo, Marcos Curty, and Bing Qi. Measurement-device-independent quantum key distribution. *Physical Review Letters*, 108(13):130503, 2012.

[97] Nicolo Lo Piparo, Neil. Sinclair, and Mohsen Razavi. Memory-assisted quantum key distribution resilient against multiple-excitation effects. *arXiv preprint arXiv:1707.07814*, 2017.

[98] Marco Lucamarini, Zhiliang L Yuan, James F Dynes, and Andrew J Shields. Overcoming the rate–distance limit of quantum key distribution without quantum repeaters. *Nature*, 557(7705):400, 2018.

[99] David Luong, Liang Jiang, Jungsang Kim, and Norbert Lütkenhaus. Overcoming lossy channel bounds using a single quantum repeater node. *arXiv preprint arXiv:1508.02811*, 2015.

[100] Norbert Lütkenhaus, John Calsamiglia, and Kalle-Antti Suominen. Bell measurements for teleportation. *Physical Review A*, 59(5):3295, 1999.

[101] Alexander I Lvovsky, Barry C Sanders, and Wolfgang Tittel. Optical quantum memory. *Nature Photonics*, 3(12):706–714, 2009.

[102] Xiongfeng Ma, Pei Zeng, and Hongyi Zhou. Phase-matching quantum key distribution. *Physical Review X*, 8(3):031043, 2018.

[103] Dmitri Maslov and Martin Roetteler. Shorter stabilizer circuits via bruhat decomposition and quantum circuit transformations. *IEEE Transactions on Information Theory*, 64(7):4729–4738, July 2018.

[104] Dzmitry Matsukevich, Thierry Chaneliere, Mishkatul Bhattacharya, Shau-Yu Lan, Stewart Jenkins, Brian Kennedy, and Alex Kuzmich. Entanglement of a photon and a collective atomic excitation. *Physical review letters*, 95(4):040405, 2005.

[105] Takaaki Matsuo, Clément Durand, and Rodney Van Meter. Quantum link bootstrapping using a RuleSet-based communication protocol. *Physical Review A*, 100(5):052320, 2019.

[106] Peter Maunz, David Moehring, Steven Olmschenk, Kelly Cooper Younge, Dzmitry Matsukevich, and Christopher Monroe. Quantum interference of photon pairs from two remote trapped atomic ions. *Nature Physics*, 3(8):538, 2007.

[107] Peter Christian Maurer, Georg Kucsko, Christian Latta, Liang Jiang, Norman Ying Yao, Steven D Bennett, Fernando Pastawski, David Hunger, Nicholas Chisholm, Matthew Markham, et al. Room-temperature quantum bit memory exceeding one second. *Science*, 336(6086):1283–1286, 2012.

[108] Mariella Minder, Mirko Pittaluga, George Lloyd Roberts, Marco Lucamarini, James Dynes, Zhiliang Yuan, and Andrew J Shields. Experimental quantum key distribution beyond the repeaterless secret key capacity. *Nature Photonics*, 13(5):334–338, 2019.

[109] Christopher Monroe and Jungsang Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.

[110] William Munro, KA Harrison, Ashley Stephens, Simon Devitt, and Kae Nemoto. From quantum multiplexing to high-performance quantum networking. *Nature Photonics*, 4(11):792, 2010.

[111] William J Munro, Koji Azuma, Kiyoshi Tamaki, and Kae Nemoto. Inside quantum repeaters. *Selected Topics in Quantum Electronics, IEEE Journal of*, 21(3):1–13, 2015.

[112] William J Munro, Ashley Stephens, Simon J Devitt, KA Harrison, and Kae Nemoto. Quantum communication without the necessity of quantum memories. *Nature Photonics*, 6(11):777–781, 2012.

[113] Sreraman Muralidharan, Linshu Li, Jungsang Kim, Norbert Lütkenhaus, Mikhail D Lukin, and Liang Jiang. Optimal architectures for long distance quantum communication. *Scientific reports*, 6, 2016.

[114] Gláucia Murta, Filip Rozpędek, Jérémy Ribeiro, David Elkouss, and Stephanie Wehner. Key rates for quantum key distribution protocols with asymmetric noise. *Physical Review A*, 101(6):062321, 2020.

[115] Kae Nemoto, Michael Trupke, Simon J Devitt, Burkhard Scharfenberger, Kathrin Buczak, Jörg Schmiedmayer, and William J Munro. Photonic quantum networks formed from NV-centers. *Scientific reports*, 6:26284, 2016.

[116] Christian Nguyen, Denis Sukachev, Mihir Bhaskar, Bartholomeus Machielse, David Levonian, Erik Knall, Pavel Stroganov, Cleaven Chia, Michael Burek, Ralf Riedinger, et al. An integrated nanophotonic quantum register based on silicon-vacancy spins in diamond. *Physical Review B*, 100(16):165428, 2019.

[117] Christian Nguyen, Denis Sukachev, Mihir Bhaskar, Bartholomeus Machielse, David Levonian, Erik Knall, Pavel Stroganov, Ralf Riedinger, Hongkun Park, Marko Lončar, et al. Quantum network nodes based on diamond qubits with an efficient nanophotonic interface. *Physical review letters*, 123(18):183602, 2019.

[118] Naomi H Nickerson, Joseph F Fitzsimons, and Simon C Benjamin. Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links. *Physical Review X*, 4(4):041041, 2014.

[119] Naomi H Nickerson, Ying Li, and Simon C Benjamin. Topological quantum computing with a very noisy network and local error rates approaching one percent. *Nature communications*, 4(1):1–5, 2013.

[120] Andrea Olivo and Frédéric Grosshans. Investigating the optimality of ancilla-assisted linear optical Bell measurements. *Phys. Rev. A 98, 042323*, 2018.

[121] Christiana Panayi, Mohsen Razavi, Xiongfeng Ma, and Norbert Lütkenhaus. Memory-assisted measurement-device-independent quantum key distribution. *New Journal of Physics*, 16(4):043005, 2014.

[122] Mihir Pant, Hari Krovi, Dirk Englund, and Saikat Guha. Rate-distance tradeoff and resource costs for all-optical quantum repeaters. *Physical Review A*, 95(1):012304, 2017.

[123] Mihir Pant, Hari Krovi, Don Towsley, Leandros Tassiulas, Liang Jiang, Prithwish Basu, Dirk Englund, and Saikat Guha. Routing entanglement in the quantum internet. *npj Quantum Information*, 5(1):25, 2019.

[124] James L Park. The concept of transition in quantum mechanics. *Foundations of Physics*, 1(1):23–33, 1970.

[125] Rüdiger Paschotta. *Encyclopedia of Laser Physics and Technology*. John Wiley & Sons, Hoboken, New Jersey, December 2008.

[126] Wolfgang Pfaff, BJ Hensen, Hannes Bernien, Suzanne B van Dam, Machiel S Blok, Tim H Taminiau, Marijn J Tiggelman, Raymond N Schouten, Matthew Markham, Daniel J Twitchen, et al. Unconditional quantum teleportation between distant solid-state quantum bits. *Science*, 345(6196):532–535, 2014.

[127] Nicoló Lo Piparo, Mohsen Razavi, and William J Munro. Measurement-device-independent quantum key distribution with nitrogen vacancy centers in diamond. *Physical Review A*, 95(2):022338, 2017.

[128] Nicolo Lo Piparo, Mohsen Razavi, and William J Munro. Memory-assisted quantum key distribution with a single nitrogen vacancy center. *arXiv preprint arXiv:1708.06532*, 2017.

[129] Stefano Pirandola. Capacities of repeater-assisted quantum communications. *arXiv preprint arXiv:1601.00966*, 2016.

[130] Stefano Pirandola and Riccardo Laurenza. General benchmarks for quantum repeaters. *arXiv preprint arXiv:1512.04945*, 2015.

[131] Stefano Pirandola, Riccardo Laurenza, Carlo Ottaviani, and Leonardo Banchi. Fundamental limits of repeaterless quantum communications. *Nature communications*, 10 2015.

[132] Matteo Pompili, Sophie LN Hermans, Simon Baier, Hans KC Beukers, Peter C Humphreys, Raymond N Schouten, Raymond FL Vermeulen, Marijn J Tiggelman, Laura dos Santos Martins, Bas Dirkse, et al. Realization of a multinode quantum network of remote solid-state qubits. *Science*, 372(6539):259–264, 2021.

[133] John Preskill. Quantum clock synchronization and quantum error correction. *arXiv preprint quant-ph/0010098*, 2000.

[134] Mohsen Razavi and Jeffrey H Shapiro. Long-distance quantum communication with neutral atoms. *Physical Review A*, 73(4):042303, 2006.

[135] Andreas Reiserer, Norbert Kalb, Machiel S Blok, Koen JM van Bemmelen, Tim H Taminiau, Ronald Hanson, Daniel J Twitchen, and Matthew Markham. Robust quantum-network memory using decoherence-protected subspaces of nuclear spins. *Physical Review X*, 6(2):021040, 2016.

[136] Andreas Reiserer and Gerhard Rempe. Cavity-based quantum networks with single atoms and optical photons. *Reviews of Modern Physics*, 87(4):1379, 2015.

[137] Renato Renner. ETH Zurich PhD thesis. *arXiv preprint quant-ph/0512258*, 2005.

[138] Daniel Riedel, Immo Söllner, Brendan J Shields, Sebastian Starosielec, Patrick Appel, Elke Neu, Patrick Maletinsky, and Richard J Warburton. Deterministic enhancement of coherent photon generation from a nitrogen-vacancy center in ultrapure diamond. *Physical Review X*, 7(3):031040, 2017.

[139] Lucio Robledo, Lilian Childress, Hannes Bernien, Bas Hensen, Paul F. A. Alkemade, and Ronald Hanson. High-fidelity projective read-out of a solid-state spin quantum register. *Nature*, 477(7366):574–578, 2011.

[140] Filip Rozpędek, Kenneth Goodenough, Jérémy Ribeiro, Norbert Kalb, Valentina Caprara Vivoli, Andreas Reiserer, Ronald Hanson, Stephanie Wehner, and David Elkouss. Parameter regimes for a single sequential quantum repeater. *Quantum Science and Technology*, 3(3):034002, 2018.

[141] Filip Rozpędek, Thomas Schiet, David Elkouss, Andrew C Doherty, Stephanie Wehner, et al. Optimizing practical entanglement distillation. *Physical Review A*, 97(6):062333, 2018.

[142] Filip Rozpędek, Thomas Schiet, Le Phuc Thinh, David Elkouss, Andrew C. Doherty, and Stephanie Wehner. Optimizing practical entanglement distillation. *Physical Review A*, 97(6), June 2018.

[143] Filip Rozpędek, Raja Yehia, Kenneth Goodenough, Maximilian Ruf, Peter C Humphreys, Ronald Hanson, Stephanie Wehner, and David Elkouss. Near-term quantum repeater experiments with NV centers: overcoming the limitations of direct transmission. *Physical Review A 99.5 (2019): 052330.*, 2018.

[144] Liangzhong Ruan, Wenhan Dai, and Moe Z. Win. Adaptive recurrence quantum entanglement distillation for two-kraus-operator channels. *Phys. Rev. A*, 97:052332, May 2018.

[145] Erhan Saglamyurek, Neil Sinclair, Jeongwan Jin, Joshua A Slater, Daniel Oblak, Félix Bussieres, Mathew George, Raimund Ricken, Wolfgang Sohler, and Wolfgang Tittel. Broadband waveguide quantum memory for entangled photons. *Nature*, 469(7331):512, 2011.

[146] Nicolas Sangouard, Christoph Simon, Hugues De Riedmatten, and Nicolas Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33, 2011.

[147] Siddhartha Santra, Liang Jiang, and Vladimir S Malinovsky. Quantum repeater architecture with hierarchically optimized memory buffer times. *Quantum Science and Technology 10.1088/2058-9565/ab0bc2*, 2019.

[148] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Reviews of Modern Physics*, 81(3):1301, 2009.

[149] Eddie Schoute, Laura Mančinska, Tanvirul Islam, Iordanis Kerenidis, and Stephanie Wehner. Shortcuts to quantum network routing. *arXiv preprint arXiv:1610.05238*, 2016.

[150] Evgeny Shchukin, Ferdinand Schmidt, and Peter van Loock. On the waiting time in quantum repeaters with probabilistic entanglement swapping. *Phys. Rev. A 100, 032322*, 2019.

[151] James D Siverns, John Hannegan, and Qudsia Quraishi. Neutral-atom wavelength-compatible 780 nm single photons from a trapped ion via quantum frequency conversion. *Physical Review Applied*, 11(1):014044, 2019.

[152] Holger P Specht, Christian Nölleke, Andreas Reiserer, Manuel Uphoff, Eden Figueroa, Stephan Ritter, and Gerhard Rempe. A single-atom quantum memory. *Nature*, 473(7346):190–193, 2011.

[153] Timothy P Spiller, Kae Nemoto, Samuel L Braunstein, William J Munro, Peter van Loock, and Gerard J Milburn. Quantum computation by communication. *New Journal of Physics*, 8(2):30, 2006.

[154] Robert Stockill, Megan Stanley, Lukas Huthmacher, Edmund Clarke, Maxim Hugues, Aaron Miller, Clemens Matthiesen, Claire Le Gall, and Mete Atatüre. Phase-tuned entangled state generation between distant spin qubits. *Physical Review Letters*, 119(1):010503, 2017.

[155] Masahiro Takeoka, Saikat Guha, and Mark M Wilde. Fundamental rate-loss tradeoff for optical quantum key distribution. *Nature communications*, 5, 2014.

[156] Masahiro Takeoka, Saikat Guha, and Mark M Wilde. Squashed entanglement and the two-way assisted capacities of a quantum channel. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 326–330. IEEE, 2014.

[157] Kiyoshi Tamaki, Hoi-Kwong Lo, Wenyuan Wang, and Marco Lucamarini. Information theoretic security of quantum key distribution overcoming the repeaterless secret key capacity bound. *arXiv preprint arXiv:1805.05511*, 2018.

[158] Tim Hugo Taminiau, Julia Cramer, Toeno van der Sar, Viatcheslav V Dobrovitski, and Ronald Hanson. Universal control and error correction in multi-qubit spin registers in diamond. *Nature nanotechnology*, 9(3):171, 2014.

[159] Yoshiaki Tamura, Hirotaka Sakuma, Keisei Morita, Masato Suzuki, Yoshinori Yamamoto, Kensaku Shimada, Yuya Honma, Kazuyuki Sohma, Takashi Fujii, and Takemi Hasegawa. Lowest-ever 0.1419-db/km loss optical fiber. In *Optical Fiber Communication Conference*, pages Th5D–1. Optical Society of America, 2017.

[160] Charles W. Thiel, Thomas Böttger, and Rufus Cone. Rare-earth-doped materials for applications in quantum information storage and signal processing. *Journal of luminescence*, 131(3):353–361, 2011.

[161] Nuala Timoney, Imam Usmani, Pierre Jobez, Mikael Afzelius, and Nicolas Gisin. Single-photon-level optical storage in a solid-state spin-wave memory. *Physical Review A*, 88(2):022324, 2013.

[162] Emre Togan, Yiwen Chu, Alexei Trifonov, Liang Jiang, Jeronimo Maze, Lilian Childress, Gurudev Dutt, Anders Søndberg Sørensen, Philip Hemmer, Alexander Zibrov, et al. Quantum entanglement between an optical photon and a solid-state spin qubit. *Nature*, 466(7307):730–734, 2010.

[163] Kazuya Tsurumoto, Ryota Kuroiwa, Hiroki Kano, Yuhei Sekiguchi, and Hideo Kosaka. Quantum teleportation-based state transfer of photon polarization into a carbon spin in diamond. *Communications Physics*, 2(1):1–6, 2019.

[164] Lev Vaidman and Nadav Yoran. Methods for reliable teleportation. *Physical Review A*, 59(1):116, 1999.

[165] Suzanne B van Dam, Peter C Humphreys, Filip Rozpędek, Stephanie Wehner, and Ronald Hanson. Multiplexed entanglement generation over quantum networks using multi-qubit nodes. *Quantum Science and Technology*, 2(3):034002, 2017.

[166] Rodney Van Meter, Thaddeus D Ladd, William J Munro, and Kae Nemoto. System design for a long-line quantum repeater. *IEEE/ACM Transactions on Networking (TON)*, 17(3):1002–1013, 2009.

[167] Michelle Victora, Stefan Krastanov, Alexander Sanchez de la Cerda, Steven Willis, and Prineha Narang. Purification and entanglement routing on quantum networks. *arXiv preprint arXiv:2011.11644*, 2020.

[168] Scott E Vinay and Pieter Kok. Statistical analysis of quantum-entangled-network generation. *Physical Review A*, 99(4):042313, 2019.

[169] Karl Gerd H Vollbrecht and Frank Verstraete. Interpolation of recurrence and hashing entanglement distillation protocols. *Physical Review A*, 71(6):062325, 2005.

[170] Shuang Wang, De-Yong He, Zhen-Qiang Yin, Feng-Yu Lu, Chao-Han Cui, Wei Chen, Zheng Zhou, Guang-Can Guo, and Zheng-Fu Han. Beating the fundamental rate-distance limit in a proof-of-principle quantum key distribution system. *Physical Review X*, 9(2):021046, 2019.

[171] Shun Watanabe, Ryutaroh Matsumoto, Tomohiko Uyematsu, and Yasuhito Kawano. Key rate of quantum key distribution with hashed two-way classical communication. *Physical Review A*, 76(3):032312, 2007.

[172] Christian Weedbrook, Stefano Pirandola, Raul Garcia-Patron, Nicolas J Cerf, Timothy C Ralph, Jeffrey H Shapiro, and Seth Lloyd. Gaussian quantum information. *Reviews of Modern Physics*, 84(2):621, 2012.

[173] Stephanie Wehner, David Elkouss, and Ronald Hanson. Quantum internet: A vision for the road ahead. *Science*, 362(6412):eaam9288, 2018.

[174] Stephan Welte, Bastian Hacker, Severin Daiss, Stephan Ritter, and Gerhard Rempe. Photon-mediated quantum gate between two neutral atoms in an optical cavity. *Physical Review X*, 8(1):011018, 2018.

[175] Mark M Wilde. *Quantum information theory*. Cambridge University Press, 2013.

[176] Mark M Wilde and Haoyu Qi. Energy-constrained private and quantum capacities of quantum channels. *arXiv preprint arXiv:1609.01997*, 2016.

[177] Raja Yehia. Available on demand.

[178] Hussain A Zaidi and Peter van Loock. Beating the one-half limit of ancilla-free linear optics Bell measurements. *Physical review letters*, 110(26):260501, 2013.

[179] Robert Andrzej Żak, Beat Röthlisberger, Stefano Chesi, and Daniel Loss. Quantum computing with electron spins in quantum dots. *arXiv preprint arXiv:0906.4045*, 2009.

[180] Sebastian Zaske, Andreas Lenhard, Christian A Keßler, Jan Kettler, Christian Hepp, Carsten Arend, Roland Albrecht, Wolfgang-Michael Schulz, Michael Jetter, Peter Michler, et al. Visible-to-telecom quantum frequency conversion of light from a single quantum emitter. *Physical Review Letters*, 109(14):147404, 2012.

[181] Xuanqiang Zhao, Benchi Zhao, Zihe Wang, Zhixin Song, and Xin Wang. LOCCNet: a machine learning framework for distributed quantum information processing. *arXiv preprint arXiv:2101.12190*, 2021.

CURRICULUM VITÆ

Kenneth DUDLEY GOODENOUGH

26-07-1992 Born in Kempton Park, South-Africa.

EDUCATION

2009–2013 Bachelor Physics
The Hague University of Applied Sciences
Delft, Netherlands

2013–2014 Bridging programme in Physics
Technische Universiteit Delft

2015-2016 Master in Applied Physics & Casimir Pre-PhD Track
Kavli Institute of Nanoscience & Leiden Institute of Physics (LION)

2017-2021 PhD. Quantum Information
Technische Universiteit Delft, QuTech
Thesis: Distributing entanglement in quantum networks
Promotor: Prof. dr. S. D. C. Wehner

List of Publications

Thesis papers

4. S. Jansen*, K. Goodenough*, S. De Bone, D. Gijswijt, and D. Elkouss *Enumerating all bilocal Clifford distillation protocols through symmetry reduction*, arXiv:2103.03669 (2021) (**Shared first author**)

3. K. Goodenough, D. Elkouss, and S. Wehner *Optimizing repeater schemes for the quantum internet*, (2021) Physical Review A 103, 032610 (2021)

2. 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 (5), 052330 (2019) (**Shared first author**)

1. F. Rozpędek*, K. Goodenough*, J. Ribeiro, N. Kalb, V. Caprara Vivoli, A. Reiserer, R. Hanson, S. Wehner and D. Elkouss, *Parameter regimes for a single sequential quantum repeater*, Quantum Science and Technology 3 (3), 034002 (2018) (**Shared first author**)

Previous publications

5. S. de Bone; R. Ouyang, K. Goodenough, D. Elkouss, *Protocols for Creating and Distilling Multipartite GHZ States With Bell Pairs*, IEEE Transactions on Quantum Engineering (2020)

4. T. P. W. Cope, K. Goodenough, S. Pirandola, *Converse bounds for quantum and private communication over Holevo-Werner channels*, Journal of Physics A: Mathematical and Theoretical 51 (49), 494001 (2018)

3. M. Jerger, Y. Reshitnyk, M. Oppliger, A. Potočnik, M. Mondal, A. Wallraff, K. Goodenough, S. Wehner, K. Juliusson, N. K. Langford, A. Fedorov, *Contextuality without nonlocality in a superconducting quantum system*, Nature communications 7 (1), 1-6 (2016)

2. B. Hensen, N. Kalb, M. S. Blok, A. E. Dréau, A. Reiserer, R. F. L. Vermeulen, R. N. Schouten, M. Markham, D. J. Twitchen, K. Goodenough, D. Elkouss, S. Wehner, T. H. Taminiau, R. Hanson, *Loophole-free Bell test using electron spins in diamond: second experiment and additional analysis*, Scientific reports 6, 30289 (2016)

1. K. Goodenough, D. Elkouss, S. Wehner, *Assessing the performance of quantum repeaters for all phase-insensitive Gaussian bosonic channels*, New Journal of Physics 18 (6), 063005 (2016)