Simulation Model for Atomic Ensemble based Quantum Repeaters and the Optimization of their Positioning

---

Master Thesis Applied Physics

Julian Rabbie   4307720

January 22, 2020

COMMITTEE MEMBERS: PROF. DR. S. WEHNER (SUPERVISOR)
PROF. DR. G. STEELE
DR. D. ELKOUSS

Contents

1 Introduction

The European Union has recognised the need for investing in quantum technologies with a 1 billion euro investment embodied in the Quantum Flagship [1] as part of Horizon 2020 [2]. One of the 20 projects that are part of the Quantum Flagship is the Quantum Internet Alliance (QIA) [3], which has the goal of forming a quantum network in which any two points in Europe can be connected via quantum communication. In the future this network could even be extended to a worldwide quantum internet, in which any two parties can be connected across the globe, similar to the classical internet.

A quantum internet [4] will allow for communication via qubits, the quantum mechanical analogy of regular (classical) bits. Its currently known applications range from clock synchronization [5], blind quantum computing [6] and an alternative to the global positiong system [7] to quantum key distribution [8]. The latter uses the laws of quantum mechanics to encrypt messages. Even though it is still intractable to break classical encryption such as the Rivest–Shamir–Adleman cryptosystem [9] with quantum computers [10], Google's recent quantum supremacy experiment [11] has been a big step forward towards this security liability nonetheless. Furthermore, with initiatives such as the recently held first ever pan-European quantum internet hackathon, new exciting applications will surely be found in the near future.

As with many emerging research projects, the initial results are focused to be obtained on a smaller scale. Since QIA has only started in October of 2018, one of the main goals for the coming years is to realize a quantum internet across a few nodes, where the Netherlands is chosen as test bed due to its manageable size. While experimental physicists are working on the many challenging parts of realizing quantum communication in general, a diverse team of theoretical physicists, mathematicians and software engineers, the NLBlueprint team, is focusing on creating a detailed and highly modular simulation model of this Dutch quantum network. This will allow for the investigation of which parameters are required for reliable quantum communication, such that experimental physicists can be supplied with concrete near-term goals to work towards. Additionally, it will allow for looking ahead by running simulations of complex, large-scale and even hybrid systems in order to pinpoint what is needed to realize a European-wide quantum internet.

During this thesis project, we have participated in the NLBlueprint team and contributed mostly to the simulations of quantum repeater chains based on atomic ensembles. An essential component of a quantum repeater is quantum memory, of which numerous types exist and are developed all across Europe. In this thesis we will give an overview of two types of atomic ensemble based quantum memories and discuss their current experimental status. Additionally, we will use an existing analytical model to investigate the performance of both types of memories within quantum repeater chains.

On the other hand, regardless of which type of memory is used, another part of this thesis is focused on finding the optimal positioning of repeater nodes within a quantum network. In particular we will focus on the suitable mathematical formulations for this problem and use these in order to investigate the optimal repeater positioning in realistic network topologies.

The exciting research related to building a quantum internet covers a wide range of disciplines. This is also reflected in this thesis, which includes subjects from quantum information science, quantum optics, software engineering and discrete optimization. The focus has therefore also been on breadth by combining the knowledge of these different fields into new models.

In a nutshell, this thesis contains

- The description of the versatile and elaborate simulation model we created for simulating multiplexed quantum repeater chains based on atomic ensembles. This can be used to test the effects of a large number of parameters, such as the number of repeaters and end-to-end distance as well as those related to detectors, memories and photon sources. We validate this model by comparing it to an analytical one, after which we analyze the effect of additional sources of noise that cannot be taken into account in the analytical model.
- A detailed theoretical overview as well as the experimental status quo of two quantum memory protocols: atomic frequency comb and electromagnetically induced transparency. We furthermore compare the two memory types by analyzing their performance when embedded in a repeater chain with the use of an existing analytical model.
- Two linear programming formulations for the optimization of the placement of quantum repeaters within a two-dimensional network. We highlight the strengths and limits of these formulations and demonstrate their use by finding the optimal repeater allocation combined with the optimal communication paths on a European-scale network topology.

This thesis is structured as follows. We will start with an overview of one-dimensional quantum networks in Section 2. After a brief overview of the building blocks of a quantum repeater chain, a frequency multiplexed protocol and the application we use to measure its performance, we continue with a detailed theoretical description of two types of atomic ensemble based quantum memory, which is an essential component of quantum repeaters. Next, in Section 3 we will describe our simulation model and in particular emphasize on the parts on which we contributed. We will also show some simulation results as well as a characterization of the two memory types. Then, in Section 4, we will describe the optimization of the positioning of quantum repeaters in a two dimensional setting. Here we give two types of binary linear programming formulations which include multiple constraints and objective functions. This is all concluded in Section 5, in which we also propose some future research topics.

2 Quantum Networks

In order realise a quantum internet, two parties will need to have a reliable way of sharing entanglement. As entanglement is relatively easy to create locally, the next step would be to determine which medium is best to use for transferring such an entangled state from one party to another. The main challenge is here that the fragile quantum state should be protected from noise and loss as much as possible during its travel. The most commonly used way of transferring entanglement is therefore by the use of photons in optical fiber. An additional benefit of photons is that we have the freedom to encode our qubit in its numerous available degrees of freedom, such as its polarization or arrival time. Even though these massless particles hardly interact with the environment in optical fiber, the loss over distance is still exponential. In particular, the probability that a photon reaches a distance of L km, also called the tansmittance η_t, is given by

$$\eta_t = 10^{-\alpha L/10}, \tag{2.1}$$

where α is the attenuation coefficient in dB/km. At the optimal transmission wavelength of 1550 nm, the attenuation coefficient is approximately 0.2 dB/km in high quality optical fiber, which results in a $99\,\%$ loss of the initial signal intensity over a distance of 100 km. Therefore this direct transmission quickly becomes inadequate to use in a European-scale quantum internet. In the classical internet, where photons are also used to transmit information, signal boosters can be used at intermediate distances to enhance the signal quality in order to overcome this exponential loss. However, the no-cloning theorem [12] forbids the copying of an arbitrary quantum state, such that we have to resort to another type of signal enhancement: the quantum repeater.

We will start in Section 2.1 with an abstract and idealized description of quantum repeaters where we do not yet consider any imperfections or noise to simplify the explanation. Next, we will describe the benefits of multiplexing and present the main protocol for our simulation model as well as the performance measure we use. Then, in Section 2.2, we continue with a detailed theoretical overview of two types of quantum memory (QM), which is an essential component of a quantum repeater. For both types we give a theoretical description including the multiplexing capabilities and discuss the experimental status quo.

2.1 Quantum Repeaters

The first abstract quantum repeater setup was described in 1998 by Briegel et al. [13]. A few years later Duan, Lukin, Cirac and Zoller (DLCZ) realized that all the requirements for this setup could be met by using atomic ensembles and linear optics [14]. We will give a brief description of this protocol where we focus on the abstract building blocks rather than the underlying physical implementation.

The basic idea of the DLCZ protocol is that in order to create entanglement over a long distance L, one can divide this distance into smaller segments, referred to as *elementary links*. An elementary link consists of two nodes which both have access to a photon source and a QM. The key requirement of the latter is to allow for a reversible mapping between light and matter. In the DLCZ protocol, the source and memory are the same entity. However, it was later realized that this could be emulated by separating these. In that case we have an entangled-photon pair source (PPS) of which one of the photons is stored in a QM, while the other one is used to generate remote entanglement [15]. For the remainder of this section, we assume that each node has access to the latter type of setup.

Entanglement generation in an elementary link proceeds as follows. Two nodes A and B, separated by a distance L, simultaneously trigger their PPS which emits an Einstein Podolsky Rosen (EPR) pair

$$\left|\Phi^{+}\right\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle + |11\rangle\right), \tag{2.2}$$

i.e. an equal superposition between emitting vacuum and a single photon pair. They store one half of the entangled-photon pair in their own local QMs and send the other photon over to a remote station, positioned in the center at a distance of $L/2$. The remote station interferes incoming photons with a $50/50$ beam splitter and measures them with two independent detectors, see Figure 2.1a). The detection of exactly one photon at the remote station projects the remaining quantum state onto the Bell state

$$\left|\Psi^{\pm}\right\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle_A |1\rangle_B \pm |1\rangle_A |0\rangle_B\right), \tag{2.3}$$

where $|0\rangle_A$ $(|1\rangle_B)$ represents the absence (presence) of a matter qubit in the memory at A (B) and the phase depends on which of the two detectors clicked. The remote station hence performs a Bell state measurement (BSM) and we will henceforth refer to it as the BSM station. Once a successful measurement is performed, a classical message is sent back to both A and B, heralding their entanglement. The BSM is inherently probabilistic and succeeds with a $50\,\%$ probability. When two or no photons are detected at the BSM station, the entanglement generation was unsuccessful and the two nodes try again. Note that the simultaneous arrival of two photons at the beam splitter can result in a single detector click due to photon bunching, which we will explain in Section 2.1.1.

Now consider the case in which we have two elementary links AB and CD, in which entanglement has successfully been established. By converting the matter qubits at B and C back to photons and performing another BSM with these photons, we can teleport the entanglement over to the memories at A and D. This procedure is called an *entanglement swap* and again only succeeds when a single photon is registered, see Figure 2.1b).

By combining elementary links and entanglement swaps, one can create entanglement over a large distance in the procedure that is illustrated in Figures 2.1c) to 2.1e). This is called a

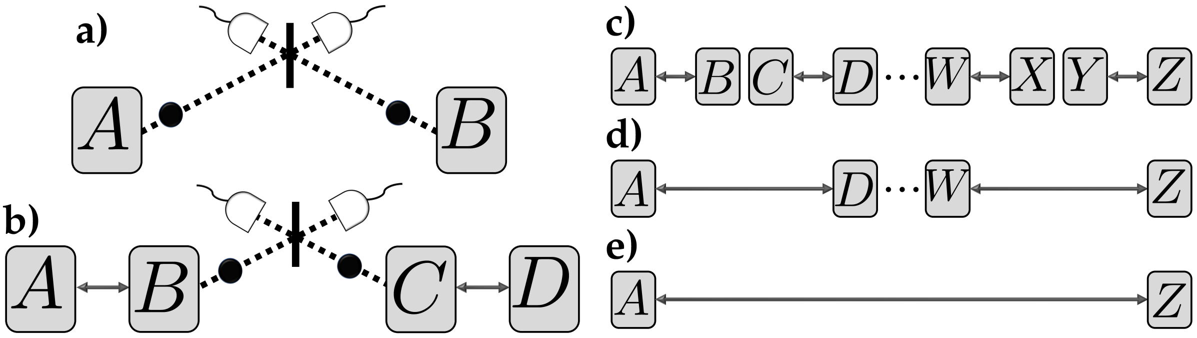

Figure 2.1: Schematic illustration of long distance entanglement generation with quantum repeaters. **a)** An elementary link, where nodes A and B generate entanglement by sending a photon (circle) over to a remote station positioned half way between them. Here, a beam splitter (thick vertical bar) and two photon detectors are used to perform a BSM, and its result is sent back to both nodes via classical communication. **b)** An entanglement swap can be performed in a similar fashion, in which a BSM is performed with the photons that were stored at nodes B and C, which are part of two separate elementary links AB and CD. **c)** Entanglement can be distributed over a long distance by first generating entanglement between the elementary links AB, CD, WX, etc. **d)** Once entanglement has been generated in two neighbouring elementary links, the first level of entanglement swaps is performed. **e)** After the final level of swaps, entanglement has been generated between the end nodes A and Z. All figures are similar to those in [16].

(one dimensional) quantum repeater chain. At first, remote entanglement is created between pairs of nodes by creating elementary links. When two neighbouring elementary links have successfully established entanglement, we can perform an entanglement swap to teleport the entanglement to the two outer nodes. This procedure can be repeated at each nesting level until we have created long range entanglement between two end nodes A and Z. When we assume that the remote BSM stations of elementary links are always placed exactly in the middle of two nodes, the use of $N = 2^n$ elementary links separated by a distance $L_0 = L/N$ requires us to send photons over a distance of just $L_0/2$, hence significantly reducing the exponential transmission loss η_t. When a photon is lost nonetheless, we can simply try to generate entanglement again while our QM stores the entanglement of a neighbouring elementary link.

So far we have considered a highly idealized setting in which all of our components were completely free of noise and errors. However, in reality this is definitely not the case. Not only do the photon detectors and QMs have a limited efficiency, but one should also consider the photon indistinguishability at the BSM station and the probability for a detector to register a dark count, i.e. a click in a detection window when no photon is present. The QMs additionally have a limited storage time, such that they will not be able to store quantum states for arbitrary long periods of time while they are waiting for two neighbouring elementary links to be successful, and last but not least, a very large source of noise arises from the emission of multiple-photon pairs from the PPS.

Altogether, these imperfections strongly deteriorate the entanglement generation rate as well as the fidelity of the quantum state that the end nodes share. The performance of quantum repeaters with the DLCZ protocol is therefore relatively poor [16]. A lot of improvements were proposed over the years and a very promising one was found in the use of multiplexing, which we will describe next.

Multiplexing

The success probability P_{el} of successfully generating entanglement in an elementary link depends on the single-photon pair emission probability, detector efficiency, detector dark count probability, the photon indistinguishability, the distance spanned by the elementary link and the attenuation coefficient. The success probability is therefore generally very small. This implies that it can take a great amount of time to establish entanglement, as this process has to be repeated a lot of rounds until a classical message is sent back by the BSM station that heralds success. We define a round as the process of resetting the memory and the source, sending a photon to the BSM station and receiving back the classical message. The time it takes to complete a round is bounded from below by $L_0/\bar{c}$, where $\bar{c}$ is the reduced speed of light in optical fiber of approximately $2 \cdot 10^5$ km/s.

Multiplexing alleviates the problem of having a small success probability by attempting to generate entanglement multiple times in a single round. This can be done by using a degree of freedom of the photons that is not used to generate entanglement, such as a temporal, spatial or spectral mode. With temporal multiplexing, a train of photons is sent per round rather than a single one, and with spatial multiplexing the repeater chain is for example copied many times in parallel. However, in this thesis we will focus on spectral multiplexing, which is incorporated in the quantum repeater protocol of Tittel et al. [17]. In this case the PPS emits photons in M distinct frequency modes. The BSM should then be able to measure all the modes independently and transmit back which of the frequency modes, if any, yielded a successful joint measurement. We therefore only require for a single mode to be successful, such that the overall success probability for an elementary link increases to

$$P_{\mathrm{el}}^{\mathrm{mp}} = 1 - (1 - P_{\mathrm{el}})^M. \tag{2.4}$$

When a sufficiently large number of modes are used such that the success probability reaches unity, we can assume that each round is successful. This implies that we can apply all the entanglement swaps simultaneously and hence we directly go from the situation in Figure 2.1c) to Figure 2.1e). Additionally, this greatly reduces the requisite storage time of the QMs to approximately $L_0/\bar{c}$, instead of them necessarily being forced to store their entanglement until the neighbouring link is successful, which is required in hierarchical quantum repeater protocols [16].

Photon Indistinguishability

As we have seen, the BSMs are a vital part of quantum repeater chains. When we have perfect photon detectors, the probability that a BSM succeeds is $50~\%$ if we are restricted to using linear optics. This can in principle be brought arbitrarily close to $100~\%$ with the use of ancilla photons [18], but this is very challenging to achieve experimentally. However, even when we can be sure that our BSM succeeds, the quality of our BSM is also determined by the extent at which our photons are completely indistinguishable.

When two photons arrive at the beam splitter they should be identical in all degrees of freedom, i.e. frequency, spectral shape, arrival time, phase and polarization, in order to delete any 'which-way' information. If this is not the case, one could (partially) determine where the photon came from, which destroys the entanglement. In other words, the projection on the Bell state $|\Psi^{\pm}\rangle$ will be noisy.

Consider the setup in Figure 2.2a) in which two photons arrive at the beam splitter from input modes a and b, with properties j and k respectively. These properties denote the values of all the degrees of freedom of the photon. The beam splitter acts as a unitary operation on the input state and in case of a perfect reflectivity of $1/2$, it holds that [19]

$$\hat{a}_j^\dagger |0\rangle \rightarrow \frac{1}{\sqrt{2}}\left(\hat{c}_j^\dagger + \hat{d}_j^\dagger\right)|0\rangle \tag{2.5}$$

$$\hat{b}_k^\dagger |0\rangle \rightarrow \frac{1}{\sqrt{2}}\left(\hat{c}_k^\dagger - \hat{d}_k^\dagger\right)|0\rangle\,, \tag{2.6}$$

where $\hat{x}_y^\dagger$ denotes the creation operator of a photon in mode x with properties y. The output state after this operation is

$$|\psi\rangle_{\text{out}} = \frac{1}{2}\left(\hat{c}_j^\dagger\hat{c}_k^\dagger + \hat{c}_k^\dagger\hat{d}_j^\dagger - \hat{c}_j^\dagger\hat{d}_k^\dagger - \hat{d}_j^\dagger\hat{d}_j^\dagger\right)|00\rangle\,. \tag{2.7}$$

If the properties j are not the same as k, we can see that the coincidence probability P_{coin} of getting a detection in both detectors is $1/2$. However, when the photons are perfectly indistinguishable such that $j = k$, the second and third terms cancel in (2.7), which implies that $P_{\text{coin}} = 0$. The photons thus bunch and both traverse to one of the two detectors, which is called the Hong-Ou-Mandel (HOM) effect [20].

In order to make the HOM effect more concrete, let us consider the case in which the arrival time of the photons is the only relevant property, while all other degrees of freedom are exactly equal. If we assume that the photon in mode b is slightly delayed by τ_{del} relative to the one in a, we can plot the coincidence probability versus this delay as shown in Figure 2.2b). Here we see that for $\tau_d = 0$ the coincidence probability drops to 0, which is known as the HOM dip. However, when the two incoming photons are not exactly identical, we can see that this dip is only partially visible. The visibility V of the HOM dip is defined by [19]

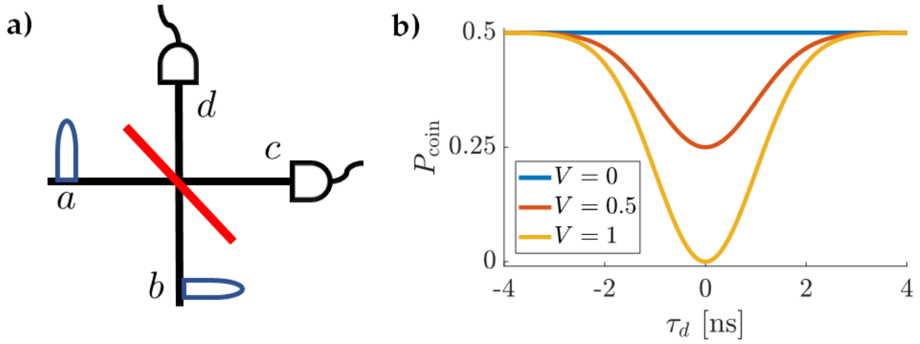

Figure 2.2: Hong-Ou-Mandel (HOM) effect. **a)** The setup in which the HOM effect occurs requires two incoming photons in modes a and b that interfere at a beam splitter and get detected afterwards in modes c and d. **b)** The coincidence probability P_{coin} of getting a detection in both c and d as a function of the delay between the arrival time of the two photons τ_d. If this delay is zero and the photons are also equal in all other degrees of freedom, the photons become indistinguishable which results in the HOM dip. If the photons are not exactly the same in all degrees of freedom, we cannot (fully) see this dip, which results in a visibility V of less than 1.

$$V = 1 - 2\min\left(P_{\text{coin}}\right). \tag{2.8}$$

In the classical limit of $V = 0$, the photons become perfectly distinguishable and the HOM dip disappears. This destroys the quantum correlations required for our application (see Section 2.1.2).

Note that regular threshold photon detectors will register two incoming photons as a single click. They can therefore not distinguish the case in which the photons have bunched from the case in which a single photon arrived, which leads to false positives in terms of the BSM outcomes. The use of number resolving detectors can circumvent this, since these count the number of photons that arrive. In our simulation model we have included both types of detectors, see Section 3.1.

2.1.1 Frequency Multiplexed Protocol

Due to the very significant improvements of multiplexing, we will describe the frequency multiplexed protocol of Tittel et al. in more detail. It is illustrated in Figure 2.3 for two elementary links. Two parties A and B again both have access to a PPS and a separate QM. These are able to respectively emit and store photons in many distinct frequency modes, where the bandwidth of the source should be smaller than or equal to the bandwidth of the memory. We will give a brief description of the physics of a PPS that can be used in this protocol below, while we devote Section 2.2 to a more elaborate description of the QM. When the photons of A and B arrive at the remote BSM station, all the spectral modes are independently measured and the successful mode (if any) is reported back to the two parties. This process is executed simultaneously on the other elementary link CD. Next, all the photons

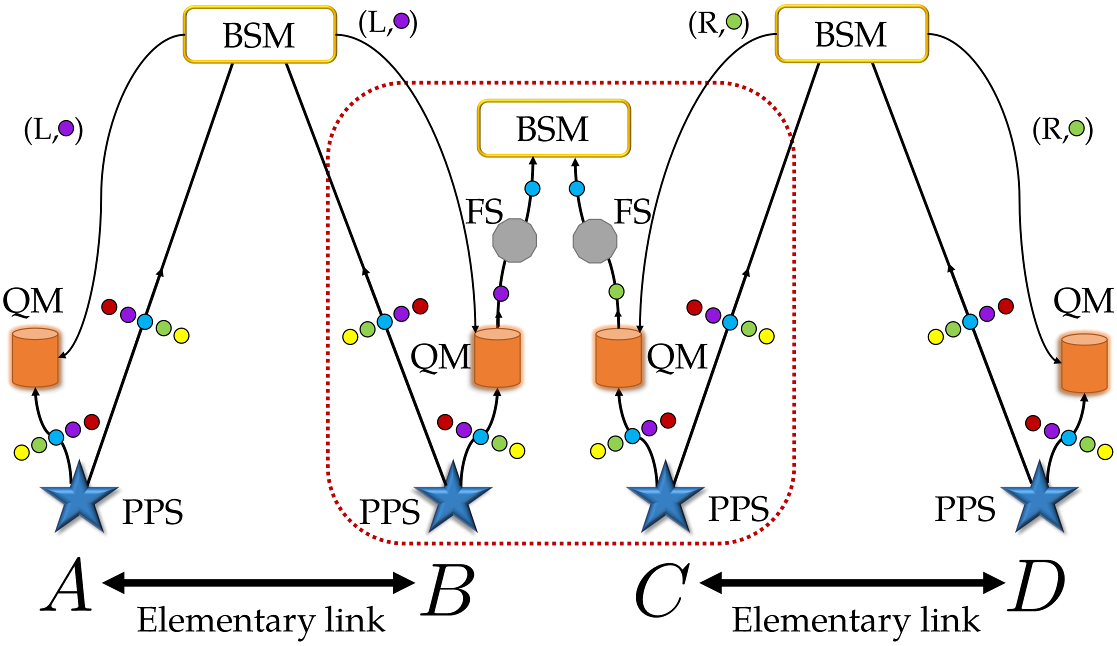

Figure 2.3: Illustration of a repeater chain with a single repeater formed by two independent elementary links AB and CD. On each side of this link, a photon pair source (PPS) emits a frequency multiplexed entangled-photon pair, of which one half is stored in the quantum memory (QM) and the other half is sent over to a remote Bell state measurement (BSM) station. Once the measurement is performed, the BSM station sends a classical message back to the nodes of the elementary link, which contains whether the left (L) or right (R) detector has clicked, as well as the frequency of the successful mode. These modes are then extracted from the memory and shifted to a common frequency by a frequency shifter (FS). The resulting photons are subsequently used to perform an entanglement swap with another BSM. The latter can be done locally, such that the nodes B and C can be contained at the same physical location called a repeater node, indicated by the red dotted line. A successful BSM on all three stations yields the entanglement of the QMs of A and D.

are extracted from the memories at B and C and spectrally filtered such that only the successful mode is ultimately transmitted to another BSM station for the entanglement swap. Note that (one of) the two photons also needs to be shifted to a common frequency in order to make them indistinguishable. The two parties B and C that perform the entanglement swap can be physically placed at the same location, such that their BSM can be performed locally. We therefore regard this as a single node, which we call a *repeater node*.

Photon Pair Source

The PPS that is used in this protocol is based on spontaneous parametric down-conversion (SPDC) [21]. This uses a combination of a non-linear crystal and a pump laser to emit an entangled-photon pair, which consists of a signal and an idler photon. There are three types of SPDC crystals, which are defined on the basis of the polarization of the emitted photons.

Since the wavelength at which the QM operates is not necessarily the same as the optimal wavelength of photons in optical fiber, we would like the signal and idler photon to be emitted at different wavelengths. This is possible by using a periodically poled lithium niobate (PPLN) crystal [22]. By carefully selecting the pump-photon frequency and changing the temperature of the crystal, one can ensure that one of the photons of the pair is emitted at around 1550 nm, the optimal wavelength for photons in optical fiber, while the other one has a frequency which is required for best storing photons in the QM.

In the DLCZ protocol, entanglement is generated by performing a BSM in which a single click heralds the entanglement of two QMs, at which point it is not possible to know where the photon is present. This is therefore referred to as *presence-absence* encoding. However, in the protocol of Tittel et al. the photon pairs are generated with a *time-bin* encoding. This implies that ideally the state which is generated locally by each node is $(\left|e\right\rangle_i \left|e\right\rangle_s + \left|l\right\rangle_i \left|l\right\rangle_s)/\sqrt{2}$, where $\left|e\right\rangle$ ($\left|l\right\rangle$) means the presence of a photon in the early (late) time bin and the subscript i (s) denotes the idler (signal) photon. This type of entanglement can be generated by pumping an SPDC crystal with a laser which has first passed through an unbalanced Mach-Zehnder interferometer, see Figure 2.4a). Note that the type of SPDC crystal must be chosen appropriately, since the polarization of the photons should be the same when they arrive at the BSM station in order for them to be completely indistinguishable. With time-bin encoding the BSM is successful if exactly one photon is detected in each time window, also called the double-click scheme. In this case both QMs of the elementary link store a single photon, in which case we cannot know which of the two has its photon stored in which time bin without measuring it. Faulty detectors cause significantly more errors in the double-click scheme, since we need to perform an additional BSM compared to the single-click scheme used for presence-absence encoding. However, due to the two-photon detection, we no longer require the optical fibers that traverse to the remote BSM to be phase stabilized [16], and with time-bin encoding it is easier to perform quantum key distribution (QKD), which we will explain below.

The probability $P_{\text{em}}^{\text{SPDC}}(n)$ of emitting a photon pair with n photons for an SPDC with a mean output photon number of μ can be derived from a squeezing Hamiltonian, which results in the distribution [21]

$$P_{\text{em}}^{\text{SPDC}}(n) = (n+1)\frac{\mu^n}{(\mu+1)^{(n+2)}}, \tag{2.9}$$

This distribution is plotted in Figure 2.4b), where we can see that the probability of generating a multiple-photon pair increases as μ increases. Multi-photon errors are detrimental for both the entanglement generation rate and the fidelity, which we will explain below. On the other hand, when we choose a very low mean photon number, we will emit no photons in most frequency modes, which also reduces our entanglement generation rate. In general, one therefore has to find a balance between fidelity and rate while choosing a suitable value for μ.

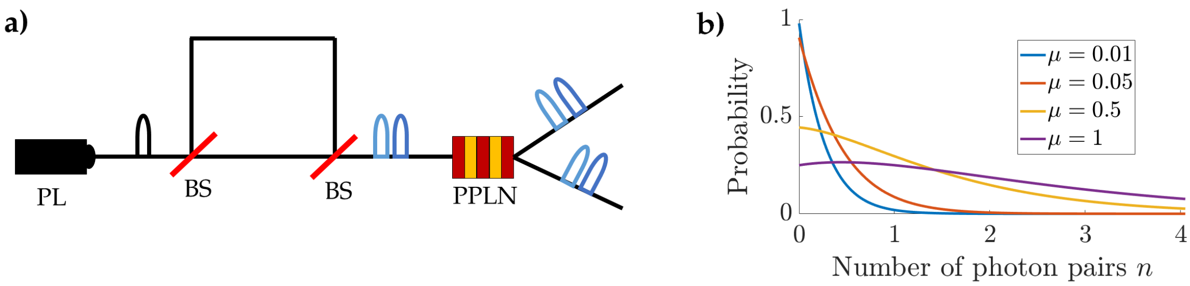

Figure 2.4: SPDC characteristics. **a)** A pump laser (PL) emits a photon, which traverses through an unbalanced Mach-Zehnder interferometer, consisting of two arms of different lengths and two beam splitters (BSs). This results in a photon that is in a superposition of an early and late time window, separated by a few ns in general. These photons then continue to the periodically poled lithium niobate (PPLN) crystal, which creates a time-bin encoded photon pair due to SPDC. **b)** Probability distribution function of an SPDC source. We see that a high value of μ increases the probability of generating the desired single-photon pair, but also increases the probability of emitting a multiple-photon pair, which is detrimental for the fidelity.

2.1.2 Application and Performance Measure

There are numerous applications for the bipartite entanglement that the end nodes aim to achieve. However, in this thesis we will only focus on QKD and the corresponding secret key rate, both of which are briefly covered in this subsection.

Quantum Key Distribution

With QKD, two end nodes Alice and Bob measure the photons that are not transmitted to the BSM station in a random basis, chosen from a set of orthogonal bases. In combination with classical communication, the two parties can then distill a secret key out of these measurements, which is inherently protected from any adversary by the laws of quantum mechanics. This is in contrast to prime factorization, a conjectured NP-complete problem from number theory which is widely used in for example RSA encryption [9] and can be deciphered in polynomial time with a quantum computer [10]. The rate at which this secret key can be established, the *secret key rate*, is an interesting figure of merit for repeater chains, since it combines the entanglement generation rate and fidelity of our end-to-end state in a single statistic, which we will explain below.

Finding the best performing and most secure QKD protocol in a certain scenario is a research field on its own. Here we will only focus on the BB84 protocol, devised by Benett and Brassard in 1984 [8]. In this protocol a secret key is constructed by performing measurements in two orthogonal bases, which we will call X and Z. In each round, Alice and Bob independently and uniformly at random pick one of these two bases in which they measure their photons. The corresponding measurement outcome is stored as a single classical bit. After a certain number of rounds have completed, they communicate their basis choices in each round via a classical channel. The secret key can then be extracted from the bits of the measurement outcomes of the rounds in which Alice and Bob measured in the same basis, after they have performed error correction and privacy amplification. A simple extension of BB84 is the six-state protocol [23], in which measurements in a third orthogonal Y basis are additionally performed. This protocol is more resilient against adversaries, at the cost of a lower secret key rate.

In case of time-bin encoding, measurements in the Z basis are performed by measuring the photons in both time bins separately. After a successful BSM, the entangled state stored in the memories of an elementary link will be projected onto $|\Psi^+\rangle$ or $|\Psi^-\rangle$, such that measurements in the Z basis should always be anti-correlated. This implies that if one party measures its photon in the early time bin, the other party should measure its photon in the late time bin. However, if we want to perform measurements in the X (and Y) basis, we will need to be able to interfere the two time bins. This can be done with an interferometer which has exactly the same difference in arm lengths as the one we used to create our time-bin encoded photons (see Figure 2.4a)). Measurements in the X basis can be both correlated or anti-correlated, depending on the BSM outcomes. The latter namely includes which of the two detectors clicked and thus whether a photon picked up a phase by getting reflected by the beam splitter or not. We will elaborate on the way we model the end node detectors in Section 3.1.

If we would use presence-absence encoding, the end nodes cannot perform measurements in more than one basis, since they do not have access to two photons which they can interfere. We would therefore need a whole second repeater chain to perform measurements in the other two bases [16], so with time-bin encoding it is much easier to perform QKD.

Secret Key Rate

The secret key rate represents the speed at which two end nodes can generate a secret key which they can use for secure communication. In case of the protocol of Tittel et al., the secret key rate S with BB84 when using N elementary links can be expressed by [24]

$$S = \frac{1}{2}\left(P_{\text{el}}^{\text{mp}}\right)^N \left(P_{\text{swap}}\right)^{N-1} P_{\text{meas}} R(Q) f_{\text{PPS}}. \tag{2.10}$$

Here P_{swap} is the probability of successfully performing an entanglement swap. Both $P_{\text{el}}^{\text{mp}}$ and P_{swap} depend on a lot of the parameters involved in the repeater chain and lies at the core of analyzing its performance. Next, P_{meas} denotes the probability that both end nodes get at least one detector click from their measurement of the end-to-end state, which can be used to generate a secret key. If an end node registers a double click, it uniformly at random assigns it to a single click in one of the two detectors. The factor $1/2$ represents the probability that Alice and Bob choose the same measurement basis. Based on the outcomes of the BSMs, the end nodes can derive which end-to-end state they expect and therefore whether their measurement result should be correlated or anti-correlated. However, noise such as multi-photon emissions can cause the wrong measurement outcome to be recorded in either the BSM stations or the end node detectors. This results in the probability Q_i that Alice and Bob obtain mismatched bits in their key when a measurement is performed in the basis $i \in \{Z, X\}$, also called the quantum bit error rate (QBER). This is an indirect measure of the fidelity of the end-to-end state. The function $R(Q) = 1 - H(Q_x) - H(Q_z)$, where $H(Q)$ is the binary entropy function, represents the number of useful key bits in case of the BB84 protocol. The binary entropy function is a measure of uncertainty in case of a Bernoulli trial, i.e. a random process with two possible outcomes. When the QBER is 0, we have a guarantee that the secret key is generated by a perfectly correlated random process, but a non-zero QBER leads to uncertainty. When $H(Q_x) + H(Q_z) \geq 1$ we are not able to generate a secret key, which occurs for example if $Q_x = Q_z \approx 0.11$ or if one of the two QBERs reaches $1/2$. The latter corresponds to the situation in which the key bits that are extracted from measurements in that basis are completely uncorrelated. In other words, Alice and Bob share the completely mixed state in this case. Finally, we should scale the number of extracted useful bits with the frequency of our photon source f_{PPS}, such that we get the rate in bits per second.

Multi-photon errors negatively affect both our entanglement generation rate and the fidelity of the end-to-end state. If multiple photons arrive at a BSM station and they traverse to two different detectors, the BSM will fail and we have to try again, hence lowering both P_{el} and P_{swap}. On the other hand, if the photons all go to the same detector and we do not use photon-number resolving detectors, the BSM will wrongfully register a success. When the photons are subsequently measured by interfering the two bins at the end node, they can get an incorrect measurement outcome such that the QBER in the X basis increases. It is thus crucial to keep the probability of emitting a multiple-photon pair as low as possible. Note that a non-unit visibility also increases the QBER in the X basis, since it causes a loss in phase coherence. If $V = 0$, the measurement outcomes in this basis will even become completely uncorrelated, which implies that $H(Q_x) = 1$ and we are thus not able to generate any secret key.

In order to derive analytical expressions for $P_{\mathrm{el}}^{\mathrm{mp}}$, P_{swap} and the QBER, one needs to know the quantum states at each point in time. It is possible to derive these under certain noisy settings, although it quickly becomes intractable. One of the goals of our simulation model is therefore to investigate additional sources of noise such as the introduction of multi-photon errors, beam splitter visibility or time-dependent retrieval loss in the quantum memory.

To analyze whether a given repeater setup performs better than simply using direct transmission of photons with a multi-mode optical fiber, we can compare their secret key rates. In particular, we will use a fundamental upper bound on the secret key rate achievable with lossy optical channels called the PLOB bound, after its authors Pirandola, Laurenza, Ottaviani and Banchi [25]. For M multiplexing modes this is given by

$$S_{\mathrm{PLOB}} = M \log_2 \left(\frac{1}{1 - \eta_t} \right) f_{\mathrm{PPS}}. \tag{2.11}$$

The goal of any repeater protocol is to surpass this upper bound, hence outperforming direct transmission. One could say that for distances at which this occurs, a repeater actually repeats.

2.2 Quantum Memories

Quantum memory is one of the essential components of most quantum repeaters. Just as any form of memory, it is used to store information and retrieve it at a later time, but in this case quantum information is stored. In multiplexed quantum repeater chains they are required to store an entangled quantum state until the remote BSMs have succeeded, after which the repeater can perform its entanglement swap. Since the entanglement is transmitted by using photons, we require optical QMs. The end nodes in principle need not to hold a QM in case of QKD, since they can measure their half of the entangled-photon pair directly [24]. However, for applications such as distributed quantum computing the end nodes would need memory. Note that there also exist all-photonic quantum repeater protocols that do not use any form of memory [26], but these require very large entangled photon states which makes these protocols experimentally very challenging.

In multiplexed quantum repeater protocols the QMs must be able to store a large number of modes in either time, space or frequency, or possibly a combination of these. Additionally we need to be able to retrieve the (arbitrary) successful mode at the midpoint from these memories with which we perform the entanglement swap at a repeater node. In this section we will analyze two possible types of atomic ensemble based QMs.

In general there are two types of memories that can be used in quantum repeater protocols: emissive and absorptive QMs. The first is used in the DLCZ protocol, where one atomic ensemble is used simultaneously as a photon source and a memory. The initial emission of a photon, called the Stokes photon, is used to perform the remote BSM and it simultaneously heralds the creation of a delocalized spin excitation within the ensemble. The memory can be read out at a later point in time by using a laser at a different frequency, which will cause the emission of an anti-Stokes photon that can be used for the local entanglement swap. On the other hand, absorptive quantum memories are separate entities that are coupled to a PPS. As the name suggests these memories absorb a photon, and re-emit it at a later time, which can either be after a fixed time duration or on-demand. Optical fiber can in principle be used as a type of absorptive quantum memory, although this suffers from an exponential decay of the signal due to transmittance as the required storage time, and hence the fiber length, increases [17].

Absorptive memories can in turn also be classified into two types, which are either based on optical control or engineered absorption [27]. The first requires a control laser to enable the storage and retrieval of photons, while the second uses an inhomogeneously broadened material combined with a photon-echo process. The use of a control laser makes the optically controlled memories inherently on-demand in terms of storage time, while the additional optical field also introduces noise to the retrieved signal. On the other hand, the storage time in memories based on engineered absorption and photon echo is in principle fixed, while the lack of a control field should result in relatively less noise. Furthermore the latter type of memory is generally more suitable for spectral and temporal multiplexing, while optically controlled memories are mostly limited to spatial multiplexing [28]. There also exist hybrid memory protocols that aim to combine the advantage of both types, of which we will give an example below.

Even though a wide range of atomic ensemble based QMs exist, in the remainder of this section we will limit our scope to two types of absorptive memories. We will start with the description of the memory type that is used in the protocol of Tittel et al. based on engineered absorption in Section 2.2.1 followed by an alternative absorptive memory based on optical control in Section 2.2.2. In both cases we will start with the theory behind the protocol after which we discuss the multiplexed storage capabilities and the experimental status quo. Hereafter, in Section 3.2.4 we will use an existing analytical model to compare the performance of these QMs in the context of quantum repeater chains.

2.2.1 Atomic Frequency Comb

The type of QM that is used in the protocol of Tittel et al. is based on atomic frequency comb (AFC) [29]. It requires an inhomogeneously broadened material, which naturally exists in rare earth ion (REI) doped crystals [30], such as thulium or erbium. These trivalent ions have the convenient property of being rather isolated systems, making them relatively protected from noise and externally controllable [31]. Due to local inhomogeneities such as strain, each ion will have a slightly distorted local environment, which causes the absorption frequency to shift. In total, this causes the material as a whole to have a broad absorption linewidth with a full width at half maximum (FWHM) of Γ as illustrated in Figure 2.5a). The AFC protocol relies on manually modifying this absorption spectrum in a particular way, after which a photon is absorbed and re-emitted after a fixed amount of time.

Theoretic Description

Consider an inhomogeneously broadened medium with spectral width Γ, which is split into equidistant peaks with FWHM Υ and separation distance Π as illustrated in Figure 2.5b). These splittings form a comb structure in the frequency domain, hence the name AFC. This structure is generally created by spectral hole burning [29], which optically pumps atoms from the ground state $|g\rangle$ to a long-lived metastable shelving state $|s\rangle$ in a three-level Λ-system as shown in Figure 2.6a). Note that the system is cooled to cryogenic temperature beforehand, such that we can assume that all atoms are initially in the ground state.

This modified medium is subsequently able to absorb a photon pulse with a FWHM that exceeds multiple comb widths Υ, but is smaller than the total inhomogeneous line width Γ, see Figure 2.5b), which is on resonance with the transition from the ground state $|g\rangle$ to the excited state $|e\rangle$. After the absorption of a photon, the atomic system at time $t = 0$ will be in a collective delocalized excited state, also called a Dicke state, given by

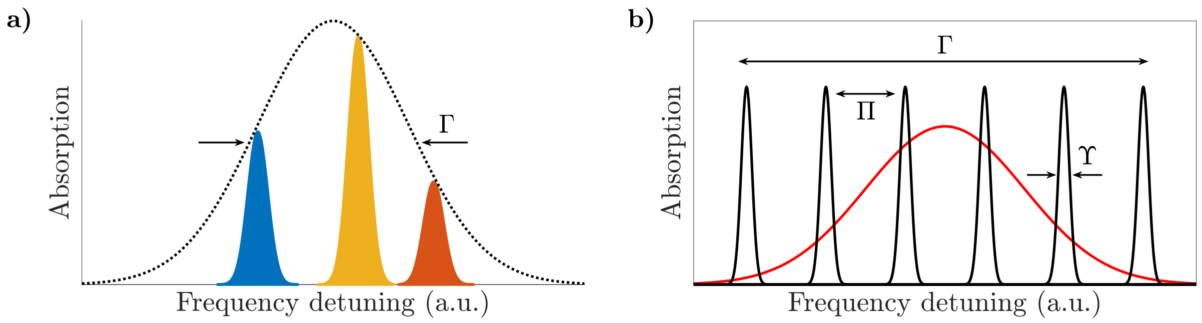

Figure 2.5: Illustration of the setup for the AFC protocol. **a)** An inhomogeneously broadened medium with a FWHM of Γ is created by a combination of homogeneously broadened absorption peaks, which are shifted due to changes in their local environment. **b)** Delocalized storage in which a probe photon (red envelope) is absorbed by an inhomogeneously broadened medium with spectral width Γ, which is split up into equidistant peaks with spacing Π that have a FWHM of ϒ. Both figures are replicas from [31].

$$\frac{1}{\sqrt{N}}\sum_{j=1}^{N} a_j e^{i2\pi\delta_j t} e^{-ikz_j} \left|g_1, g_2, \ldots, e_j, \ldots, g_N\right\rangle,$$

where a_j, δ_j and z_j are the frequency-dependent probability amplitude of the excitation, frequency detuning and position of atom j respectively, and k is the wave number of the input photon. This photon will then be absorbed by a large number of AFC modes, which are initially in phase, but will now rapidly start to dephase due to the absorption. This dephasing does not lead to a strong re-emission of the photon due to the relative detuning $\exp(i\delta_j t)$. If we assume that the comb peaks are approximately Dirac delta peaks, then δ_j is always an integer multiple of the intercomb spacing Π, i.e. $\delta_j = m\Pi$ where m is a positive integer. This results in the effect that after a fixed time $t = 1/\Pi$, the system will rephase and emit a photon while traversing back to the collective ground state. This process is illustrated in Figure 2.6b).

The retrieval efficiency of the AFC protocol depends on a number of parameters and is given by [31]

$$\eta_{\text{mem}}^{\text{AFC}} = \left(\frac{d}{F}\right)^2 e^{-\frac{d}{F}} e^{-\frac{7}{F^2}} e^{-d_0}, \tag{2.12}$$

where d is the optical depth of the medium, $F \equiv \Pi/\Upsilon$ is the comb finesse, and d_0 is the residual optical depth, caused by left-over atoms that are not correctly transferred to the shelving level after the spectral hole burning. The first two terms in (2.12) form a trade-off with respect to d/F and respectively arise due to collective re-emission and re-absorption in the atomic medium. The third term represents the loss due to dephasing, where a high value of F validates the approximation that $\delta_j = m\Pi$. Finally, imperfections in the AFC preparation process are captured in the term $\exp(-d_0)$, which leads to absorption of photons in valleys of the comb. Even when the latter term becomes negligible, i.e. $d_0 \to 0$, the maximum efficiency is still at most $54\ \%$, which implies that AFC memories would be of little practical use. Furthermore, the storage time is fixed and limited by the optical coherence lifetimes of REI doped crystals.

Luckily there are methods to resolve both of these issues. First, the efficiency can be boosted up to $100\ \%$ by forcing the photons to be re-emitted in the backward direction. In this case, the storage efficiency is modified to [29]

$$\tilde{\eta}^{\mathrm{AFC}}_{\mathrm{mem}} = \left(1 - e^{-\frac{d}{F}}\right)^2 e^{-\frac{7}{F^2}}, \tag{2.13}$$

from which we can see that for large values of d and F we can approach unit efficiency. Since d is generally a fixed experimental parameter, the storage efficiency can be maximized by carefully tailoring the medium for the optimal value of F, which forms a trade-off between absorption and dephasing.

The storage time in the AFC protocol can be increased by adding a fourth atomic state $|s'\rangle$ to our system, which is a long-lived spin-state that can for example be created by inducing a Zeeman splitting. After a photon is absorbed in the AFC medium, the atoms in the excited state can be transferred to this spin-state by an optical π-pulse resonant on the $|s'\rangle \leftrightarrow |e\rangle$ transition. If there is no spin-inhomogeneous broadening present in the material, the phase will be locked in this state [31]. Then, after an adjustable storage time T_s, the same resonant π-pulse in the backward direction can be used to continue the rephasing, after which the photon can be retrieved after a total time $1/\Pi + T_s$. This protocol, also referred to as the full AFC protocol, in principle leads to an on-demand quantum memory with a retrieval efficiency up to $100\ \%$. Note that this is an example of a hybrid protocol, since it combines an engineered absorption medium with optical control.

Alternatively, one can also increase the AFC efficiency by embedding the atomic ensemble in an asymmetric cavity [32]. When the impedance matching condition is attained, which implies that the cavity decay rate is equal to the absorption rate of the ensemble, one has the guarantee that a photon will be absorbed. This makes the storage efficiency independent of the optical depth, and approximately only the term $\exp\left(-7/F^2\right)$ in (2.12) remains, such that unit efficiency can be reached by solely increasing the finesse F. This cavity-enhanced AFC can furthermore still be combined with the on-demand storage by using the long-lived spin-state as we described above.

Multiplexing Capabilities and Experimental Status Quo

The main power of absorptive QMs based on AFC is that the number of possible storage modes is independent on the optical depth. It is only limited by the bandwidth of the inhomogeneous broadening of the REI and the pulse width of the entering photons, provided that the AFC itself can be perfectly prepared. Ignoring all the very complex experimental limitations, this is currently the only known memory protocol capable of storing the 10^6 modes referred to in [33], which is necessary for reaching a near-unit success probability for generating entanglement with an elementary link over large distances.

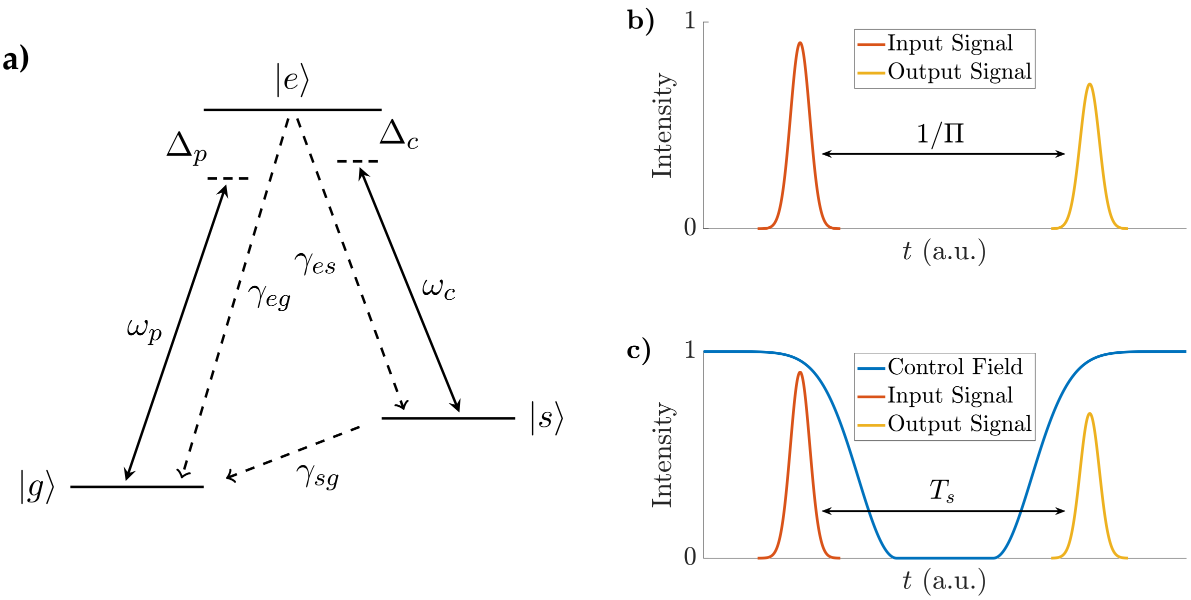

Figure 2.6: Atomic configuration and field intensities for AFC and EIT. **a)** Three-level Λ-system with a ground state $|g\rangle$, a slightly higher energetic metastable state $|s\rangle$ and a single excited state $|e\rangle$. For AFC the state $|s\rangle$ is used as a shelving level to store the atoms by spectral hole burning. In case of EIT an electromagnetic probe field with frequency ω_p and control field with frequency ω_c are detuned from their respective transitions with detunings Δ_p and Δ_c. A state $|i\rangle$ can decohere to a lower energetic state $|j\rangle$ with decay rate γ_{ij}. **b)** Illustration of the signal transmission in the AFC protocol. After a fixed time $1/\Pi$ the input signal is transmitted due to photon echo. **c)** Signal transmission for EIT, where an input pulse is trapped by adiabatically turning of the control beam and retrieved after an adjustable time T_s by the reverse process.

Unfortunately, up until this point the experimental realizations of AFC protocols are less optimistic. The maximum number of storage modes of approximately 130 is achieved in [34] by combining temporal and spectral multiplexing, with an average efficiency of just 8.5 % and on-demand storage time of 3.5 µs with the full AFC protocol. By purely using temporal multiplexing, 100 temporal modes have been stored, coincidentally also with an efficiency of 8.5 % but a larger, fixed storage time of 51 µs [35]. The highest storage efficiency with a single mode is achieved with the cavity enhanced AFC protocol, resulting in an efficiency of 56 % after a fixed storage time of 1.1 µs [36]. Very recently a storage time in the order of a second (0.53 ms) has been achieved by using dynamical decoupling in the full AFC protocol, although the storage efficiency was limited to 0.5 % [37]. In the near future, the Tittel group at the Delft University of Technology plans on performing an experiment that aims to achieve 10 frequency modes and 50 % efficiency per mode [38].

2.2.2 Electromagnetically Induced Transparency

The second type of memory we will describe is based on electromagnetically induced transparency (EIT) [39–42], introduced in 1990 by Harris and colleagues [43]. This is an effect where an opaque sample of atomic ensembles becomes transparent by occupying a dark state that arises from the interaction with two distinct electromagnetic fields. This transparency is accompanied by a very steep dispersion, which in turn affects the group velocity of an incoming light pulse. Under certain conditions this effect can cause light to be slowed or even completely stopped and later re-emitted, hence its use as a quantum memory [27].

A three-level atomic configuration is required as well as two optical fields that operate at different resonance frequencies, see Figure 2.6a). Combined with an adiabatic change in light intensity, the atoms can be trapped in a meta-stable spin-state that has a slightly increased energy compared to the ground state. This time-reversible process can subsequently be used to coherently map the atomic spin excitation back to photons that exit the medium.

About twenty years ago the first realization of the EIT protocol led to a cover page article in Nature, in which the ground-breaking experiment caused light to travel at the speed of a cyclist [44]. This spurred further research into EIT, which has lead to its application in phenomena such as lasing without inversion [45], optical switches [46], frequency converters [43] and of course QM.

EIT shares some similarities with QMs based on stimulated Raman emission, which is used in the DLCZ protocol for example. Even though they both use the same three-level system and two optical fields to store a photon, in EIT these optical fields must be (close to) resonance, while the Raman scheme is based on an off-resonant transition. This implies that one will also need a higher-powered laser to control the storage for the latter, which generally leads to an increase in noise and unwanted side effects such as spontaneous four-wave mixing [27]. However, the Raman scheme has a broad spectral bandwidth and therefore allows for spectral multiplexing, in contrast to EIT based memories which need to operate close to a single resonant frequency, as we will describe below. The storage time of both memories is limited by the spin-decoherence time of the system, which depends on a lot of factors such as the type of atoms and the state of matter in which they are.

We will start with the general physical requirements needed for EIT, after which we describe the effects in two distinct pictures that highlight different aspects of EIT. Finally, we will briefly describe some recent experimental realizations as well as the multiplexing capabilities of EIT based memories.

Eigenstates

Consider a three-level system in a Λ configuration as illustrated in Figure 2.6a). Here $|g\rangle$ is a stable ground state, while $|s\rangle$ is a meta-stable spin-state with a slightly higher energy level, which can for example be generated by lifting a spin degeneracy by inducing a Zeeman splitting. The state $|e\rangle$ represents the excited state of the electron and can be used to transition from $|g\rangle$ to $|s\rangle$, since this is not possible directly due to a dipole blockade. The latter implies that the states have the same parity of the angular momentum quantum number such that the selection rules in quantum mechanics forbid a transition [47].

We continue by introducing the two electromagnetic fields that interact with this atomic three-level system. The so called pump (control) field has a frequency ω_p (ω_c), detuned by Δ_p (Δ_c) relative to the transition $|g\rangle \to |e\rangle$ ($|s\rangle \to |e\rangle$), with a corresponding Rabi frequency Ω_p (Ω_c). In QMss, the probe field is the single photon we would like to store, while the control field can be generated by using a laser.

If we apply a transform to a reference frame rotating with the frequencies of these optical fields, the corresponding combined interaction Hamiltonian is given by [40]

$$H = \hbar \left[\Delta_p \sigma_{ee} + (\Delta_p - \Delta_c)\,\sigma_{ss}\right] - \frac{\hbar}{2}\left[\Omega_p\,(\sigma_{eg} + \sigma_{ge}) + \Omega_c\,(\sigma_{es} + \sigma_{se})\right], \tag{2.14}$$

where $\sigma_{ij} \equiv |i\rangle\langle j|$. Note that we have assumed here that both Rabi frequencies are real. Solving for the dressed states, i.e. the eigenstates of this Hamiltonian that arises due to its interaction with the light fields, in terms of the atomic bare states generally cannot be done analytically. However, if we make the assumption that the probe and control lasers have equal detuning, such that $\Delta_p = \Delta_c \equiv \Delta$, the eigenstates are given by [40]

$$\left|B^+\right\rangle = \sin(\theta)\sin(\phi)\,|g\rangle + \cos(\theta)\sin(\phi)\,|s\rangle + \cos(\phi)\,|e\rangle \tag{2.15}$$

$$\left|B^-\right\rangle = \sin(\theta)\cos(\phi)\,|g\rangle + \cos(\theta)\cos(\phi)\,|s\rangle - \sin(\phi)\,|e\rangle \tag{2.16}$$

$$|D\rangle = \cos(\theta)\,|g\rangle - \sin(\theta)\,|s\rangle\,, \tag{2.17}$$

where

$$\tan(\theta) = \frac{\Omega_p}{\Omega_c} \tag{2.18}$$

$$\tan(2\phi) = \frac{\sqrt{\Omega_p^2 + \Omega_c^2}}{\Delta}. \tag{2.19}$$

The probability amplitudes of states (2.15) and (2.16) have a component in the excited state $|e\rangle$ and can therefore be detected by fluorescence measurements, hence they are referred to as bright states. Their eigenenergies are

$$\frac{\hbar}{2}\left(\Delta \pm \sqrt{\Delta + \Omega_p^2 + \Omega_c^2}\right). \tag{2.20}$$

On the other hand, the dark state (2.17) with eigenenergy 0, has no contribution in the excited state and can therefore not emit a photon under spontaneous emission. Furthermore, the transition probability of going to the excited state from the dark state under interaction with this Hamiltonian is zero, i.e. $\langle e|\,H\,|D\rangle = 0$.

Now consider the case where all atoms are initially in the dark state, which can conveniently be done by only applying the control laser such that $\Omega_c \gg \Omega_p$, which implies that all electrons populate the state $|D\rangle \approx |g\rangle$. Note that this can also be seen as a form of optical pumping. The adiabatic theorem states that the system will remain in this eigenstate as long as any pertubation occurs slowly, i.e. adiabatically, provided that there is a gap between the eigenenergies of the other states [48]. We can use this to carefully turn off the control beam while the probe pulse enters the atomic medium such that $\Omega_p \gg \Omega_c$ and therefore $|D\rangle \approx -|s\rangle$. This sequence is illustrated in Figure 2.6c) and *coherently* converts the light pulse into a metastable atomic spin-state, where any additional photon energy is absorbed into the control field. After a given amount of time the control laser can be turned on again, which coherently maps the spin excitation back into a photon without altering its state or shape, which then exits the medium along the same propagation direction as it entered. This procedure can thus be used to completely stop, store and and retrieve light, where the storage time is only limited by the spin decoherence rate γ_{sg}.

Susceptibility

After considering the interactions from the point of view of the atoms, we can also investigate the propagation of the electromagnetic probe field in more detail. In this case, the key of the underlying physics lies in the (linear) electric susceptibility $\chi(\omega_p)$. This can be interpreted as the response function of the polarization under an interaction with a (slowly varying) electromagnetic field [47]. Furthermore, in the context of EIT, its real and imaginary parts directly relate to the intensity transmission coefficient (ITC) $T(\omega_p)$ and the refractive index $n(\omega_p)$ by

$$T(\omega_p) = e^{-\operatorname{Im}[\chi(\omega_p)]L_s} \tag{2.21}$$

$$n(\omega_p) = \sqrt{1 + \operatorname{Re}[\chi(\omega_p)]}., \tag{2.22}$$

where L_s is the length of the sample. The ITC is defined as the fraction of the remaining intensity of the electric field that is left after propagation through an atomic ensemble of length L_s, which is also known as the Beer-Lambert law. In the context of the latter, $\operatorname{Im}[\chi(\omega_p)]$ can be interpreted as the linear attenuation coefficient. On the other hand, the refractive index describes the dispersive properties of the field within the atomic medium. This in turn also directly relates to the group velocity of the electromagnetic field [40]

$$v_g(\omega_p) = \frac{c}{n(\omega_p) + \omega_p \frac{\mathrm{d}n}{\mathrm{d}\omega_p}}, \tag{2.23}$$

where c is the speed of light in vacuum. This represents the velocity of the electromagnetic field envelope, in contrast to the phase velocity $v_p = c/n$ that describes the oscillatory velocity of a point on the wave. Now let us investigate the qualitative behaviour of all these quantities

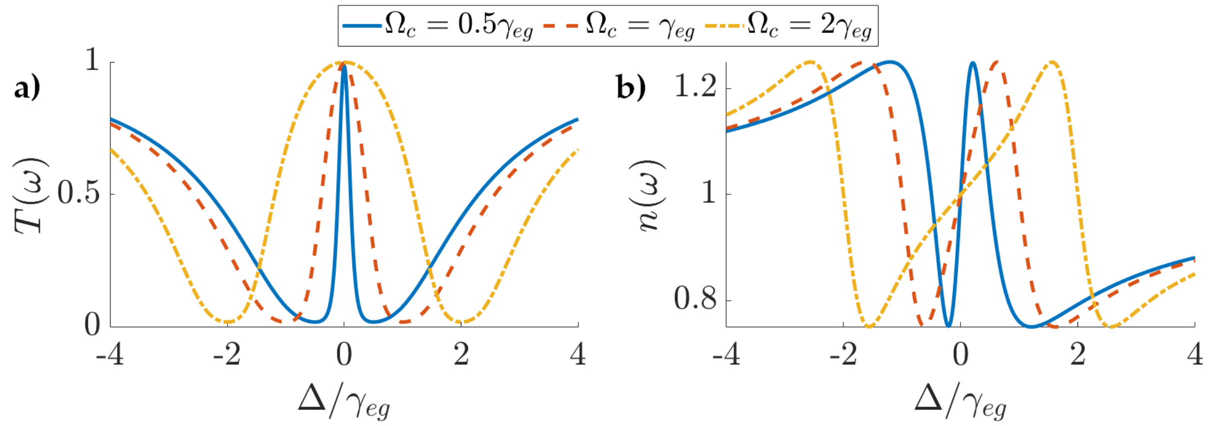

Figure 2.7: Absorptive and dispersive properties of an EIT medium as a function of the detuning Δ for three values of Ω_c in units of γ_{eg}, in an atomic ensemble with optical depth $a_0 L_s = 4$ and $\gamma_{sg}/\gamma_{eg} = 10^{-3}$. **a)** The intensity transmission coefficient $T(\omega)$. When the frequency is close to resonance, an incoming photon is transmitted rather than absorbed in a region called the transmission window. We can see that an increase of the the Rabi frequency of the control field causes the transmission window to broaden. **b)** Refractive index $n(\omega)$, where we see that as the control laser intensity decreases, the slope of the refractive index near resonance gets larger, which in turn causes a decrease in group velocity, see (2.23). When $\Omega_c \to 0$, this will become a discontinuity such that $v_g \to 0$ and hence the photon will be put to a halt.

in the context of the three-level system of Figure 2.6a). In this case, again under the assumption that the probe and control laser have equal detuning Δ, the linear electric susceptibility is given by [47]

$$\chi(\omega_p) = a_0 \frac{i\gamma_{eg}}{\gamma_{eg} - i\Delta + |\Omega_c|^2(\gamma_{sg} - i\Delta)^{-1}}, \tag{2.24}$$

where $a_0 \equiv \sigma_0 \varrho$ is the resonant absorption coefficient, defined by the product of the atomic absorption cross section σ_0 and the atomic density ϱ. The resulting ITC and refractive index are both illustrated in Figure 2.7.

Using this illustration we can continue with describing how a photon pulse is affected and possibly stored in an atomic medium with EIT. First and foremost we see that due to a non-zero control field we get a small range close to resonance, called the transmission window, in which the medium becomes transparent. Hence the name electromagnetically induced transparency. This is in contrast to a two-level system in which we get absorption on resonance. The transmission window has a spectral width of [47]

$$\delta\omega_{\text{tw}} \approx \frac{|\Omega_c|^2}{\gamma_{eg}\sqrt{2a_0 L_s}}. \tag{2.25}$$

In case the probe field is a photon pulse, it should thus hold that the bandwidth of the pulse should be smaller than $\delta\omega_{\text{tw}}$. Furthermore, the full temporal pulse width τ_p should be captured within the atomic medium, such that $\tau_p v_g \ll L_s$. This limits the amount of temporal multiplexing that is possible, which we will elaborate further on below.

In contrast to AFC, an inhomogeneously broadened sample would be disastrous for storing photons, since the transmission window would not be the same for all atoms and therefore the probe pulse would not be stored accurately. However, there will always be inhomogeneous broadening naturally present in the medium due to individual atoms moving at different velocities and therefore experiencing their own distinct Doppler shift in frequency. Luckily this effect can be compensated by applying the two optical fields colinearly, given that the energy levels of $|g\rangle$ and $|s\rangle$ are closely spaced, which is the case in most alkali atoms [47]. For example in [49] the two lasers are applied with a relative angle shift of just 1° to compensate for this Doppler shift. Additionally, one can cool the atomic ensemble to cryogenic temperatures to freeze the atoms, thereby reducing their velocities.

On the other hand, if we focus on the refractive index, we see that on resonance the refractive index is 1, such that the phase velocity remains unaffected. However, this is accompanied by a very steep slope, which even becomes increasingly steeper if we decrease the control field intensity. In fact, when we combine (2.22), (2.23) and (2.24), one can derive that $v_g \propto |\Omega_c|^2/a_0$ [47], which shows that we can manipulate the group velocity with our control field. In line with the eigenstate picture given above and the field intensities of Figure 2.6c), we start with $\Omega_c \gg \Omega_p$, which implies that we have a relatively large transparency window, but also a relatively high group velocity. Then, as the pulse is fully captured within the transparency window, the control beam intensity is adiabatically reduced to zero, which then also causes the envelope of our electromagnetic probe wave to come to a halt. As we have seen in the previous picture, the atomic state is then transferred from the ground state $|g\rangle$ to the metastable state $|s\rangle$, such that we have stopped our photon and coherently stored it in a long-lived atomic spin-state. At a given point later in time, we can turn the control beam on again, which transfers the atom back to its ground state, while simultaneously increasing the group velocity of the photon back to a positive value as well as opening up the transparency window, such that the photon propagates out of our atomic ensemble.

Multiplexing Capabilities and Experimental Status Quo

Multiplexing with EIT based memories is quite difficult. In [27] the multi-mode capacity is even characterized as ‘low’, in contrast to the ‘high’ multi-mode capacity of AFC based QMs. The main problem for EIT is the fact that the number of (temporal) storage modes depends on the optical depth, in contrast the multiplexing capabilities of AFC memories, which do not independent on the optical depth. Theoretical analysis shows that the number of temporal modes scales with the square root of the optical depth, compared to a linear scaling with optical depth for other memory types such as those based on controlled reversible inhomogeneous broadening or gradient echo [50]. This is also reflected in (2.25), where we can directly see that the temporal width of the transmission window, which is approximately the inverse of its spectral width, is inversely proportional to the square root of the optical depth $d = a_0 L_s$. Even though ultra high optical depths of more than 1000 have been experimentally realized [51], this would theoretically only allow for just over 30 temporal modes.

In case of spectral multiplexing all pulses have to be contained within the EIT transparency window, which implies that if we need to have access to a photon source that is capable of emitting extremely narrow pulses in frequency. However, these pulses will then be very broad in time, such that they might not be fully captured within the medium before the control beam is turned off. Additionally, the further away these pulses are from the resonance frequency, the more they are absorbed which can be seen in Figure 2.7. Increasing the spectral bandwidth of our transmission window would come at the cost of a lower initial group velocity, which could cause pulses to escape the medium before their attempted storage.

One alternative is to use spatial multiplexing, where either the light pulse is split over a number of transverse modes, i.e. perpendicular to the direction of propagation, or an array of parallel atomic ensembles is used to store photons simultaneously. In the first case, up to 2 spatial modes have been realized [49] with an efficiency of 68 %, while in the second case experimental violations of a Bell inequality with 12 parallel spatial modes have been reported [52]. However, we should note that an array of parallel atomic ensembles is of little practical use in large scale quantum repeater setups compared to the other types of multiplexing.

Both frequency and angular multiplexing with up to two modes have been achieved by using spectral hole burning in REI doped crystals [53]. With angular multiplexing, storage and retrieval of the probe field is done by two separate control fields that operate under a different angle. This also causes the atomic spin waves to be stored with different wave vectors, which can selectively be addressed by phase-matching conditions for retrieval. However, in both cases the efficiency of retrieving the classical probe beam was in the order of 1 %, which makes it a proof of concept rather than a currently realistic alternative. Nonetheless, in the theoretical analysis of [54] it is noted that for a set of highly optimistic parameters, it should be possible to achieve in the order of 100 angular modes with EIT based QMs.

The non-classical storage of photons emitted from an SPDC source in an EIT memory based on rubidium atoms has also been experimentally demonstrated [55], although the storage time and efficiency were limited to 0.4 μs and 14 % respectively. This is mainly caused by the fact that SPDC sources naturally have a large bandwidth (THz), while the required bandwidth for EIT is much smaller (MHz), such that lossy filters or optical cavities are required to reduce the SPDC output bandwidth. The best experimental realization of single qubit storage in an EIT based memory has an efficiency of 86 % for a storage time of 1 μs [56]. At the moment of writing, the same group as in [49] is working on a very promising experiment with which they aim to achieve a QM capable of storing 15 spatial modes combined with a a 90 % efficiency per mode [57].

3 Simulation Model

With a broad range of physical systems, protocols, photon encodings and a lot of free parameters we can tune, it quickly becomes intractable to derive closed form analytical expressions for metrics like the secret key rate. Some extensive work has been done in [17, 24, 33] to derive analytical expressions for the entanglement generation probability and secret key rate for the spectrally multiplexed repeater chain protocol discussed in Section 2.1.1. However, by focusing on analytical expressions, we are ultimately limited to idealizations of how the setup will perform in reality. Moreover, these expressions are quite rigid and the inclusion of for example another type of PPS or time dependent noise could require a full revision of the model, if it is at all still possible.

The NLBlueprint team is therefore currently working on a full-scale simulation model for repeater chains not only based on atomic ensembles, but also for trapped ions and nitrogen-vacancy centers in diamond [58]. For the latter, an overview of the simulation model can be found in [59], from which we also used some designs such as the BSM stations. In the following we will mainly focus on the additional modeling we have done in order to realize the simulation model for atomic ensemble based repeater chains.

The basis of our simulation, such as the low-level modeling of qubits, operators and noise models, is handled by NetSquid [60]. This is a discrete event simulator for quantum networks, written in Python and C++ and based on DynAA [61]. It is actively being developed by QuTech, a collaboration between the Delft University of Technology and the Netherlands Organisation for Applied Scientific Research (TNO). The goal of our work on the simulation model was to come up with a set of modular and tunable NetSquid based components, which can be combined in a pick-and-choose manner in order to simulate and numerically evaluate the performance of quantum repeater chains based on atomic ensembles. We would like to note that a full and more detailed overview of the simulation model, including the local protocols per node and more elaborate results such as parameter optimization, can be found in [62]. A snippet of the full code will also be made publicly available.[1] In the remainder of this section, we will limit the scope to a brief overview of its capabilities and the modeling of a few components on which we mainly contributed.

3.1 Components

The power of our simulation model lies in the fact that the modeling of abstract components for the channels, memories, photon sources and nodes allows for a pick-and-choose manner of creating a repeater chain. One of the main parameters is the number of repeaters involved in the chain, which can be set to any non-negative integer without significantly affecting the

---

[1] For current access to the repository, please contact d.j.maier@tudelft.nl.

required computer memory for running a simulation. In this subsection we will go over all the other building blocks of a quantum repeater chain one-by-one.

Channels

Both the quantum and classical channels that we respectively use to send the entangled-photon pairs and receive the BSM outcomes are standard components in NetSquid. These can have delay, loss and noise models attached to them, which affect the messages that are sent through. We can for example have loss of photons due to fiber coupling and subsequent length-dependent transmission loss in a quantum channel. One can also include depolarizing or dephasing noise as well as fixed or Gaussian delay models to analyze how these affect the performance of a quantum repeater chain. In general one can assume that the classical channels are noiseless due to their reliability in the real world, but nonetheless a user is able to for example add bit flip noise to the classical channel. Additionally, a user is able to change the length of these channels for each elementary link. This conveniently allows for the relaxation of the frequently used assumption that the remote BSM station is always positioned exactly in the middle of an elementary link [16].

Qubits and Photons

Our simulation also allows for the use of multiple degrees of freedom of photons. Next to the time-bin encoded photons required for the protocol of Tittel et al., we have also included the presence-absence encoding such that users can analyze the performance of protocols that rely on the single-click scheme at the BSM stations [62].

Spectral multiplexing is done by combining multiple qubits in a list after which we measure them sequentially. When the first successful mode is found, we discard the rest of the photons. Since the discrete event simulator allows this to occur at the same instance of simulation time, it realistically emulates the parallel measurement of multiplexed photons.

One of the largest sources of noise that deteriorates the quality of a repeater chain, arises from a realistic photon source in the form of multiple-photon pair emissions. However, the modeling of multi-pair emissions would require the use of quantum states with more than two dimensions, which are not available in NetSquid. We therefore choose to use a binary-to-decimal encoding for the modeling of these higher dimensional photon-number states. In particular, we use two qubits to represent the number of photons at any point in time. This implies that we can have four possible photon numbers, where the vacuum state is represented by $|00\rangle$ and the largest photon state with three photons is $|11\rangle$. This allows a user to analyze the effect of multi-pair emissions upto a certain extent, but at the cost of creating significantly larger quantum states.

Quantum Memory

Similar to the channels, QMs are a standard abstract component in NetSquid. In principle they store a list of memory positions, of which empty ones can store qubits. A user can include both noise and loss models, which can for example be T_1 or T_2 noise or (time dependent) retrieval loss η_{mem}. This enables the easy modeling of any type of quantum memory such as those based on AFC or EIT. We assume that there is no intrinsic change in the way the QM stores multi-photon states compared to single photons. This is motivated by the fact that experimentally the same setup can be used to store both (weak) coherent states and single photons, see for example [55]. However, we did have to make adjustments to the loss and noise models such that they would work with our binary-to-decimal encoded photon states.

Photon Source

An important component of our simulation is the entangled PPS. A wide range of possible physical systems are available to generate the photons, which all have their own distinct photon emission probability distribution. In order to include a general PPS in our simulation model, a user is able to give a list of probabilities $P_{\text{em}}(n)$ of emitting a photon pair with n photons, truncated to at most three photons. A perfect PPS would deterministically emit a single-photon pair each round, but we could include a simple form of noise by making this probabilistic. In that case, the source would emit either nothing or a single-photon pair with probability $P_{\text{em}}(1)$, which can be set by the user. The probability distribution would then be $P_{\text{em}}(0) = 1 - P_{\text{em}}(1)$ and $P_{\text{em}}(2) = P_{\text{em}}(3) = 0$. It is important to note that we can then also resort to back to qubit states with a dual-rail encoding for the time bins [63]. We have additionally included the SPDC source probability distribution (2.9), such that a user can also give a mean photon number μ as input. The time-bin encoded entangled-photon pair output state is given by [64]

$$|\psi\rangle_{\text{PPS}} = \sum_{n=0}^{3} \sqrt{P_{\text{em}}(n)}\,|\psi_n\rangle\,, \tag{3.1}$$

where

$$|\psi_n\rangle = \frac{1}{\sqrt{n+1}} \sum_{m=0}^{n} (-1)^m\,|n-m, m; n-m, m\rangle\,. \tag{3.2}$$

Here we use the notation that the state $|i, j; k, l\rangle$ refers to a state with i (j) photons traveling to the left in the early (late) time bin and k (l) photons traveling to the right in the early (late) time bin. Combined with the binary-to-decimal encoding of our photons as described above, the highly entangled output state of our PPS is thus

$$
\begin{aligned}
|\psi\rangle_{\mathrm{PPS}} = {} & \sqrt{P_{\mathrm{em}}(0)}\, |00,00;00,00\rangle + \sqrt{\frac{P_{\mathrm{em}}(1)}{2}} \left(|01,00;01,00\rangle - |00,01;00,01\rangle\right) \\
& + \sqrt{\frac{P_{\mathrm{em}}(2)}{3}} \left(|10,00;10,00\rangle - |01,01;01,01\rangle + |00,10;00,10\rangle\right) \\
& + \frac{\sqrt{P_{\mathrm{em}}(3)}}{2} \left(|11,00;11,00\rangle - |10,01;10,01\rangle + |01,10;01,10\rangle - |00,11;00,11\rangle\right), \quad (3.3)
\end{aligned}
$$

where $P_{\mathrm{em}}(3) = 1 - P_{\mathrm{em}}(2) - P_{\mathrm{em}}(1) - P_{\mathrm{em}}(0)$. We thus need eight qubits in total to describe this state.

Magic

Such large eight-qubit states also affect the way we should model our noise and loss models. For example, when the state $|10\rangle$ loses a photon, it should decay to $|01\rangle$, which differs from applying a loss model to the two qubits separately. These more elaborate noise models combined with the large qubit states strongly affect the required computation time of our simulation. To alleviate this, we also implemented a sampling method for our elementary links, which we call magic. Note that this is completely unrelated to magic states in the context of error correction [65].

We distinguish between two types of magic: analytical and sampled. For analytical magic, we simply implement a full analytical model for P_{el} and the corresponding states that are generated for a given set of simulation parameters, similar to what is done in [24]. We can then 'magically' distribute a quantum state over the memories of an elementary link after a certain number of rounds, rather than simulating the full emission, loss, noise and measurement of the entangled-photon pairs. We can use this to quickly generate data for a set of parameters, but this also somewhat defeats the purpose of our elaborate simulation model. Moreover, deriving analytical expressions quickly becomes intractable once we for example introduce multi-pair emissions from our source [33].

For sampled magic, we first run our full simulation to generate a (large) data set that contains the states, the corresponding BSM outcome and the number of rounds for a single elementary link. We can then use this data set to sample the states for any other elementary link in the same or an independent repeater chain. Note that this only works for exactly the same simulation parameters as those that were used to create the data set. Especially for parameters such as detector or memory efficiencies that can take any real value between 0 and 1, a small change instantly renders the data set unusable. Furthermore, we will introduce additional statistical errors in our sampling due to the finite size of the data set. Nonetheless, this strategy can significantly improve the computation time of our simulation, especially if we use a large number of repeaters and modes.

BSM Stations and End Node Detectors

The photon detectors are another crucial part of our simulation as they perform the measurements required for the entanglement generation, entanglement swaps and QKD.

The BSM stations consist of a single beam splitter and two photon detectors. For the latter we have modeled both single-photon (threshold) detectors as well as photon-number resolving detectors. We use the same detector stations for the BSMs at the midpoint of an elementary link as described in [59] for single photon (qubit) states. However, we have extended these to also work for our binary-encoded higher photon-number states. The relevant parameters for all the detectors in our setup are the efficiency η_{det} and dark count probability P_{dark}. These respectively correspond to the probability of losing a photon and the probability of wrongfully detecting one.

One additional component we have included in our model are the end node detectors that are used for QKD, which can perform measurements in three orthogonal bases and hence are able to interfere photons at the end nodes. With these three measurement bases a user is in principle able to implement the six-state protocol, an extension of the BB84 protocol with an additional basis, which is more resilient to adversaries. In the case of presence-absence encoding, we need two parallel repeater chains to perform QKD and we can use the same components that are used for the BSMs in order to perform measurements in two bases [62]. However, in the case of time-bin encoded qubits, we will need a separate type of detector at the end nodes.

Measurements in the Z basis are rather straight forward in this case, since we can measure the photons in both time bins separately at the end nodes. However, if we want to perform measurements in the X and Y bases, we will need to be able to interfere the two time bins with an unbalanced interferometer. The length difference of the arms of the latter should be exactly equal to the one that was used to create the initial time-bin encoded photons. The interference can be achieved in two ways: either *actively* with an optical switch, or *passively* with an additional beam splitter. Both setups use threshold detectors and are shown in Figure 3.1.

With the fast optical switch, the early (late) mode is always directed along the long (short) path. The photons will then always arrive at the beam splitter at the same time, such that we have a guarantee that they will interfere. This would result in measurements in the X basis, while the phase modulator can add an additional phase of $\pi/2$ to one of the arms in order to perform measurements in the Y basis. Note that after the switch and phase modulator, the photons are measured with a BSM station (one beam splitter and two detectors). Although ultra-fast optical switches have been experimentally demonstrated with a switch time in the order of ps, the efficiency is rather limited at around $1\,\%$ [66].

Alternatively, if we do not have access to an efficient fast optical switch, we could resort to using a second beam splitter to create the interference. In this case, there are six possible bins in which a photon could end up, so we would need two detectors that must be capable of performing detections in three distinct time windows instead of one. In principle we would

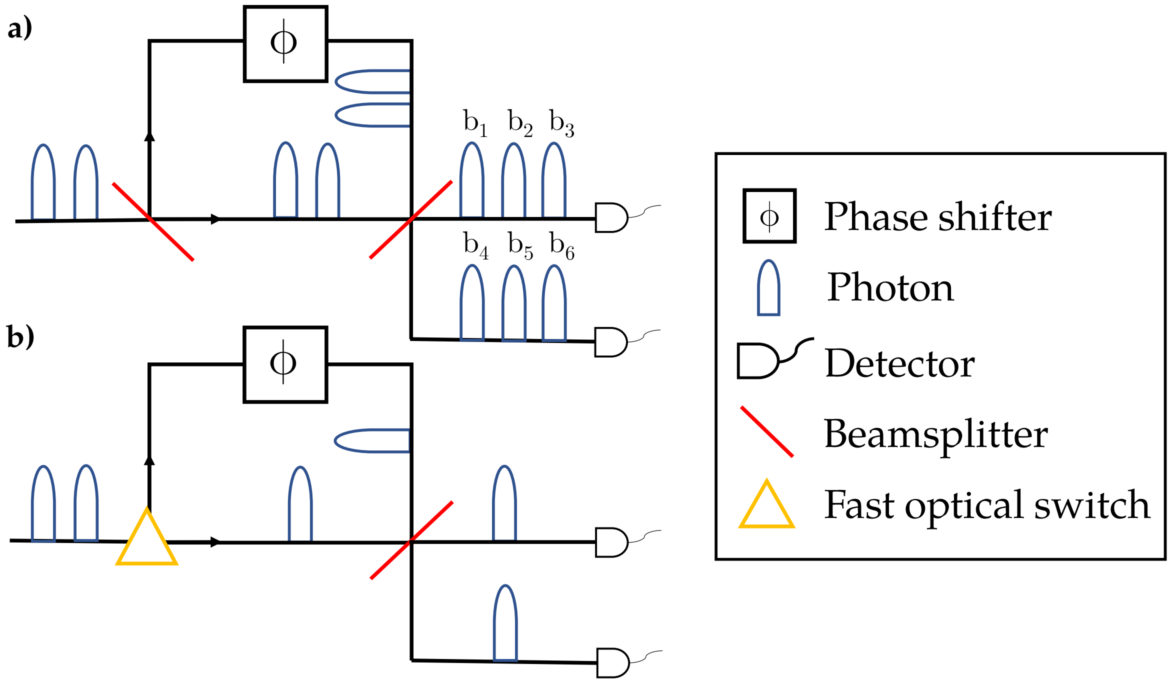

Figure 3.1: Illustration of our end node detectors, where the two time bins are interfered in order to perform measurements in the X basis. A phase shifter can be used to apply a phase shift of $\pi/2$ which allows for measurements in the Y basis. **a)** Passive setup with two beam splitter, in which our photon ends up in one of six detection bins $b_1, \ldots, b_6$, while we are only interested in clicks in bins b_2 or b_5 for our QKD measurements, since interference has occurred only in these bins. **b)** Active setup in which the early (late) photon is always routed to the long (short) arm with a fast optical switch, such that they will always correctly interfere at the beam splitter.

then only be interested in single clicks in of the middle time bins, which would imply that we have to throw away a large portion of our data due to photons unfortunately ending up in the incorrect measurement bin. Luckily we can still use cross-clicks by implementing the squashing model described in [67]. This gives a complete and secure classical post processing scheme for the passive setup. By using this scheme we can for example uniformly at random assign cross-clicks in the middle and late time windows to a usable measurement outcome. In the end, a user can easily switch between these two types of detectors to have a trade off between a higher secret key rate or a setup that is experimentally more realistic, which is another example of the power and modularity of our simulation.

3.2 Results

With the translation from physical to NetSquid-based components in place, let us continue with our results. We will start with the validation of our simulation model by comparing it to the analytical model of Guha et al. [24]. Once our model is validated, we can subsequently analyze the effects of parameters that are not taken into account for this analytical model, such as the visibility and the use of an SPDC source with multi-photon errors. Finally, we will perform a comparison of the quantum memories described in Section 2.2, in order to characterize their performance in a quantum repeater setup.

3.2.1 Validation

We have verified the correctness of each component individually with a number of unit tests, aimed to check whether a given input results in the expected output. These for example include the verification of whether our PPS emits the expected state and whether our end node detectors are able to correctly perform measurements in three bases. We have also included some integration tests which check whether the communication between components operates as expected. Once the correctness in these cases is validated, we can continue with generating results of fully fledged quantum repeater chains by building and connecting the individual components and nodes as well as include all the local protocols per node, which are described in [62].

Let us start by validating that our model works correctly for different number of repeaters. Since we are working with a modular and versatile simulation model where tasks have to be executed sequentially and handled correctly, it is vital to confirm that this is done in the right way. The validation in this case is based on the output of our model in terms of secret key rate, QBER and the number of required attempts (rounds) per success for a varying end-to-end distance L. We run all of our simulations up until a fixed number of successes have been achieved, where a success is defined as at least one click on both end node detectors. We choose to use the same measurement basis for both end nodes in order to generate data faster, such that we have to scale our final secret key rate by a factor of $1/2$. We first measure our photons in the Z basis until half the successes are generated, after which we switch to the X basis.

The results are shown in Figure 3.2 for a setup with a single elementary link, one repeater and two repeaters. For now we use a perfect PPS, i.e. $P_{\mathrm{em}}(1) = 1$ and $P_{\mathrm{em}}(0) = P_{\mathrm{em}}(2) = P_{\mathrm{em}}(3) = 0$ and the other parameters that we use are shown in the caption of Figure 3.2. Note that these parameters are hugely optimistic and can probably not be realized in the near future, but the goal here is to only validate our simulation model.

We see that, aside from a few statistical outliers, our simulated secret key rate accurately follows the analytical model of Guha et al. For every N, the secret key rate remains at a constant value as the end-to-end distance increases, which is caused by $P_{\mathrm{el}}^{\mathrm{mp}}$ being very close to 1 due

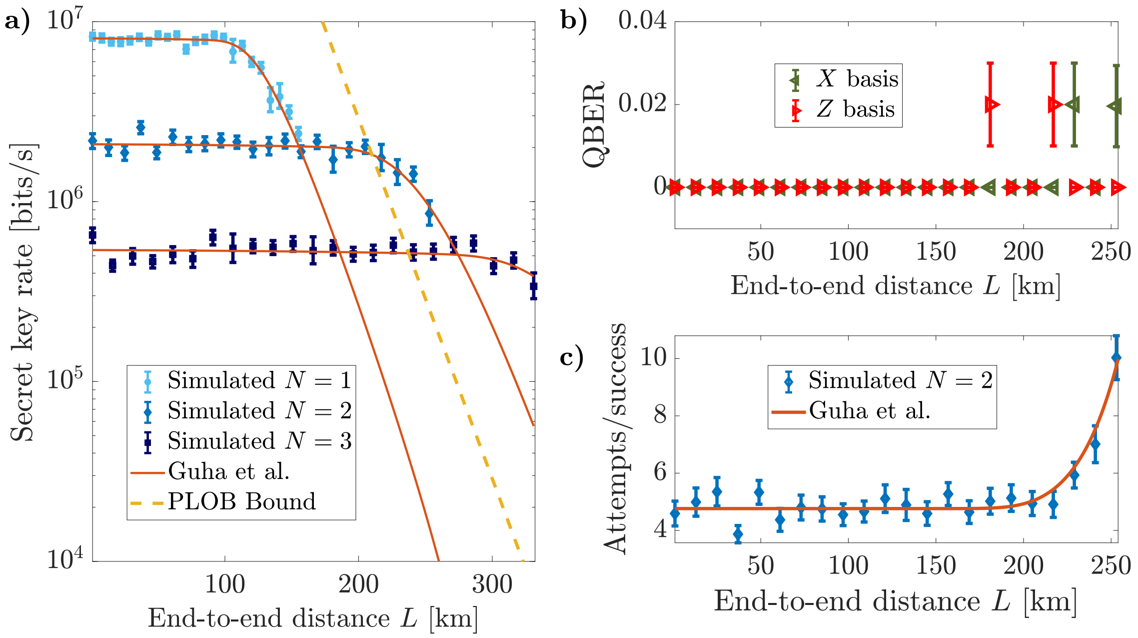

Figure 3.2: Validation of our simulation model for a single elementary link ($N = 1$), one repeater ($N = 2$) and two repeaters ($N = 3$), compared to the analytical model of Guha et al. [24]. Each data point corresponds to 100 successes and the error bars represent either upper and lower bounds or one standard deviation, see Appendix A. **a)** Secret key rate. The solid red line represents the analytical model and the blue points are the simulated data. We can see that for each value of N there is a regime where the large number of modes causes $P_{\text{el}}^{\text{mp}}$ to be (close to) unity, such that the secret key rates remains constant over distance. For $N = 1$ and $N = 2$, the quantum repeater chain beats the upper bound on the secret key rate that is achievable with direct transmission. **b)** The QBER for $N = 2$. It remains zero for in the flat secret key rate range since we are dealing with a perfect PPS and a low dark count probability. When the exponential decrease of the secret key rate kicks in, dark counts start to give a small amount of false positives, hence the increase in QBER. For $N = 1$ and $N = 3$, the QBER is qualitatively equal. **c)** Number of attempts per success for $N = 2$. For distances beyond about 200 km the number of modes is not sufficient to guarantee a unit success probability anymore and the number of required attempts to successfully generate entanglement increases exponentially There is similar agreement for $N = 1$ and $N = 3$. All data is generated with an attenuation coefficient $\alpha = 0.2$ dB/km, PPS frequency $f_{\text{PPS}} = 20$ MHz, the ideal PPS distribution $P_{\text{em}}(1) = 1$ and $P_{\text{em}}(0) = P_{\text{em}}(2) = P_{\text{em}}(3) = 0$, $M = 10^3$ frequency modes, dark count probability $P_{\text{dark}} = 3 \cdot 10^{-5}$, detector efficiency $\eta_{\text{det}} = 0.9$, visibility $V = 1$ and fixed memory efficiency $\eta_{\text{mem}} = 0.8$.

to the large number of modes. We then only have to rely on the local BSM to succeed and on the two end nodes to measure in the same basis and get at least one click, which occurs with probability of $(P_{\text{swap}})^{N-1} P_{\text{meas}} \approx 0.2$ for $N = 2$. This is indeed accurately reflected in Figure 3.2c) with an average number of attempts per success of approximately 5. After a certain distance $P_{\text{el}}^{\text{mp}}$ gets significantly smaller than 1 due to the transmittance η_t, such that the number of attempts increases exponentially, which causes a sharp drop in the secret key rate. This decline occurs at varying lengths for different number of repeaters, due to the difference in the distance L/N that the elementary links span. The QBER remains practically zero for all distances, since we use a PPS that deterministically emits a single-photon pair each time it is triggered. The only cause of non-zero QBER is detector dark counts in the range where the number of modes is not sufficient anymore to guarantee that $P_{\mathrm{el}}^{\mathrm{mp}}$ is close to unity. Last but not least, if we use one or two repeaters, we are able to surpass the PLOB bound for the secret key rate such that our repeater actually repeats.

The generation of a data point, which corresponds to running the simulation up to 100 successes, in the 'flat' secret key rate region takes a about 3 minutes for $N = 1$ on our laptop with a 2.8 GHz processor and 8 GB of RAM. However, this increases to more than 30 minutes for a setup with $N = 3$. Additionally, once the number of attempts starts to increase exponentially, so does our computation time. We therefore only simulated our data in a limited region for all values of N. However, our simulation is quite efficient in terms of memory usage, as about 85 MB of RAM is required to run a simulation for $N = 3$.

3.2.2 Photon Indistinguishability

After the validation of our simulation model in place, we can continue to analyze effects that cannot be taken into account in the model of Guha et al. Let us start with the photon visibility. We use the same value for V in both our BSM stations and our end node detectors.

The result for $V = 0.9$ and $N = 2$ is shown in Figure 3.3. We can clearly see the negatives effects of a non-unit visibility, since there is a drop in the secret key rate as well as in the statistical uncertainty of our data points. It is also not certain anymore whether we are able to surpass the PLOB bound. This is all caused by the big increase in QBER in the X basis, as we would expect. The distinghuisability of the photons causes a loss of phase coherence and causes a mismatch in bits that are generated from measurements in the X basis. If we would choose to perform our measurements in the X and Y basis for the BB84 protocol, the secret key rate will drop to 0 as the QBERs in both bases would exceed the threshold of $R(Q)$. The number of attempts remains qualitatively unchanged with respect to $V = 1$, which is used in the model of Guha et al., and the computation time for generating 100 successes also remains practically the same.

3.2.3 Realistic Photon Pair Sources

Thus far we have used a perfect PPS in our setup which deterministically emits a time-bin encoded single-photon pair in each round. In this subsection we will deviate from this ideal case and instead use a PPS which emits photons according to an SPDC probability distribution (2.9). We choose to use a mean photon number μ of 0.05 in order to keep the probability of emitting multiple-photon pairs as low as possible. However, this also strongly reduces the probability of emitting an single-photon pair $P_{\mathrm{em}}(1)$ to approximately 0.09.

We would like to note that in [33] the authors also analyze the effects of an SPDC source

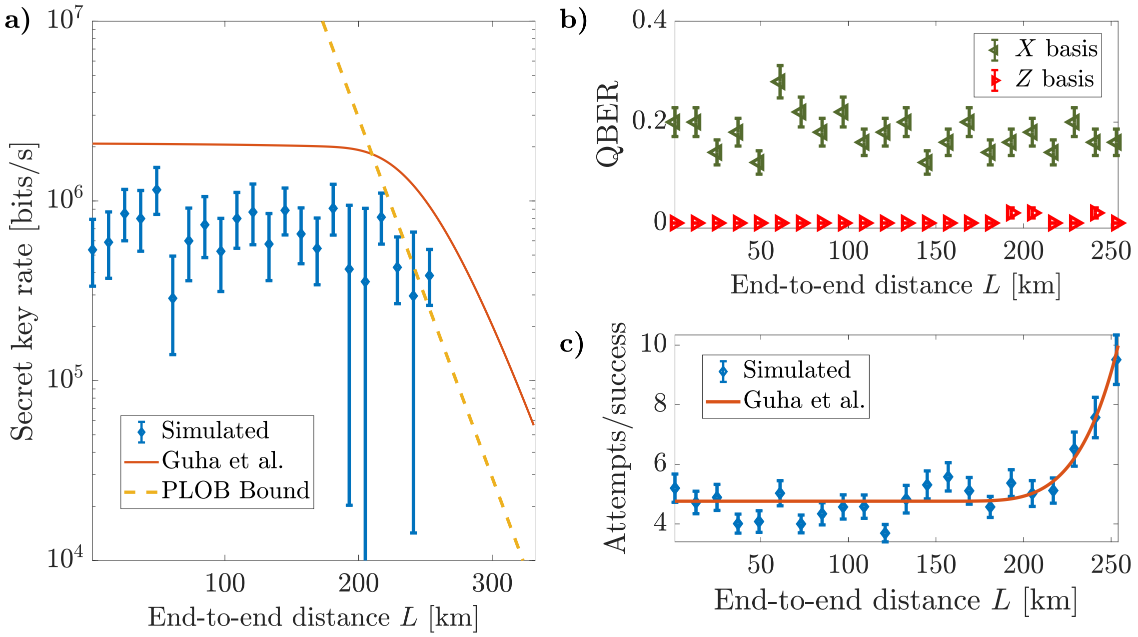

Figure 3.3: Single repeater setup with visibility $V = 0.9$. Each data point is generated for 100 successes and the errorbars represent one standard deviation or upper and lower bounds, see Appendix A. **a)-b)** Secret key rate and QBER. We can see that there is a drop in the secret key rate relative to $V = 1$, which is used in the model of Guha et al., and an increase of the size of the errorbars due to a higher QBER in the X basis. For one data point at 205 km, the lower bound on the secret key rate is 0, since the upper bound on the QBER exceeds the threshold as discussed in Section 2.1.2, i.e. $R(Q) = 0$ in (2.10). We cannot surpass the direct transmission rate for distances smaller than 250 km within one the range of the errorbars. **c)** The number of attempts is qualitatively the same as with $V = 1$, such that the drop in secret key rate can solely be attributed to the increase in QBER. The other parameters that we used are the same as in Figure 3.2.

on the protocol of Tittel et al., by using an analytical model. However, there are some key differences between their setup and ours that do not makes us able to compare the results. First of all, the state that their PPS emits takes up to two-photon pairs into account, while in our simulation model we include states with up to three-photon pairs. On top of that they use an incorrect distribution for their SPDC source. If we would try to compare our results to this nonetheless, we would have to simulate a million modes, which is the value of M that the authors use. In our simulation model this would require approximately 33 GB of RAM per PPS, which makes it infeasible to run (on our laptop). Additionally they use photon-number resolving detectors at the end nodes in order to classically post select on the rounds in which single photon states were emitted, while we use threshold detectors at the end nodes.

Since we are now using the larger eight-qubit states, we decrease the number of successes to simulate to 50, which still takes about 3 to 8 hours to complete, depending on the value of L. Moreover, generating a single success at a distance of 150 km takes more than 1 hour in this setup, such that we have generated data for a smaller length range compared to the previous

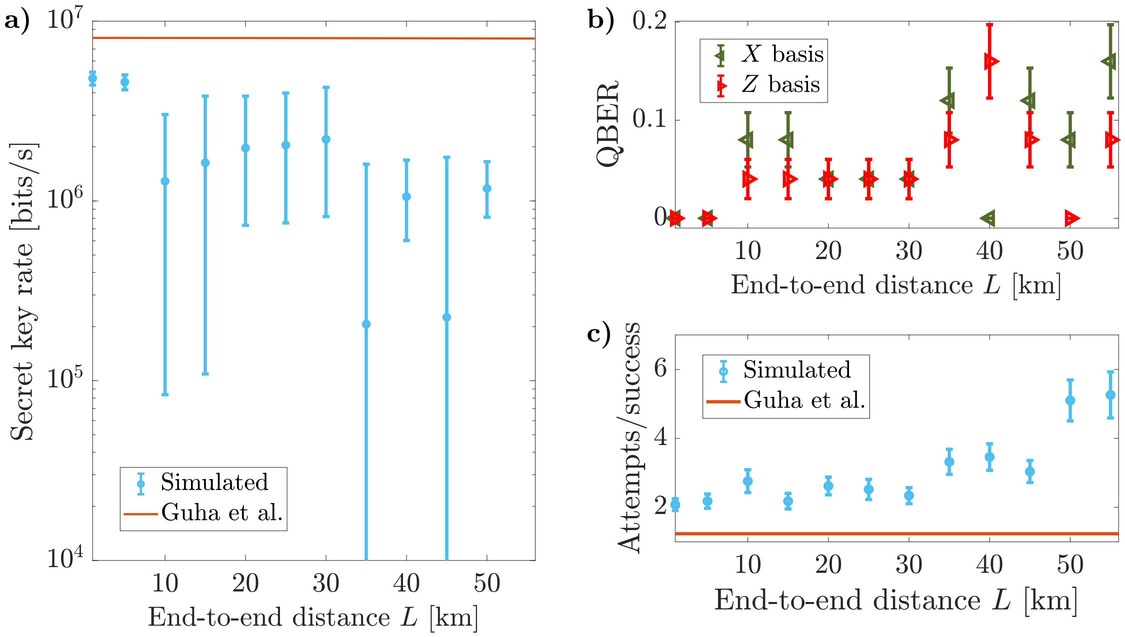

Figure 3.4: Single elementary link setup in combination with two SPDC sources that operate with $\mu = 0.05$, which corresponds to $P_{\text{em}}(0) = 0.907029$, $P_{\text{em}}(1) = 0.086383$, $P_{\text{em}}(2) = 0.006170$ and $P_{\text{em}}(3) = 0.000416$. Each data point is generated for 50 successes and the errorbars represent one standard deviation or upper and lower bounds, see Appendix A. **a)-b)** Secret key rate and QBER. We can see that there is a significant drop in secret key rate compared to the perfect PPS that is used in the model of Guha et al. The errobars are also bigger, which can be attributed to the increase in QBER. We think this is caused by the low probability of emitting a single-photon pair, such that the probability of wrongfully detecting one gets similar. For $L = 55$ km, the mean of the mean secret key rate is 0 due to the high QBER, such that the data point cannot be shown on a logarithmic scale. **c)** The number of attempts per success seems to increase exponentially at around 40 km already, instead of the 100 km that can be seen in Figure 3.2 for $N = 1$. This implies that $M = 1000$ is not sufficient anymore for $P_{\text{el}}^{\text{mp}}$ to be close to 1. The other simulation parameters are the same as in Figure 3.2.

two cases. We use photon-number resolving detectors at the BSMs in order to partially filter out multi-photon pair emissions, provided that the photon is not lost.

The results are shown in Figure 3.4. We can see that the secret key rate has significantly dropped compared to the perfect photon pair source of Guha et al. This is also reflected by the exponential increase in the number of attempts per success at just 40 km, which implies that $P_{\text{el}}^{\text{mp}}$ is not close to 1 in this range. The QBER is also quite high in both bases, which we think is caused by either vacuum or multi-photon states which are incorrectly detected as single photons. Since we use threshold detectors for the end node measurements, we cannot post-select on the rounds in which we have single-photon pairs. It would be interesting to analyze what the effect of photon-number resolving detectors would have in this case. All in all, further investigation is required in order to determine the exact cause.

A very promising way of resolving the huge increase in computation time has been found with the use of sampled magic [62]. If we generate a large data set of single-mode simulation data for an elementary link, we can use this to sample the states for the elementary links in a separate repeater chain simulation in which $M \gg 1$, by using the fact that all modes are independent. Note that this assumes that there is no cross-talk between the modes. We would then only need to draw a single random number to determine in which round the entanglement generation has succeeded, after which we sample one of the single-mode states and place it in the QMs of the elementary link. Preliminary results show that this could speed up our simulation by a factor 1000, although creating more statistical errors due to the use of sampled data.

There are a lot of other interesting results one could generate with our simulation model due to the large availability of tunable parameters and components. These can (partly) be found in the work of [62] and in one or multiple papers that will be published based on this simulation model.

3.2.4 Memory Comparison

Due to the fact that we have the analytical model of Guha et al. to our disposal for the validation, we can directly use it to generate data for the memory comparison. We thus do not include multi-photon errors or visibility and therefore we also do not have to use our NetSquid based simulation model, which makes gathering the data much faster and without any statistical errors.

We have condensed the experiments mentioned in Section 2.2 with the highest efficiency (HE), largest multi-mode capacity (MM) and projected to be performed in the near future (PR) in Table 3.1. Since the lowest storage time τ in these experiments is 1 µs , we can analyze their performance in a setup where the elementary links have a length of at most $L_0 = \tau \bar{c} = 0.2$ km.

We have plotted the secret key rate for a setup with one and three repeaters as a function of M and η_{mem} in Figure 3.5. First when we look at the color scales, we can see that for three repeaters (four elementary links) the secret key rate significantly drops with respect to the same setup with one repeater, which is in line with our expectations (see Figure 3.2). Due to the fact

Table 3.1: Overview of the experiments that are mentioned in Section 2.2 for both AFC and EIT based memories in terms of retrieval efficiency η_{mem}, number of possible storage modes M and storage time τ. We distinguish between the experiment with the highest efficiency (HE), largest multi-mode capacity (MM) and a promising experiment that is projected to be performed in the near future (PR). A dash indicates that the value of the storage time is not currently known.

	AFC_{HE} [36]	AFC_{MM} [34]	AFC_{PR} [38]	EIT_{HE} [56]	EIT_{MM} [49]	EIT_{PR} [57]
η_{mem}	0.56	0.085	0.5	0.86	0.68	0.9
M	1	130	10	1	2	15
τ (µs)	1.1	3.5	-	1	1.2	-

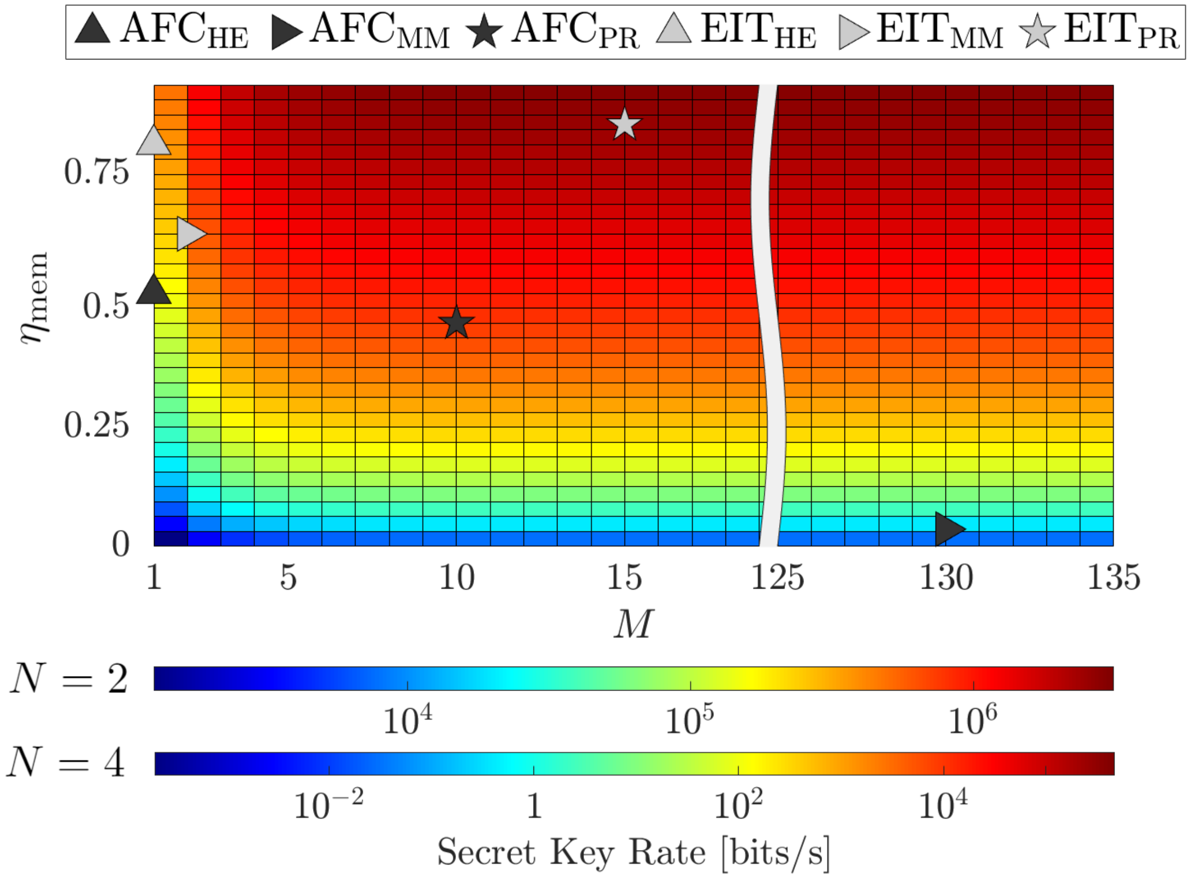

Figure 3.5: Secret key rate in bits per second on a logarithmic scale as a function of memory efficiency η_{mem} and number of multiplex modes M for one repeater ($N = 2$) and three repeaters ($N = 4$) and a fixed elementary link length $L_0 = 0.2$ km. As the number of repeaters increases, the maximum achievable secret key rate drops, but the qualitative behaviour remains unchanged. The grey symbols correspond to the AFC and EIT memories with the currently reported highest efficiency (HE), largest number of multiplexing modes (MM) and the projected (PR) outcome of two experiments that are planned to be performed in the coming years, as mentioned in Sections 2.2.1 and 2.2.2. We can see that the current EIT experiments outperform the AFC experiments due to the higher memory efficiency. Additionally, the projected EIT experiment of [57] has by far the best overall performance. The other parameters that we have used are the same as in Figure 3.2.

that the elementary links only span a distance of 200 m, the number of modes does not affect the secret key rate very significantly. It is sufficient to have around 5 spectral modes in order to be in the region where $P_{\text{el}}^{\text{mp}} \approx 1$, given the other parameter values.

Regarding the type of QM, we can see that the EIT memories that are currently experimentally realized would clearly outperform the AFC memories in terms of secret key rate. This shows that even though AFCs have a higher multi-mode capacity, the current very low retrieval efficiency limits them from having much practical use. When we look ahead, the projected AFC memory experiment will be able to outperform the currently realized EIT memories, but this will only hold up until the projected EIT experiment is performed, which will have the best overall performance.

When the memory lifetimes of both types of memories would increase, we would be able to analyze the secret key rates for greater lengths and the mode dependency would therefore also increase. However, even if we compare the highest reported memory lifetimes EIT still has the upper hand with storage times in the order of 1 minute realized more than 6 years ago [68], compared to the recently reported lifetime in the order of a second for AFC [37]. Note that both of these experiments were preformed with a single mode and a very low storage efficiency of less than $1\ \%$, such that we can reasonable compare them only based on storage time.

Additionally, the fact that EIT memories are inherently on-demand makes them more resilient against BSM outcomes arriving earlier or later than expected compared to AFC memories with a fixed storage time. When the BSM message arrives earlier, we can simply also apply our control beam earlier in order to retrieve the photons with EIT memories. However, with AFCs we would need to wait until the fixed storage time has passed. On the other hand, when the classical message from the BSM station arrives later than expected, the photons will be completely unusable when AFC memories are used since we did not know the successful frequency mode to filter on. With on-demand memories, we can again just apply the control beam slightly later, allowing us to still perform the entanglement swap. For practical purposes we would thus prefer the on-demand full-AFC memory protocol to be used, but boosting the retrieval efficiency becomes even harder here due to the additional noise introduced by the two π-pulses that need to be applied.

We would like to conclude by noting that even though the high-efficiency EIT based memories will currently have a better performance when embedded in quantum repeater chains, it is not clear how they will scale up the multi-mode storage capabilities in the future. The most promising candidate of storing 10^6 modes, which is a requirement for reliable long-range quantum communication, are still AFC based QMs with combined forms of multiplexing, although increasing the retrieval efficiency still remains a challenging aspect here. We think that in both cases a lot of promising work is still left to be done and for now only time can tell which type will ultimately cause the overall best quantum repeater performance.

4 Repeater Location Optimization

In the context of repeater chain analysis, irrespective of whether atomic ensembles or another type of system is used, the focus is mostly on optimizing the parameters for a one dimensional chain of elementary links. However, if we want to proceed in building a large scale quantum network and move from one to two dimensional topologies, there will be more parties than just Alice and Bob that we want to connect. In particular, if we have a set of end nodes that all want to communicate with each other via bipartite entanglement, it would not be very cost-effective if we consider pairs of end nodes independently and start to build repeater nodes in an optimized way per pair. Therefore we will consider a more elaborate setting, where we minimize the overall number of repeaters in a two dimensional topology with a set of source-destination pairs that we want to connect. A natural mathematical way of modeling this problem is by using linear programs, which is the focus of this section.

We will start with a detailed description of the problem setting we consider, together with the necessary assumptions we make. Next, we will describe two formulations we have found to be suitable for modeling this problem. On the one hand we consider an intuitive set-covering formulation where the variables represent all possible end-to-end paths in a given graph. While this allows for a wide range of possible constraints, the number of variables scale exponentially with the input size, i.e. the number of nodes in the graph. On the other hand we also give a formulation where the variables represent elementary links, of which the number scales polynomially with the input size. However, we are rather limited in terms of constraints here, since we can only define a cost per elementary link and are not able to consider end-to-end figures of merit like the secret key rate. Finally, we will use two realistic data sets and a linear programming solver in order to analyze the performance of our model.

4.1 Problem Description

Consider a graph $\mathcal{G}(\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{E}$ is the set of undirected edges that connect (a subset of) the nodes. In order to model end nodes separately from repeater nodes, we partition the set $\mathcal{N}$ into a set of consumers, i.e. end nodes, $\mathcal{C}$ and possible repeater node locations $\mathcal{R}$, such that $\mathcal{R} = \mathcal{N} \setminus \mathcal{C}$. We choose to use any node in $\mathcal{R}$ in one of two ways: either we place a quantum repeater, or nothing is done here and the incoming elementary link is simply extended. We will refer to the latter as *link extension*. Although this assumption might be somewhat simplistic, it is in principle realizable with an optical switch that routes incoming photons to the correct destination based on classical information accompanied in the wave packet. We furthermore (weakly) assume that our quantum repeaters are based on atomic ensembles and that the spectrally multiplexed protocol of Tittel et al. is used.

One can imagine that building, calibrating and maintaining a quantum repeater can be a costly endeavour, while link extension is practically free. Therefore, the overall goal of this problem setting will be to minimize the total number of repeaters under a set of constraints. These constraints relate to the requests that a realistic source-destination pair can impose on their connection quality, such as that the end-to-end fidelity and entanglement generation rate should satisfy a certain minimum. This can be translated to a restriction on the number of entanglement swaps and on the maximum length of the elementary links that a repeater chain from source to destination can have.

We could for example use a metric from the analytical model of Guha et al. for the maximum length of an elementary link, which the authors refer to as L'. This denotes the maximum length at which the entanglement generation probability is close to 1, under the assumption that there are no detector dark counts. It is given by

$$L' = \frac{\log\left(\frac{M\eta_e^2}{2}\right)}{\alpha}. \tag{4.1}$$

This is a quite conservative metric since it should hold that $M\eta_e^2 \geq 2e$ in order to get a non-zero length where $P_{\text{el}}^{\text{mp}} \approx 1$. For example, for 10 spectral modes with perfect photon detectors and an attenuation coefficient of 0.2 dB/km, we only get a maximum length of 8 km. Nonetheless, one could use this metric to translate the component parameters to a maximum length for the elementary links which guarantees a near-unit success probability per elementary link. In Figure 3.2 we have seen that the more repeaters are used, the less the maximum achievable secret key rate is, and furthermore if we have imperfect detectors and visibility, we can expect that each entanglement swap will degrade the fidelity of our end-to-end state. Therefore, the maximum elementary link length and the number of swaps, i.e. repeaters, can indirectly impose a minimum on the entanglement generation rate and fidelity. We furthermore assume that all the component parameters for a repeater chain as discussed in the previous sections are fixed and given in the rest of this section.

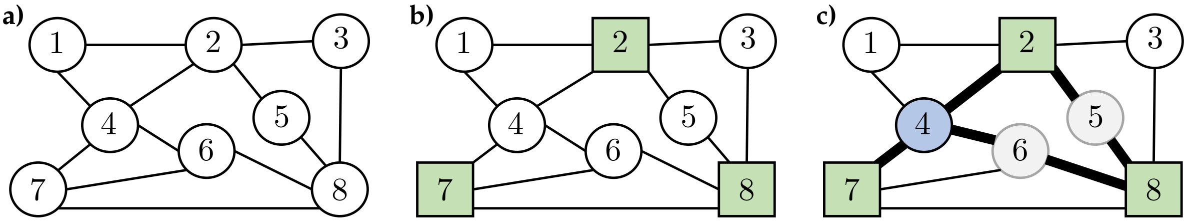

Figure 4.1: Example of our problem setting and solution. **a)-b)** We start with a graph $\mathcal{G}(\mathcal{N}, \mathcal{E})$, after which we partition $\mathcal{N}$ into a set of consumers $\mathcal{C} = \{2, 7, 8\}$ (squares) and possible repeater node locations $\mathcal{R} = \mathcal{N} \setminus \mathcal{C}$. We then solve the problem of connecting all the customers while minimizing the number of repeaters with some additional constraints, see Section 4.2. **c)** Example of a solution, where the consumers are connected via the highlighted paths. Node 4 is chosen as a repeater for connecting node 7 to nodes 2 and 8, while nodes 5 and 6 are chosen as link extension nodes. Nodes 1 and 3 can in principle be removed from the graph as they are not used in any of the paths.

A graphical illustration of our problem setting and an example of a solution is displayed in Figure 4.1. Given a set of nodes, we first choose the partitioning of $\mathcal{N}$ into $\mathcal{C}$ and $\mathcal{R}$, depending on the graph and the input. In case of the small-scale Dutch quantum network, we could for example partition the nodes of a graph of the Netherlands such that the consumers are $\mathcal{C} = \{\text{Amsterdam}, \text{Delft}, \text{Leiden}, \text{The Hague}\}$ and the possible repeater locations $\mathcal{R}$ are the intermediate nodes that are used to connect these cities.

Our problem setting and formulations share some similarities with the work in [69], although there the authors only consider the possibility of using either a single type of classical signal booster or a direct link between source destination pairs. In contrast, we allow a single pair to use multiple repeaters. Furthermore, in [69] a cost is associated to building a repeater, and the objective is to minimize the combined cost of opening a repeater and traversing to and away from it, while we minimize the total number of used repeaters. Our problem setting also shares some similarities with the uncapacitated facility location problem [70], where the main difference is that we simultaneously require the connection of a set of source-destination pairs.

4.2 Linear Programming Formulations

In this section we will give two types of linear integer programming formulations that can be used to solve the repeater location problem. The variables either represent end-to-end paths or elementary links, and they both have their trade-offs in terms of scalability and freedom of choice for the costs and objective function. However, both share some sets and parameters we will define here.

First, we can convert our set of consumers $\mathcal{C}$ to a set of unique *pairs* of end nodes $\mathcal{Q}$. We will denote the number of unique pairs as $n_q = |\mathcal{C}|(|\mathcal{C}| - 1)/2$, while the number of intermediate nodes is $n_r = |\mathcal{R}|$. We assume that any end nodes that are not in a given pair can only be used for link extension and not as repeater nodes. A pair $q \in \mathcal{Q}$ consists of a tuple with a fixed source s and destination t, and we assume that communication in the reverse direction follows the same path. In order to model the yes-no choice for the placement of a repeater at node $i \in \mathcal{R}$, we use the binary variables y_i, where a 1 represents that the node is used as a repeater and a 0 implies the extension of an elementary link. The only relevant cost parameter in our graph is the length c_{ij} of edge $(i, j) \in \mathcal{E}$ in km. Finally, we define $R^q_{\max}$ and $L^q_{\max}$ as respectively the maximum number of repeaters and the maximum length of an elementary link a pair $q \in \mathcal{Q}$ imposes. These indirectly represent the minimum fidelity and entanglement generation rate that a pair might require as we discussed above. Each consumer pair can individually choose $R^q_{\max}$ and $L^q_{\max}$, which can affect both the feasibility of a solution and the chosen paths.

With the help of a linear programming solver, we then seek the minimum number of repeaters such that all source-destination pairs are connected. The optimal solution subsequently directly assigns all nodes in $\mathcal{R}$ to be either repeaters or link extension nodes. One can then both derive where we should physically place our repeaters as well as which paths the consumers should take to communicate with one another. Note that if a link extension node is not actually used on any of the chosen paths, the node is left unused which implies that it can be deleted from the solution.

4.2.1 Path-based Formulation

Let us start with a formulation for the allocation of repeater nodes based on set-covering. Say we have a set $\mathcal{P}$ that contains all feasible paths p which start at $s \in q$ and end at $t \in q$ for all $q \in \mathcal{Q}$, the goal is to find the optimal subset $\mathcal{P}^*$ that optimizes the objectives function in the presence of a set of linear constraints. A path in itself is simply defined as a set of nodes and edges that connect these nodes. We assign a binary variable x_p to each path p, which has value 1 if it is used in a solution of our decision problem and 0 otherwise. The power of this formulation lies in the freedom of defining the paths and assigning costs to them, while it lacks polynomial scalability with respect to the input size of the problem.

In order to fully construct our formulation, let us define some additional parameters. Let c_p be the cost of using path $p \in \mathcal{P}$, which can simply be its total length $\sum_{(i,j)\in p} c_{ij}$, or a non-linear cost like the secret key rate, which also depends on the number of repeaters on the path. Since our path starts and ends at the end nodes of a pair by construction, these costs can conveniently cover any repeater usage in between. Since we want to keep track of the total number of used repeaters in our graph, let r_{ip} be a binary parameter that has value 1 if we use node $i \in \mathcal{R}$ as a repeater on path $p \in \mathcal{P}$ and is 0 otherwise. Note that $\sum_{i\in R} r_{ip}x_p$ will then correspond to the total number of repeaters in the path. Furthermore, to ensure that a pair is connected by a single path, we define the binary parameter a_{pq} that has value 1 if and only if path $p \in \mathcal{P}$ connects a given pair $q \in \mathcal{Q}$. Next, let l_p be the length of the largest elementary link that is used in a path $p \in \mathcal{P}$, which we can use to disallow the usage of some paths. Finally, in order to ensure multiple paths can use the same repeater we define the dummy parameter $D \geq n_q$ which can be an arbitrarily large number. We will explain its use below. Our path-based formulation (PBF) is given by

$$
\begin{aligned}
\min \quad & \sum_{i \in \mathcal{R}} y_i && && (4.2)\\
\text{s.t.} \quad & \sum_{p \in \mathcal{P}} a_{pq} x_p = 1 && \forall q \in \mathcal{Q} && (4.3)\\
& \sum_{p \in \mathcal{P}} r_{ip} x_p \leq D y_i && \forall i \in \mathcal{R} && (4.4)\\
& \sum_{p \in \mathcal{P}} a_{pq} l_p x_p \leq L^q_{\text{max}} && \forall q \in \mathcal{Q} && (4.5)\\
& \sum_{i \in \mathcal{R}} \sum_{p \in \mathcal{P}} a_{pq} r_{ip} x_p \leq R^q_{\text{max}} && \forall q \in \mathcal{Q} && (4.6)\\
& x_p \in \{0, 1\} && \forall p \in \mathcal{P} && (4.7)\\
& y_i \in \{0, 1\} && \forall i \in \mathcal{R} && (4.8)
\end{aligned}
$$

The objective function (4.2) minimizes the total number of quantum repeaters that are used. Constraints (4.3) ensure that each pair $q \in \mathcal{Q}$ is connected exactly once, since it only holds when a single path p is chosen ($x_p = 1$) that starts at $s \in q$ and ends at $t \in q$ ($a_{pq} = 1$) for a given $q \in \mathcal{Q}$. Next, Constraints (4.4) link the two types of decision variables and allows for the shared usage of repeaters. For example if there is a certain (s, t) pair that uses a path with a repeater placed at node $i \in \mathcal{R}$, then the corresponding y_i must also have value 1. Conversely, when a certain y_i variable has value 1, then any other (s, t) pair can use a the same repeater node at $i \in \mathcal{R}$. Note that setting $D = n_q$ is sufficient and yields that at most every consumer pair can use the same node as repeater. Next, Constraints (4.5) and (4.6) respectively put an upper bound on the length of the largest elementary link and the number of repeaters in a path, per source-destination pair. If a path is chosen that connects a certain source-destination pair, given by the product $a_{pq}x_p$, then the maximum elementary link length l_p and number of used repeaters on this path $\sum_{i \in \mathcal{R}} r_{ip}$ are both bounded from above. Note that if L^q_{max} is smaller than the length of the shortest edge that leaves $s \in q$, we know a priori that no feasible solution is possible. Finally, Constraints (4.15) and (4.16) simply represent that both our decision variables are binary.

Constraints (4.5) and (4.6) indirectly form a lower bound on the secret key rate that a consumer pair might have. However, with this formulation we are in principle able to directly impose restrictions on the secret key rate. For example, when c_p denotes the secret key rate of path $p \in \mathcal{P}$, then we can include the constraints

$$
\sum_{p \in \mathcal{P}} a_{pq} c_p x_p \geq S^q_{\text{min}} \quad \forall q \in \mathcal{Q}, \tag{4.9}
$$

where S^q_{min} is a lower bound on the secret key rate that the pair $q \in \mathcal{Q}$ would like to achieve. This implies that the PBF can also be used for general and even hybrid quantum repeater chains, including those based on nitrogen-vacancy centers in diamond or trapped ions, as long as one can define the cost per path c_p. However, as we will see in Section 4.2.2, this is not possible with our second formulation.

Complexity

Unfortunately this formulation becomes intractable to use on larger networks due to its scaling. We have to enumerate every possible path in our graph per pair, which is $\mathcal{O}\left(n_q n_r!\right)$ [71], and for each path we should additionally consider every possible permutation of repeater placements from 0 up to $R^q_{\max}$. For a fully connected graph with just 20 nodes and a single source-destination pair, we would thus have more than 10^{18} variables.

However, this approach of enumerating all paths is definitely not the most efficient one, and we could improve this by focusing on finding 'promising' paths. This is the essence of column generation [72], which roughly works as follows. We start with a small subset of paths $\tilde{\mathcal{P}}$, which contains for example only the shortest (s,t) paths that are feasible for the given values of $R^q_{\max}$ and $L^q_{\max}$ (if possible). This enables us to start with a feasible, but probably non-optimal solution. Then, provided this solution we solve a sub-problem based on the variables of the dual formulation which yields one or more paths we can add to our set $\tilde{\mathcal{P}}$ to get closer to the global optimum. This process can be iterated until a certain condition is met, after which we have reached the global optimal solution without enumerating all possible paths, and hence $\mathcal{P}^* \subseteq \tilde{\mathcal{P}}$. However, this sub-problem is generally an NP-hard problem in itself, such that heuristics are mostly used to find good solutions. For now we will consider the inclusion of column generation as beyond the scope for a physics thesis, and we will not use the PBF to generate results. We have presented it nonetheless due to its intuitive interpretation and flexibility in costs and constraints.

4.2.2 Link-based Formulation

Fortunately we have also found an alternative formulation for minimizing the number of repeaters. In this case we use the same idea of a one dimensional repeater chain, where the repeaters connect two neighbouring elementary links, by defining the binary variables $\tilde{x}^q_{ij}$. These have value 1 if source-destination pair q uses an elementary link from $i \in \mathcal{N}$ to $j \in \mathcal{N}$, and value 0 otherwise. Here we make the assumption that such an elementary link is formed by taking the *shortest path* from i to j, where $i \neq j$, which could in principle visit multiple other nodes. We are agnostic about whether the BSM is performed at one of the intermediate nodes (if any) or always exactly half way between the two nodes that span the elementary link. It is convenient to then also define a set of virtual edges $\mathcal{E}_v$, where each edge from i to j represents an elementary link formed by the shortest path from i to j. We define the cost of an elementary link as $\tilde{c}_{ij}$, which is simply the sum of the length of all the edges that are used in forming the link. With these decision variables, parameters and costs we define our link-based formulation (LBF) as

$$\min \quad \sum_{i\in\mathcal{R}} y_i \tag{4.10}$$

$$\text{s.t.} \quad \sum_{j\in\mathcal{R}\cup\{t\in q\}} \tilde{x}^q_{ij} - \sum_{j\in\mathcal{R}\cup\{s\in q\}} \tilde{x}^q_{ji} = \begin{cases} 1, & \text{if } i = s \in q \\ -1, & \text{if } i = t \in q \\ 0, & \text{if } i \in \mathcal{R} \end{cases} \qquad \forall q \in \mathcal{Q} \tag{4.11}$$

$$\sum_{q\in\mathcal{Q}} \sum_{i\in\mathcal{R}\cup\{s\in q\}} \tilde{x}^q_{ij} \leq D y_j \qquad \forall j \in \mathcal{R} \tag{4.12}$$

$$\tilde{c}_{ij}\tilde{x}^q_{ij} \leq L^q_{\max} \qquad \forall (i,j) \in \mathcal{E}_v, q \in \mathcal{Q} \tag{4.13}$$

$$\sum_{i\in\mathcal{R}\cup\{s\in q\}} \sum_{j\in\mathcal{R}} \tilde{x}^q_{ij} \leq R^q_{\max} \qquad \forall q \in \mathcal{Q} \tag{4.14}$$

$$\tilde{x}^q_{ij} \in \{0,1\} \qquad \forall (i,j) \in \mathcal{E}_v, q \in \mathcal{Q} \tag{4.15}$$

$$y_i \in \{0,1\} \qquad \forall i \in \mathcal{R} \tag{4.16}$$

The objective function (4.10) remains unchanged relative to the PBF and minimizes the total number of quantum repeaters that are used. Constraints (4.11) ensure that for every source-destination pair $q \in \mathcal{Q}$ there is always exactly one elementary link that starts at $s \in q$ and one that ends at $t \in q$, while an elementary link that ends at a repeater node is always followed by another elementary link. For example, when $i = s$ it should always hold that $\sum_{j\in\mathcal{R}\cup\{t\in q\}} \tilde{x}^q_{sj} = 1$, while $\sum_{j\in\mathcal{R}\cup\{s\in q\}} \tilde{x}^q_{js} = 0$ and vice versa for $i = t$. Otherwise, if $i \in \mathcal{R}$ each incoming elementary link should be connected by an outgoing elementary link, such that the LHS of (4.11) should always be $1 - 1 = 0$. Next, Constraints (4.12) link the two types of decision variables and allow for the shared use of repeater nodes. The dummy variable D has the same interpretation as in the PBF and it is sufficient to set $D = n_q$. Constraints (4.13) and (4.14) respectively impose restrictions on the maximum length of an elementary link and the number of repeaters, i.e. entanglement swaps, that a pair q demands. Note that for a path with N elementary links, we have $N - 1$ repeaters, such that we exclude any elementary link that ends at our destination $t \in q$ in the summation on the LHS of (4.6). Finally, Constraints (4.15) and (4.16) again represent that both types of decision variables are binary.

Complexity

The power of this formulation lies in its scalability. The total number of variables is namely

$$n_r + n_q(n_r(n_r+1)+1), \tag{4.17}$$

which comes from the fact that next to the n_r variables y_i, we must find the shortest path for every node in $\mathcal{R}$ to any other node in $\mathcal{R}$ plus the sink $t \in q$, as well as the path directly from $s \in q$ to $t \in q$. Our LBF thus has a polynomial scaling of $\mathcal{O}\left(n_q n_r^2\right)$ with respect to the input size of the problem. Furthermore, in order to find the shortest paths for the elementary links and hence construct the set $\mathcal{E}_v$ from $\mathcal{E}$, we can use Dijkstra's algorithm, which also results in a polynomial time complexity of $\mathcal{O}\left(n_r^4\right)$ for finding all the variables. However, we cannot impose a lower bound on an end-to-end metric like the highly non-linear secret key rate in this case, since we are restricted to linear constraints related to elementary links.

With this formulation it is possible to get cyclic paths in an optimal solution that are not used in the path from s to t, in case the number of repeaters on a path has some slack compared to $R^q_{\max}$. For example, consider the case in which the maximum number of repeaters for a certain (s,t) pair is 6, the optimal path from s to t uses 2 repeaters and the overall minimum number of repeaters is 4. The solution for this pair could then also contain an elementary link from $i \in \mathcal{R}$ to $j \in \mathcal{R}$ and two or more links that form a path from $j \in \mathcal{R}$ to $i \in \mathcal{R}$, where i and j are both not in the path from s to t, simply because there is no penalty imposed on doing so. A solution to this problem would be to introduce a weighted cost in the objective function on the elementary link usage, which we will elaborate on in the next subsection. Without this adjusted objective, we can just delete all cyclic paths in our solution a posteriori. Before we conclude this subsection, we would like to note that the comparison and proving the equivalence of the LBF and PBF are also beyond the scope of this work, and are aimed to be discussed in an upcoming paper.

4.2.3 Multi-objective Function

Finding the minimum number of repeaters in a graph is an objective on itself, although in general the optimal solution will be highly degenerate. When a solution is found for our LBF, the path formed by the elementary links between the source-destination pairs might not necessarily be the shortest ones, since there is no incentive for the solver to choose these instead of longer paths. However, there is a way we can add this incentive by modifying the *single-objective* function (4.10) to a *multi-objective* function

$$\min \quad \sum_{i \in \mathcal{R}} y_i + \beta \sum_{q \in \mathcal{Q}} \sum_{(i,j) \in \mathcal{E}_v} \tilde{c}_{ij} \tilde{x}^q_{ij}, \tag{4.18}$$

where β is a tunable parameter of which the value should be carefully chosen. If it is set to a value which is too large, such that $\beta \sum_{q \in \mathcal{Q}} \sum_{(i,j) \in \mathcal{E}_v} \tilde{c}_{ij} \tilde{x}^q_{ij} > 1$, it can influence the minimum number of repeaters which we want to prevent. On the other hand, if we choose a very small value for β the solver could encounter problems in the numerical precision of the objective coefficients. We will therefore have to empirically choose a suitable value of β which forms a balance between these two extremes. Note that the use of this multi-objective function should also result in the removal of the cyclic paths, since the usage of these will now result in a higher objective value. One could also use a multi-objective function which for example minimizes the length of the longest elementary link per source-destination pair, but we choose to not include this in our model.

4.3 Results

In order to test the performance of our LBF, we use two open source network topologies from the Topology Zoo database [73]. The first, Surfnet (see Figure 4.2a)), is a data set of a fiber network in the Netherlands from 2010 that contains 50 nodes and 68 edges, of which the length ranges from 3 to 112 km, with an average length of 32 km. We assign the consumers to be $\mathcal{C} = \{\text{Groningen}, \text{Maastricht}, \text{The Hague}, \text{Vlissingen}\}$, since these lie in the outer regions of the Netherlands. It was not interesting to choose the four cities that want to be connected in the Netherlands for the small-scale quantum internet realization as mentioned in Section 4.1, since there were not enough nodes in between to be used as repeaters in this data set. We also use a second data set of a European-scale telecom network called Colt (see Figure 4.2b)). This is a larger graph with 149 nodes and 167 edges, of which the length ranges from 3 to 1237 km, with an average length of 164 km. For this data set we will choose the consumers to be the locations of which at least one university involved in QIA is located, such that $\mathcal{C} = \{\text{Barcelona}, \text{Basel}, \text{Copenhagen}, \text{Geneva}, \text{Innsbruck}, \text{Lisbon}, \text{Paris}, \text{Stuttgart}, \text{The Hague}\}$. Note that we have assumed the location of QuTech to be in The Hague since Delft was not available as a node in the data set. Furthermore we also excluded the university located at Garching (Max Planck Institute of Quantum Optics), which is close to Munich, since there are very few intermediate nodes available in the region together with Stuttgart, Innsbruck and Basel.

We choose to use the same constraints on the maximum elementary link length and number of repeaters for all pairs, i.e. $L^q_{\max} = L_{\max}, R^q_{\max} = R_{\max} \; \forall q \in \mathcal{Q}$. To generate these results we use a laptop with a 2.8 GHz processor and 8 GB of RAM, and as linear program solver we use the Python API for IBM ILOG CPLEX version 12.9.0. Since we are dealing with a binary linear optimization problem, finding a solution is known to be NP-complete [74]. This implies that there is no (known) algorithm that is able to solve this problem in polynomial time. CPLEX therefore mainly relies on a set of highly optimized branch-and-bound-and-cut algorithms to quickly find an optimal solution [75].

Let us start with the Surfnet data set. Constructing the graph $\mathcal{E}_v$ from $\mathcal{E}$ takes just a few seconds, and the same subsequently holds for finding the global optimal solution with CPLEX. The results for varying values of $L_{\max}$ are shown in Figure 4.3a). Here we can see that as the allowed maximum length for a repeater increases, the optimal number of repeaters also decreases. This is as we would expect, since large elementary links allow for the usage of less repeaters. If $L_{\max} \to \infty$, the objective value will decrease to 0, since all source-destination pairs will simply be connected by single elementary links. We can also see that the usage of less repeaters can sometimes cause an increase in the average path length per source-destination pair, which implies that some pairs have to take longer paths than before in order to reach the repeater of which the use is shared among multiple parties. For the multi-objective function (4.18) we find that $\beta = 1/2500$ is a suitable value to use in terms of the trade-off we mention above. We can see that the average path length is always equal to or shorter compared to the

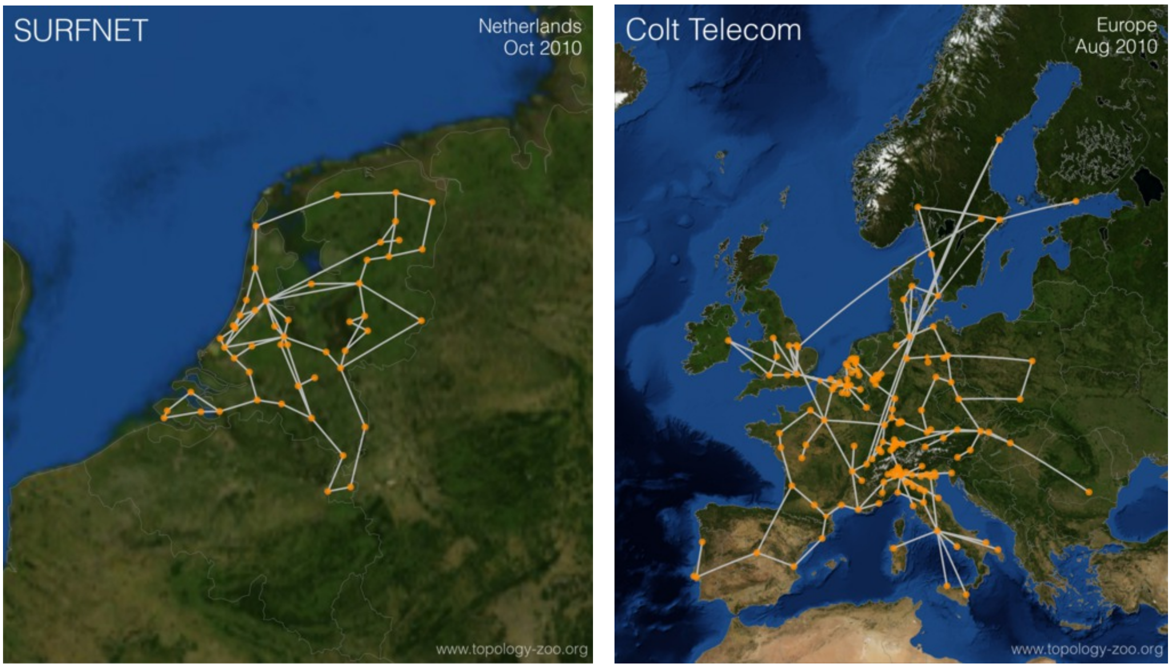

a) **b)**

Figure 4.2: Graphical representation of **a)** the Surfnet data of the Netherlands and **b)** the Colt data set of Europe. Both images are directly retrieved the Topology Zoo [73].

single-objective function, while the minimum number of repeaters remains unchanged. This implies that the solution for the minimum number of repeaters is indeed degenerate and we can find solutions with shorter paths per source-destination pair but with the same minimum number of repeaters.

The effects of changes in the values of both L_{max} and R_{max} are shown in Figure 4.3b). We can see that if the maximum elementary link length is set to a low value, we must use a large number of repeaters to generate a feasible solution. On the other hand, if we allow for large elementary links to be used, we require less repeaters for a feasible solution. Once a feasible solution has been found for a given value of L_{max} it remains unaffected when R_{max} is increased. This is in line with our expectation, since we are minimizing the overall number of repeaters.

For the Colt data set, we can see similar behaviour as shown in Figures 4.3c) and 4.3d). In this case the construction of the graph $\mathcal{E}_v$ and the linear program takes about 20 minutes and an optimal solution is mostly found within a few minutes for the multi-objective function, even though we have over $7 \cdot 10^5$ variables (see (4.17)). This shows the power of the polynomial scaling of our LBF in combination with a fast linear programming solver and its applicability to large graphs. In this case $\beta = 1/75000$ turns out to be a suitable value in order to correctly nudge our objective function towards the usage of shorter paths without affecting the minimum number of repeaters. We can see similar behaviour for the feasibility as function of L_{max} and R_{max} as in the Surfnet data set. The main difference is that for this region for L_{max}, we generally need less repeaters for a feasible solution to be found.

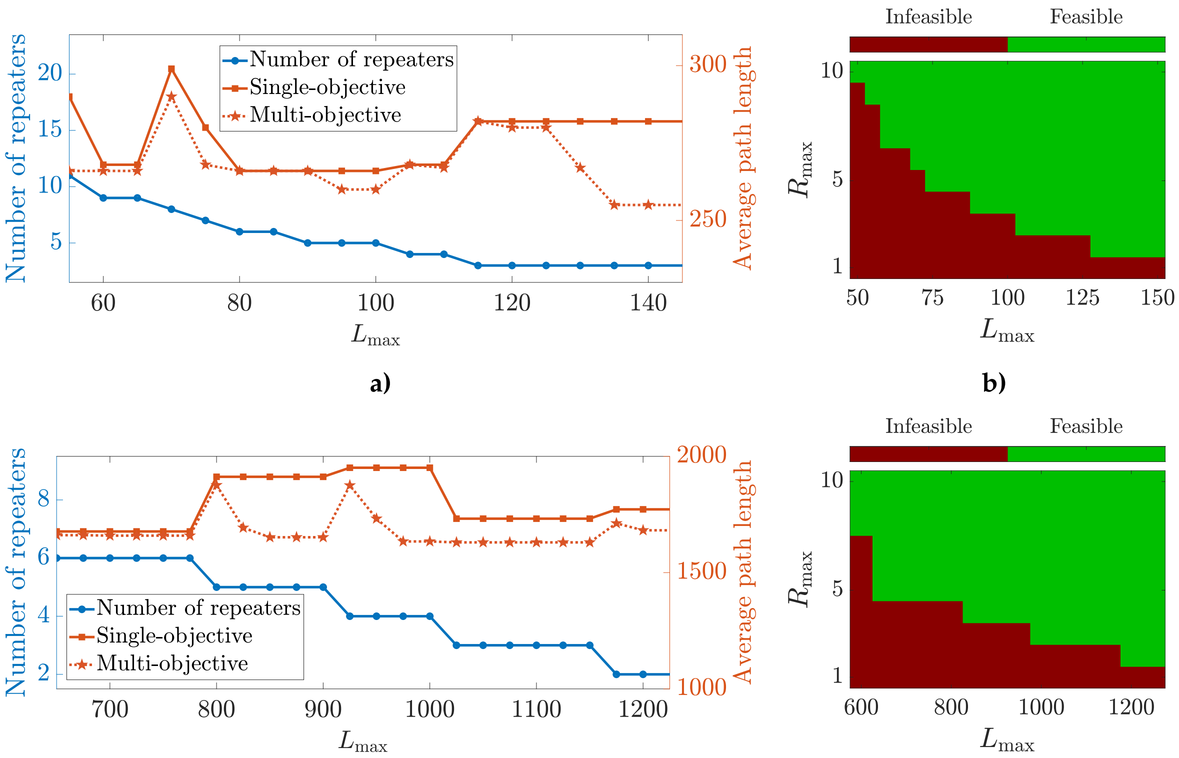

a) **b)** **c)** **d)**

Figure 4.3: Results for the LBF. **a), c)** Total number of repeaters and the average end-to-end path length for our LBF with the single and multi-objective function for the **a)** Surfnet data set with $R_{\max} = 11$ and **c)** Colt data set for $R_{\max} = 6$. Here, the average path length means an average over the number of (s, t) pairs n_q. We see that as $L_{\max}$ increases, the minimum number of repeaters decreases since longer elementary links are allowed to be used. However, the decrease in the number of repeaters can sometimes cause an increase in the average path length. The use of the multi-objective function does not affect the minimal number of repeaters (not visible), while it always causes the same or shorter paths to be chosen. **b), d)** Solution feasibility as a function of $L_{\max}$ and $R_{\max}$ for the **b)** Surfnet and **d)** Colt data set. We see that as the number of repeaters increases, we can find a feasible solution with shorter paths and vice-versa. Once we are in the feasible region for a given $L_{\max}$, the optimal objective value remains unaffected when $R_{\max}$ increases.

An optimal solution for the Colt data set is shown in Figure 4.4. Here we can see that the repeater nodes are chosen in paths that need to span larger distances, while source-destination pairs that are close to each other are directly connected by a single elementary link.

We have shown that our LBF is able to solve the optimized positioning of quantum repeaters within a large scale network in reasonable time, and that the outcomes are as expected. The next step forward would be to make our model more elaborate, where the simulation model of Section 3 could for example be used to find parameters such as $L^q_{\max}$ and $R^q_{\max}$, or even compute the secret key rate per source-destination pair based on the optimal solution of the repeater positioning in a graph.

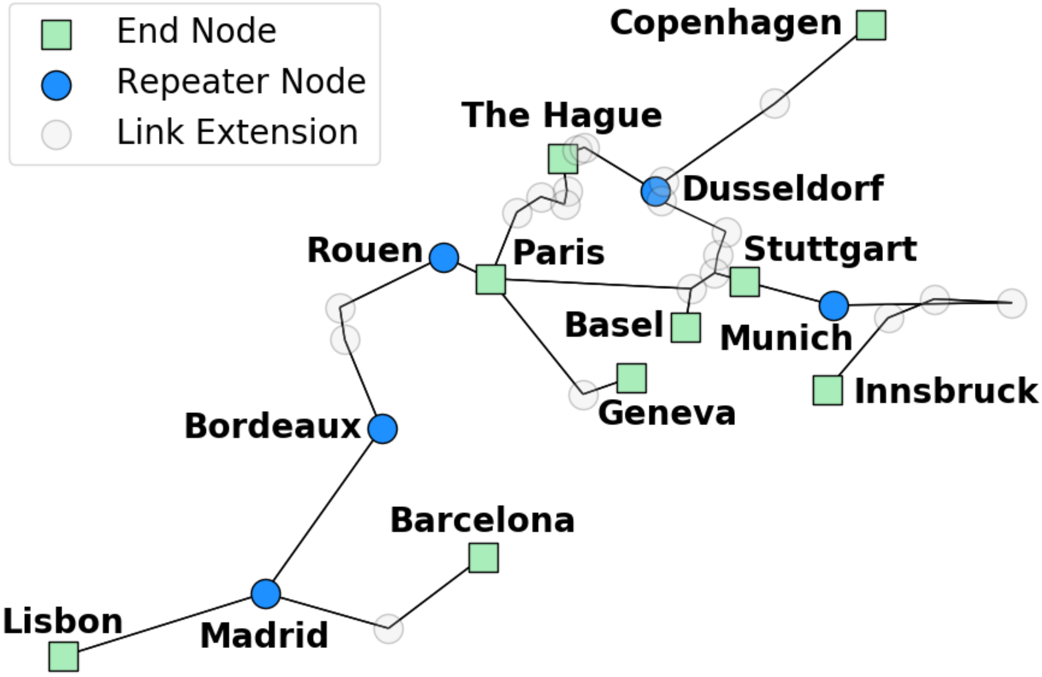

Figure 4.4: Solution of the LBF for the Colt data set for $L_{\max} = 900$, $R_{\max} = 6$ and $\beta = 1/75000$, with cities that are involved in QIA chosen as end nodes. We can see that Bordeaux, Dusseldorf, Madrid, Munich and Rouen are used as a repeater and the optimal solution thus uses 5 repeaters in total. Due to the multi-objective function, the shortest paths are taken for each source-destination pair. Note that any end node that is not contained in a pair $q \in \mathcal{Q}$ can solely be used for link extension.

5 Conclusions and Outlook

Quantum repeaters are required for the realization of a quantum internet because they make it possible to overcome the exponential loss of photons in optical fiber. In this thesis we have contributed on a versatile and modular simulation model for quantum repeater chains based on atomic ensembles. The diverse set of abstract components within this model allows for a convenient way of analyzing different types of photon encodings, channels, detectors, photon sources and memories. In this work we have shown the validation of the simulation model, as well as the effects of non-perfect photon indistinguishability and realistic photon pair sources.

Due to the modularity of our model and the use of NetSquid as its lower level engine, one can investigate a wide range of parameters and their effects on the performance of a quantum repeater chain. Future research could for example include temporal multiplexing, time dependent retrieval efficiency in the memory or emission noise on the output state of the source. One could even investigate the effect of using a completely different type of photon source. Experimental realizations have been reported of near-deterministic single photon sources based on quantum dots [76], which can be used in a presence-absence encoded setup [16].

In this thesis we have analyzed the performance of quantum repeater chains under a set of fixed parameters, but another interesting research topic would be to optimize the parameters in order to maximize the performance. However, the large number of tunable, often real-valued parameters make the combined search space huge, such that one would probably have to resort to fixing a subset of parameters after which local maxima can be sought. On the other hand, we could also combine some parameters to make our model more abstract and the search space smaller. For example, the detector efficiency, dark count rate and visibility could be combined into a single 'BSM quality' parameter and similarly for the other components. This will also make it easier to analyze the performance of hybrid repeater chains, where the end nodes could for example hold nitrogen-vacancy centers that can be used for distributed quantum computing.

We have also focused on the physics of QMs by giving an elaborate description of two atomic ensemble based protocols: AFC and EIT. While AFC memories are sometimes presented as a promising type of memory with a high multi-mode capacity, the current experimental realizations lack a practical level of efficiency. The use of EIT based memories therefore currently leads to a better performance when embedded in quantum repeater chains, due to their high efficiency, especially for parameters that are expected to be realized in the near future. However, it is not clear how EIT based memories will be able to scale up the numbers of multiplexing modes. In this thesis we have limited ourselves to just two atomic ensemble based QM protocols, while there also exist other protocols such as controlled reversible inhomogeneous broadening and gradient echo [31], of which the performance in repeater chains could also be characterized. In our case we have used the analytical model of Guha et al. [24] to generate the results for the memory comparison, so a natural extension would be to instead use our simulation model for this, with the inclusion of parameters that cannot be taken into account in the analytical one.

For the optimization of the positioning of repeaters, we have given two linear programming formulations with different flexibility and scaling. We have used the LBF to generate our results due its the polynomial scaling with respect to the input size. The results show that we are able to find the optimal repeater locations for two realistic graphs of different sizes. The use of a multi-objective function causes the optimizer to simultaneously choose the shortest paths between end nodes without significantly increasing the required computation time. We have also included a visualization of a solution for a European-scale graph, which shows the optimal locations of quantum repeaters for the QIA nodes. Another interesting use case, which could be analyzed in a follow-up research, would be to use optical fiber data of the Netherlands to find the best repeater allocation that is required for a national quantum internet.

Currently our problem setting for the optimized repeater allocation does not include any probabilistic behavior inherent to quantum mechanics. There exist methods to incorporate uncertainty in costs and demand with linear programming, which fall under the branch of robust optimization [77]. Provided a set of uncertain parameters, one could for example find an optimal solution for the worst case scenario or for the expected value of the probabilistic parameters in case historic data is known. This could ultimately lead to a coupling of our simulation model to the optimization of the repeater positioning, in order to determine the latter in a highly realistic setting.

All in all, we hope that this thesis will help with laying the groundwork for the numerical analysis of quantum repeater chains, including their positioning, and thereby contribute to establishing a European-scale quantum internet, or even beyond.

Acknowledgements

First and foremost I would like to thank all the (past) members of the NLBlueprint team. Especially the trip to Lisbon for the QIA consortium was a very cool experience. I very much enjoyed working in a team and it also definitely helped me during my thesis. In particular the members I shared a room with during most of my time at QuTech - Axel, David, Francisco, Guus and Tim - have made for a very chill and fun environment.

Stephanie's very broad and expert knowledge caused me to always have answers to any of the questions I encountered during my thesis. She also enabled me to always have clear long term goals to work towards, which I found very helpful during my thesis project. Also my daily supervisor Guus was always available for me to ask questions, no matter how small, and provided very useful feedback on my thesis report. I would also like to thank Filip for his supervision during the early months of my project.

Next, Kaushik has also been of great help during the fruitful discussions on the repeater location optimization as well as with his feedback on the corresponding chapter in my thesis. I hope we will be able to publish a nice paper in the coming month. Additionally I want to thank David and Gary for joining my thesis committee and I am looking forward to their questions at my defense.

Last but not least I want to sincerely express my gratitude towards my family and close friends for their continuous support. From providing me with food and motivation to having fun at bunker parties, it all contributed towards making this written result of a nine-month project into something of which I think I can be proud.

Appendix

A Computation of Errorbars for the Secret Key Rate

The computation of errorbars for the secret key rate is not straightforward. The two output statistics of our simulation we need in order to compute the secret key rate are $R(Q) = 1 - H(Q_x) - H(Q_z)$ and the number of attempts per success (NAPS), which should converge to

$$\left(\left(P_{\text{el}}^{\text{mp}}\right)^N \left(P_{\text{swap}}\right)^{N-1} P_{\text{meas}}\right)^{-1}, \tag{A.1}$$

if we use the same parameters as Guha et al. [24] (see Equation (2.10)). Both of these metrics contain statistical errors we should take into account. For the NAPS, we can directly compute the standard deviation σ_{NAPS} from our data, and similarly for the QBER in two bases σ_{Q_x} and σ_{Q_z}. However, computing the standard deviation of $R(Q)$ is less obvious, since $H(Q)$ is not defined for $Q = 0$ and $Q = 1$ and $R(Q)$ should never become negative, such that we actually compute $\max(R(Q), 0)$ for our secret key rate. In order to get a range of our secret key rate, we instead compute lower and upper bounds for this value.

The lower bound of the rate $\underline{R}$ is computed by

$$\underline{R} = \max\left(1 - \sum_{Q \in \{Q_x, Q_z\}} \max\left[H(Q), H(Q + \sigma_Q), H(Q - \sigma_Q)\right], 0\right), \tag{A.2}$$

and the upper bound $\overline{R}$ is computed by

$$\overline{R} = \max\left(1 - \sum_{Q \in \{Q_x, Q_z\}} \min\left[H(Q), H(Q + \sigma_Q), H(Q - \sigma_Q)\right], 0\right). \tag{A.3}$$

The standard errors of the secret key rate S are then estimated by using the product rule for the propagation of uncertainty [78]

$$\left(\frac{\epsilon_{\text{min}}}{\mu_S}\right)^2 \approx \left(\frac{\mu_R - \underline{R}}{\mu_R}\right)^2 + \left(\frac{\sigma_{\text{NAPS}}}{\mu_{\text{NAPS}}}\right)^2, \tag{A.4}$$

and

$$\left(\frac{\epsilon_{\text{max}}}{\mu_S}\right)^2 \approx \left(\frac{\overline{R} - \mu_R}{\mu_R}\right)^2 + \left(\frac{\sigma_{\text{NAPS}}}{\mu_{\text{NAPS}}}\right)^2, \tag{A.5}$$

where μ_X denotes the mean of random variable X. The errorbars in the figures of Section 3.2 are then $\mu_S + \epsilon_{\text{max}}$ and $\mu_S - \epsilon_{\text{min}}$. For example if $\underline{R} = 0$, then $\epsilon_{\text{min}} \approx \mu_S$, since $(\sigma_{\text{NAPS}}/\mu_{\text{NAPS}})^2$ is generally small for a sufficient amount of data. This implies that the lower bound on the secret key rate goes to 0. On the other hand, when $\underline{R}$ or $\overline{R}$ is equal to μ_R, the error of the secret key rate is solely determined by the standard deviation of the NAPS.

References

[1] Quantum Flagship. (2019). About - Quantum Technology, [Online]. Available: `https://qt.eu/about/`.

[2] European Commission. (2019). What is Horizon 2020, [Online]. Available: `https://ec.europa.eu/programmes/horizon2020/what-horizon-2020`.

[3] Quantum Internet Alliance. (2019). Quantum Internet Alliance, [Online]. Available: `http://quantum-internet.team/`.

[4] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead", *Science*, vol. 362, no. 6412, eaam9288, 2018.

[5] P. Komar, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, "A quantum network of clocks", *Nature Physics*, vol. 10, no. 8, p. 582, 2014.

[6] A. Broadbent, J. Fitzsimons, and E. Kashefi, "Universal blind quantum computation", in *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, IEEE, 2009, pp. 517–526.

[7] T. B. Bahder, "Quantum positioning system", *arXiv preprint quant-ph/0406126*, 2004.

[8] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing.", *Theor. Comput. Sci.*, vol. 560, no. 12, pp. 7–11, 2014.

[9] R. L. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems", *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.

[10] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer", *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.

[11] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, *et al.*, "Quantum supremacy using a programmable superconducting processor", *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.

[12] D. Dieks, "Communication by EPR devices", *Physics Letters A*, vol. 92, no. 6, pp. 271–272, 1982.

[13] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, "Quantum repeaters: The role of imperfect local operations in quantum communication", *Physical Review Letters*, vol. 81, no. 26, p. 5932, 1998.

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

[15] C. Simon, H. De Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, "Quantum repeaters with photon pair sources and multimode memories", *Physical review letters*, vol. 98, no. 19, p. 190 503, 2007.

[16] N. Sangouard, C. Simon, H. De Riedmatten, and N. Gisin, "Quantum repeaters based on atomic ensembles and linear optics", *Reviews of Modern Physics*, vol. 83, no. 1, p. 33, 2011.

[17] N. Sinclair, E. Saglamyurek, H. Mallahzadeh, J. A. Slater, M. George, R. Ricken, M. P. Hedges, D. Oblak, C. Simon, W. Sohler, *et al.*, "Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control", *Physical review letters*, vol. 113, no. 5, p. 053 603, 2014.

[18] F. Ewert and P. van Loock, "3/4-efficient bell measurement with passive linear optics and unentangled ancillae", *Physical review letters*, vol. 113, no. 14, p. 140 403, 2014.

[19] A. M. Brańczyk, "Hong-ou-mandel interference", *arXiv preprint arXiv:1711.00080*, 2017.

[20] C.-K. Hong, Z.-Y. Ou, and L. Mandel, "Measurement of subpicosecond time intervals between two photons by interference", *Physical review letters*, vol. 59, no. 18, p. 2044, 1987.

[21] A. Lamas-Linares, J. C. Howell, and D. Bouwmeester, "Stimulated emission of polarization-entangled photons", *Nature*, vol. 412, no. 6850, p. 887, 2001.

[22] S. Tanzilli, H. De Riedmatten, W. Tittel, H. Zbinden, P. Baldi, M. De Micheli, D. B. Ostrowsky, and N. Gisin, "Highly efficient photon-pair source using periodically poled lithium niobate waveguide", *Electronics Letters*, vol. 37, no. 1, pp. 26–28, 2001.

[23] H. Bechmann-Pasquinucci and N. Gisin, "Incoherent and coherent eavesdropping in the six-state protocol of quantum cryptography", *Physical Review A*, vol. 59, no. 6, p. 4238, 1999.

[24] S. Guha, H. Krovi, C. A. Fuchs, Z. Dutton, J. A. Slater, C. Simon, and W. Tittel, "Rate-loss analysis of an efficient quantum repeater architecture", *Physical Review A*, vol. 92, no. 2, p. 022 357, 2015.

[25] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, "Fundamental limits of repeaterless quantum communications", *Nature communications*, vol. 8, p. 15 043, 2017.

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

[27] L. Ma, O. Slattery, and X. Tang, "Optical quantum memory based on electromagnetically induced transparency", *Journal of Optics*, vol. 19, no. 4, p. 043 001, 2017.

[28] K. Heshami, D. G. England, P. C. Humphreys, P. J. Bustard, V. M. Acosta, J. Nunn, and B. J. Sussman, "Quantum memories: Emerging applications and recent advances", *Journal of modern optics*, vol. 63, no. 20, pp. 2005–2028, 2016.

[29] M. Afzelius, C. Simon, H. De Riedmatten, and N. Gisin, "Multimode quantum memory based on atomic frequency combs", *Physical Review A*, vol. 79, no. 5, p. 052 329, 2009.

[30] G. Liu and B. Jacquier, *Spectroscopic properties of rare earths in optical materials*. Springer Science & Business Media, 2006, vol. 83.

[31] M. Askarani, “Telecom-wavelength quantum memories in rare earth ion-doped materials for quantum repeaters”, PhD thesis, Delft University of Technology, 2019. DOI: `https://doi.org/10.4233/uuid:d418a98b-f2aa-4af3-b0e0-864875fcad2b`.

[32] M. Afzelius and C. Simon, “Impedance-matched cavity quantum memory”, *Physical Review A*, vol. 82, no. 2, p. 022 310, 2010.

[33] H. Krovi, S. Guha, Z. Dutton, J. A. Slater, C. Simon, and W. Tittel, “Practical quantum repeaters with parametric down-conversion sources”, *Applied Physics B*, vol. 122, no. 3, p. 52, 2016.

[34] A. Seri, D. Lago-Rivera, A. Lenhard, G. Corrielli, R. Osellame, M. Mazzera, and H. de Riedmatten, “Quantum storage of frequency-multiplexed heralded single photons”, *arXiv preprint arXiv:1902.06657*, 2019.

[35] P. Jobez, N. Timoney, C. Laplane, J. Etesse, A. Ferrier, P. Goldner, N. Gisin, and M. Afzelius, “Towards highly multimode optical quantum memory for quantum repeaters”, *Physical Review A*, vol. 93, no. 3, p. 032 327, 2016.

[36] M. Sabooni, Q. Li, S. Kröll, and L. Rippe, “Efficient quantum memory using a weakly absorbing sample”, *Physical review letters*, vol. 110, no. 13, p. 133 604, 2013.

[37] A. Holzäpfel, J. Etesse, K. T. Kaczmarek, A. Tiranov, N. Gisin, and M. 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.

[38] G. Amaral, Private communication, Nov. 14, 2019.

[39] J. Marangos and T. Halfmann, “Electromagnetically induced transparency”, in *Handbook of Optics by Optical Society of America*, McGraw-Hill, 2001, pp. 23–1.

[40] M. Fleischhauer, A. Imamoglu, and J. P. Marangos, “Electromagnetically induced transparency: Optics in coherent media”, *Reviews of modern physics*, vol. 77, no. 2, p. 633, 2005.

[41] M. Lukin, “Colloquium: Trapping and manipulating photon states in atomic ensembles”, *Reviews of Modern Physics*, vol. 75, no. 2, p. 457, 2003.

[42] S. Harris, “Electromagnetically Induced Transparency”, *Physics Today*, vol. 50, no. 7, p. 36, 1997. DOI: `https://doi.org/10.1063/1.881806`.

[43] S. E. Harris, J. Field, and A. Imamoğlu, “Nonlinear optical processes using electromagnetically induced transparency”, *Physical Review Letters*, vol. 64, no. 10, p. 1107, 1990.

[44] L. V. Hau, S. E. Harris, Z. Dutton, and C. H. Behroozi, “Light speed reduction to 17 metres per second in an ultracold atomic gas”, *Nature*, vol. 397, no. 6720, p. 594, 1999.

[45] H. Wu, M. Xiao, and J. Gea-Banacloche, “Evidence of lasing without inversion in a hot rubidium vapor under electromagnetically-induced-transparency conditions”, *Physical Review A*, vol. 78, no. 4, p. 041 802, 2008.

[46] H. Schmidt and R. Ram, "All-optical wavelength converter and switch based on electromagnetically induced transparency", *Applied Physics Letters*, vol. 76, no. 22, pp. 3173–3175, 2000.

[47] P. Lambropoulos and D. Petrosyan, *Fundamentals of Quantum Optics and Quantum Information*, 1st ed. Springer-Verlag Berlin Heidelberg, 2007, p. 325, ISBN: 978-3-540-34572-5. DOI: `10.1007/978-3-540-34572-5`.

[48] M. Born and V. Fock, "Beweis des adiabatensatzes", *Zeitschrift für Physik*, vol. 51, no. 3-4, pp. 165–180, 1928.

[49] P. Vernaz-Gris, K. Huang, M. Cao, A. S. Sheremet, and J. Laurat, "Highly-efficient quantum memory for polarization qubits in a spatially-multiplexed cold atomic ensemble", *Nature communications*, vol. 9, no. 1, p. 363, 2018.

[50] J. Nunn, K. Reim, K. Lee, V. Lorenz, B. Sussman, I. Walmsley, and D. Jaksch, "Multimode memories in atomic ensembles", *Physical review letters*, vol. 101, no. 26, p. 260 502, 2008.

[51] Y.-F. Hsiao, H.-S. Chen, P.-J. Tsai, and Y.-C. Chen, "Cold atomic media with ultrahigh optical depths", *Physical Review A*, vol. 90, no. 5, p. 055 401, 2014.

[52] S.-Y. Lan, A. Radnaev, O. Collins, D. Matsukevich, T. Kennedy, and A. Kuzmich, "A multiplexed quantum memory", *Optics express*, vol. 17, no. 16, pp. 13 639–13 645, 2009.

[53] G. Heinze, N. Rentzsch, and T. Halfmann, "Multiplexed image storage by electromagnetically induced transparency in a solid", *Physical Review A*, vol. 86, no. 5, p. 053 837, 2012.

[54] K. Surmacz, J. Nunn, K. Reim, K. Lee, V. Lorenz, B. Sussman, I. Walmsley, and D. Jaksch, "Efficient spatially resolved multimode quantum memory", *Physical Review A*, vol. 78, no. 3, p. 033 806, 2008.

[55] K. Akiba, K. Kashiwagi, M. Arikawa, and M. Kozuma, "Storage and retrieval of nonclassical photon pairs and conditional single photons generated by the parametric down-conversion process", *New Journal of Physics*, vol. 11, no. 1, p. 013 049, 2009.

[56] Y. Wang, J. Li, S. Zhang, K. Su, Y. Zhou, K. Liao, S. Du, H. Yan, and S.-L. Zhu, "Efficient quantum memory for single-photon polarization qubits", *Nature Photonics*, vol. 13, no. 5, p. 346, 2019.

[57] F. Hoffet, Private communication, Oct. 25, 2019.

[58] M. A. Nielsen and I. Chuang, *Quantum computation and quantum information*, 2002.

[59] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpędek, M. Pompili, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, *et al.*, "A link layer protocol for quantum networks", in *Proceedings of the ACM Special Interest Group on Data Communication*, ACM, 2019, pp. 159–173.

[60] QuTech, *NetSquid - The Network Simulator for Quantum Information using Discrete events*, `https://netsquid.org/`, 2019.

[61] J. de Oliveira Filho, Z. Papp, R. Djapic, and J. Oosteveen, "Model-based design of self-adapting networked signal processing systems", in *2013 IEEE 7th International Conference on Self-Adaptive and Self-Organizing Systems*, IEEE, 2013, pp. 41–50.

[62] D. Maier, "Scalability of a Quantum Repeater Protocol based on Atomic Ensembles", Master Thesis, Mar. 2020.

[63] T. C. Ralph and G. J. Pryde, "Optical quantum computation", in *Progress in optics*, vol. 54, Elsevier, 2010, pp. 209–269.

[64] C. Śliwa and K. Banaszek, "Conditional preparation of maximal polarization entanglement", *Phys. Rev. A*, vol. 67, p. 030 101, 3 Mar. 2003. DOI: `10.1103/PhysRevA.67.030101`. [Online]. Available: `https://link.aps.org/doi/10.1103/PhysRevA.67.030101`.

[65] S. Bravyi and A. Kitaev, "Universal quantum computation with ideal clifford gates and noisy ancillas", *Physical Review A*, vol. 71, no. 2, p. 022 316, 2005.

[66] T. Volz, A. Reinhard, M. Winger, A. Badolato, K. J. Hennessy, E. L. Hu, and A. Imamoğlu, "Ultrafast all-optical switching by single photons", *Nature Photonics*, vol. 6, no. 9, p. 605, 2012.

[67] O. Gittsovich, N. J. Beaudry, V. Narasimhachar, R. R. Alvarez, T. Moroder, and N. Lütkenhaus, "Squashing model for detectors and applications to quantum-key-distribution protocols", *Physical Review A*, vol. 89, no. 1, p. 012 325, 2014.

[68] G. Heinze, C. Hubrich, and T. Halfmann, "Stopped light and image storage by electromagnetically induced transparency up to the regime of one minute", *Physical review letters*, vol. 111, no. 3, p. 033 601, 2013.

[69] L. C. Monticone and G. Funk, "Application of the facility location problem to the problem of locating concentrators on an faa microwave system", *Annals of operations research*, vol. 50, no. 1, pp. 437–454, 1994.

[70] A. Alcouffe and G. Muratet, "Optimal location of plants", *Management Science*, vol. 23, no. 3, pp. 267–274, 1976.

[71] B. Roberts and D. P. Kroese, "Estimating the number of st paths in a graph.", *J. Graph Algorithms Appl.*, vol. 11, no. 1, pp. 195–214, 2007.

[72] G. Desaulniers, J. Desrosiers, and M. M. Solomon, *Column generation*. Springer Science & Business Media, 2006, vol. 5.

[73] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo", *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 9, pp. 1765–1775, Oct. 2011, ISSN: 0733-8716. DOI: `10.1109/JSAC.2011.111002`.

[74] R. M. Karp, "Reducibility among combinatorial problems", in *Complexity of computer computations*, Springer, 1972, pp. 85–103.

[75] IBM, "IBM ILOG CPLEX 12.9 User's Manual", 2019. [Online]. Available: `www.cplex.com`.

[76] M. Arcari, I. Söllner, A. Javadi, S. L. Hansen, S. Mahmoodian, J. Liu, H. Thyrrestrup, E. H. Lee, J. D. Song, S. Stobbe, *et al.*, "Near-unity coupling efficiency of a quantum emitter to a photonic crystal waveguide", *Physical review letters*, vol. 113, no. 9, p. 093 603, 2014.

[77] D. Bertsimas, D. B. Brown, and C. Caramanis, "Theory and applications of robust optimization", *SIAM review*, vol. 53, no. 3, pp. 464–501, 2011.

[78] Harvard University, "A summary of error propagation", 2007. [Online]. Available: `http://ipl.physics.harvard.edu/wp-uploads/2013/03/PS3_Error_Propagation_sp13.pdf`.