# Deep reinforcement learning for key distribution based on quantum repeaters

Simon D. Reiß* and Peter van Loock†
*Johannes-Gutenberg University of Mainz, Institute of Physics, Staudingerweg 7, 55128 Mainz, Germany*

This work examines secret key rates of key distribution based on quantum repeaters in a broad parameter space of the communication distance and coherence time of the quantum memories. As the first step in this task, a Markov decision process modeling the distribution of entangled quantum states via quantum repeaters is developed. Based on this model, a simulation is implemented, which is employed to determine secret key rates under naively controlled, limited memory storage times for a wide range of parameters. The complexity of the quantum state evolution in a multiple-segment quantum repeater chain motivates the use of deep reinforcement learning to search for optimal solutions for the memory storage time limits - the so-called memory cut-offs. The novel contribution in this work is to explore very general cut-off strategies which dynamically adapt to the state of the quantum repeater. An implementation of this approach is presented, with a particular focus on four-segment quantum repeaters, achieving proof of concept of its validity by finding exemplary solutions that outperform the naive strategies.

arXiv:2207.09930v1 [quant-ph] 20 Jul 2022

## I. INTRODUCTION

In the last century, the discovery and progress of quantum physics fundamentally reshaped the understanding of the world. In the 21st century, the scientific community started to utilize quantum mechanics to develop new technologies emerging as an entire field of quantum technology. In this field, quantum cryptography is of high relevance and the first technology to be on its way to broader commercialization. Common classical cryptography relies upon computational hardness assumptions to ensure the security of the transmitted information. This is an inherently vulnerable concept, since advances in the computational power can break the condition on which the security was built. Quantum cryptography, however, offers unconditional security based on fundamental laws of physics [1, 2]. In recent years, substantial research has been put into realizing practical quantum cryptography. While fiber-based quantum key distribution (QKD) over a couple of hundreds of kilometers has already been successfully implemented and employed [3, 4], achieving high transmission rates over longer distances, up to 1000 km and beyond, remains an ambition for the foreseeable future. Complementary, satellite-based quantum communication transmitting quantum information through free space is another promising approach [5, 6]. Ultimately, global quantum networks are expected to be based upon a combination of fiber and satellite links. Another interesting option is the concept of twin-field QKD reaching distances beyond 400 km with standard optical fibers [7] and up to about 800 km by using low-loss fibers (0.1419 dB/km) as was recently demonstrated in Ref. [8].

The fundamental building block to extend the range of quantum networks to larger distances is a quantum repeater [9, 10]. Similar to classical communication, the optical fibers used in quantum communication suffer from channel losses that exponentially grow with distance. Quantum repeaters are designed to overcome these losses and preserve transmission rates at long distances by dissecting the communication channel, distributing entangled quantum states over sufficiently short segments, and eventually connecting the elementary links via quantum teleportation (entanglement swapping) [10]. The ongoing development of quantum repeaters has seen significant progress in recent years in the theoretical concepts [9, 11–14], proposals for implementations [15–17], as well as the engineering of the necessary hardware [18–21].

Besides the experimental implementations, suitable strategies operating the quantum key distribution over quantum repeaters have to be developed. This turns out to be a very challenging task, since the complexity of analyzing multiple-segment quantum repeaters grows quickly with the number of repeater stations and hence the distance [14, 22].

Memory-based quantum repeaters store intermediate states in quantum memories and are currently the most experimentally feasible approach that is scalable to large distances by concatenating sufficiently many quantum repeater nodes [18, 23]. Hence, one task of the operating protocol is to manage and control the quantum states stored in the quantum memories. Physical implementations of today's quantum memories suffer from degradation of the stored quantum states [18, 24, 25].

There are at least two established, similar approaches to counteract the memory degradation and improve the fidelity of the distributed states at the cost of lower rates. One approach is based upon a so-called cut-off where quantum states are discarded when their storage times exceed a chosen threshold [12, 24–31]. An approach to simplify the computation is to use a memory buffer instead (sometimes referred to as "memory access time") where the generation of initial entangled states between adjacent repeater stations is restarted at fixed times [32, 33]. A memory buffer is distinct from a memory cut-off, since two neighboring segments, even when ready, must wait until a predetermined time resulting in an unnecessary dephasing of the states. In the case of a

* sreiss@uni-mainz.de
† loock@uni-mainz.de

cut-off, states that have waited for any duration below the cut-off are swapped as soon as possible.

The motivation of this work is to optimize the quantum repeater strategies in quantum key distribution tasks. Exact analytical or even numerical optimization approaches often seem computationally infeasible in the treatment of multiple-segment quantum repeaters [27, 29, 32–35]. In the following, we will use the typical terminology of machine learning where strategies are termed policies.

In Ref. [31] a numerical optimization of the cut-off to maximize the secret key rate is presented. In that work the cut-off is optimized per nesting level. The results presented in Ref. [31] use a fixed ("doubling") swapping scheme which differs from the present work where quantum states are "swapped as soon as possible". More importantly, the policies presented in our work are more versatile in their ability to dynamically adapt to the state of the repeater chain and hence are significantly more complex to analyze.

On the way towards optimizing large-scale quantum networks incorporating a vast parameter space, methods able to handle this level of complexity remain of particular interest. Reinforcement learning (RL) is a method capable of finding near optimal solutions to problems where an analytical treatment is infeasible. Deep reinforcement learning (DRL) extends RL by the use of artificial neural networks. In recent years, DRL has made significant advances in optimizing control tasks for problems which where previously unsolvable [36, 37]. Notable examples include training a computer to play Atari games from raw game pixels [38] and performing locomotion tasks [39, 40]. This motivates the application of these methods to quantum communication networks, offering solutions even for scenarios where the complexity exceeds what is achievable with other numerical optimization approaches.

Most recently, RL was successfully applied to some quantum information tasks. For example, in Refs. [41, 42] agents autonomously developed well-known quantum information protocols and completed quantum error correction strategies, respectively. In Ref. [43] RL was used to optimize quantum error correction codes.

The above-mentioned references as well as the present work make use of classical algorithms to solve quantum problems. This should not be confused with quantum machine learning approaches where the optimization itself has quantum aspects to it.

In the present work, DRL will be applied to QKD via quantum repeaters [44]. As the first step, we formulate a memory-based multi-segment quantum repeater as a Markov decision process (MDP). This MDP incorporates the full description of the quantum states and incorporates channel loss and Pauli errors as well as the option to discard any intermediate quantum state. Based on this, a simulation is employed, including a simple uniform memory cut-off. We present a broad range of results on the dependency of the secret key rate of the experimental parameters for the memories, the segment lengths, and the uniform cut-off parameter for the special class of four-segment quantum repeaters. The simulation also serves as the necessary groundwork for our DRL approach.

We adapt a public implementation of a proximal policy optimization (PPO) DRL algorithm to the simulation in order to find sophisticated memory policies optimizing the secret key rate. The major obstacle in this application of DRL is that the optimization merit is non-additive in terms of the fidelity of the distributed quantum states. We offer an elegant solution in proposing a generalized objective function, expanding common RL theory while maintaining the applicability of convergence improving techniques related to value functions. This improves computational feasibility compared to a simple solution via an episodic reward. The search space consists of the full memory control over discarding individual quantum states, based on the entire information available at any moment.

In this dynamic adapting of the policy lies the novel contribution of our work, compared to prior work which considered static fixed cut-offs that were assigned to nesting levels of a doubling scheme or a single point-to-point link. Furthermore, DRL is the enabling method to achieve this versatility. It has proven to excel in optimization tasks whose complexity exceeds the capabilities of other numerical approaches [36, 37], thus offering an approach for large-scale networks including numerous interleaving processes and errors. This proves to be an already non-trivial task for the four-segment quantum repeaters as considered here. In principle, the state space of our MDP modeling the quantum repeater is infinite. By setting a maximum accumulated storage time $t_m$, which will be further defined in Sec. II, one could limit the size of the state space to $t_m^9$ for a four-segment quantum repeater. Thus, even when assuming $t_m$ as low as $t_m = 10$ the number of states is $10^9$. Taking into account that there are at least two possible actions for any relevant state of the MDP and at least two possible transitions for each action, this lower estimation illustrates that analytical approaches are infeasible to solve this optimization problem. The present work can be understood as a proof of concept of a DRL-based optimization method for four-segment repeaters that is equally applicable to larger repeater chains where we expect it to be even more powerful compared with the standard approaches.

Ultimately, we find policies for the quantum memory treatment, which outperform the naive approaches used in the simulations. Therefore we demonstrate a successful proof-of-concept application of a DRL approach in a first step towards solving complex optimization tasks in quantum networks. The paper is organized as follows. In Sec. II, we will present our abstract model of a multi-segment quantum repeater chain. Then, in Sec. III, we will describe and discuss the results of simulating four-segment quantum repeaters using this model. In Sec. IV, we present our DRL approach and its results, especially in comparison with the simulations without DRL. Finally in Sec. V, we conclude with a summary and con-

clusion, followed by several appendices giving additional technical details.

## II. QUANTUM REPEATER MODEL

This section presents our model of a multi-segment quantum repeater including the relevant errors. We will use an MDP to describe the evolution of quantum states in a quantum repeater. We briefly introduce the secret key rate as a figure of merit to evaluate the performance of a quantum repeater. Lastly, we will discuss a memory cut-off policy, in which quantum states in the memories of a quantum repeater are discarded in order to improve the fidelity of the states distributed between the communicating parties.

### A. Physical model and parameters

Let us now introduce a simplified model of a multi-segment quantum repeater. This offers the possibility to obtain fairly general and conceptual results independent of any specific implementation.

A simple, generic multi-segment quantum repeater chain is depicted in Fig. 1. Quantum repeaters are used to distribute entanglement between two distant parties by segmenting their connecting channel. Initial entanglement is generated in each segment, for instance, employing a source of entangled photon pairs at a node (or placed in the middle between two memory nodes) and sending one photon to each adjacent node. At each repeater node an entanglement swapping operation, which is essentially a Bell measurement for quantum teleportation, is performed on the memory qubits transferring the entanglement step by step over the entire distance to the communicating parties.

Throughout this work, quantum repeaters with the following properties are considered:

- The quantum repeater may in principle consist of an arbitrary number of segments. For simplicity, only one-dimensional concatenated segments are treated here, i.e. so-called quantum repeater chains (an extension to multidimensional repeater networks would also be possible with our methods, but this will not be demonstrated in this work). Eventually, we will put a particular emphasis on four-segment quantum repeaters.

- Each repeater node contains one quantum memory for each adjacent node.

- The repeaters are "clocked repeaters" where the clock times are determined by the classical communication times between nodes.

- All errors of the implementation can be described as Pauli channels,

$$\mathcal{N}(\rho) = \sum_{j=0}^{3} a_j P_j \rho P_j^\dagger, \tag{1}$$

where $\rho$ is the density operator of a quantum state, $P_j$ are the Pauli operators, $\{P_j\}_{j\in\{0,1,2,3\}} = \{\mathbb{1}, X, Z, XZ\}$, and $a_j$ are real-valued, non-negative probabilities satisfying the normalization property $\sum_{j=0}^{3} a_j = 1$.

- Elements for active entanglement purification [45, 46] and more general quantum error correction are not included throughout. The only mechanism to suppress errors on the memories is a finite memory cut-off.

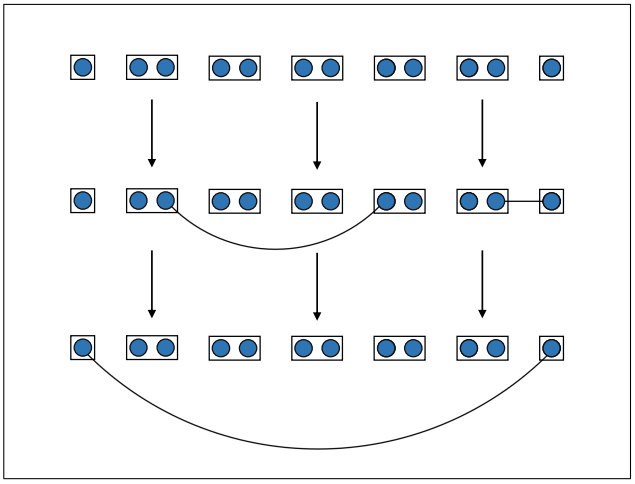

FIG. 1: A multiple-segment quantum repeater chain. Each box depicts a quantum repeater node. Entangled two-qubit states are initially distributed within the segments. As soon as two-neighboring segments are ready, i.e. their memory pairs share a successfully distributed state, the entanglement swapping is performed. After subsequent swappings, synchronized via storage at the memory nodes, eventually an entangled state is shared between the two most outer nodes.

The parameters characterizing the implementation of a quantum repeater in our model are:

- $n$: The number of segments.

- $L_0$: The length of one segment, which is the distance between adjacent repeater nodes.

- $\nu_i \in [0, 1]$: The probability of any error to occur in a given situation, in particular, the probability of a phase flip to occur on a quantum state stored in a memory.

- $c$: The signal speed for classical communication between repeater nodes (typically, the signal speed in an optical fiber, which can be employed for transmitting both quantum and classical signals).

- $L_{\text{att}}$: The attenuation length of the optical fibers for realizing the quantum channels between the repeater nodes.

The following, additional parameters can be obtained from those listed above:

- $\tau_0 = \frac{L_0}{c}$: One round of quantum and classical communication, which is the time it takes for a node to send qubits and classical information to an adjacent node.

- $\tau_c$: The memory coherence time for a bipartite quantum state, see Eq. (8). It corresponds to half of the commonly used parameter of the single-qubit dephasing channel for the quantum memories.

- $\eta = e^{-\frac{L_0}{L_{\text{att}}}}$: The transmissivity of the optical fibers connecting the repeater nodes, inducing an exponential photon loss with distance.

- $p = p_x\eta$: The probability of generating an initial two-qubit entangled state in a segment in one time step. The parameter $p_x$ incorporates the photon creation efficiency of the spin-photon or photon pair source, fiber channel in- and out-coupling and detector efficiencies, as well as memory write-in efficiencies.

  The attenuation length $L_{\text{att}}$ is a physical parameter of the optical fiber dependent on the employed wavelength of the transmitted photons. The preferred wavelength is that of telecommunication (1.55 $\mu$m, potentially requiring wavelength conversions) and a typical value for $L_{\text{att}}$ is 22 km.

  Note that the particular form of the transmissivity $\eta$ relies on the assumption of constant photon losses, which is the case for optical fibers and this would have to be changed for different channel implementations (such as fluctuating loss in a free-space channel).

### B. Markov decision process

For the purpose of adequately describing a quantum repeater, we designed an MDP to model the propagation and evolution of the quantum states stored in the quantum memories of a multi-segment quantum repeater in our physical model.

An MDP is defined as a tuple $(S, A, P, R)$ where

- $S$ is a set of states of the environment,

- $A$ is a set of actions performed by the agent on the environment,

- $P : S \times A \times S \to [0, 1]$ is the transition probability in a time step between states dependent on the applied action in this time step,

- $R : S \times A \times S \to [0, 1]$ is the immediate reward received from a transition.

In practice, besides photon losses, Pauli channels include the most common errors in quantum repeaters and networks [10]. This is fortunate from the mathematical perspective, since Pauli channels commute with the entanglement swapping operations [47, 48]. Thus, in the case of memory dephasing, which is a Pauli channel, the number of accumulated error operations is proportional to the accumulated storage time, i.e. additively propagating the storage times of quantum states through the swapping operations gives the correct accumulated dephasing. Therefore, instead of rigorously calculating the density matrix of every intermediate quantum state, it is sufficient to treat the undisturbed quantum states that are only subject to channel loss and count the accumulation of errors separately. Then all these accumulated errors can be applied to the final quantum state distributed between the communicating parties. Hence, the states of the MDP can be encoded in a triangular matrix where each entry corresponds to a pair of repeater nodes storing the accumulated errors of the bipartite quantum state. One time step of the MDP is the time it takes to send quantum and classical information between adjacent repeater nodes (corresponding to $\tau_0$ in our physical model). Note that in our treatment here we only consider the simplest, still to some extent idealized scenario with only two effects determining the final rates: channel loss and memory dephasing. This channel-loss-and-memory-dephasing-only model is similar to that of Ref. [18], however, here extended from two to four repeater segments. Additional experimental (error) parameters have been included in the analytical treatment of Ref. [14]. As opposed to the present work, Ref. [14] has no focus on optimizing the memory cut-off (for a discussion on other existing works and approaches for optimizing the memory cut-off, see Sec. II D 1).

The set of possible actions of the agent consists of arbitrary combinations of swapping and discarding operations. The action applied in one time step may perform entanglement swapping on any subset of repeater stations and discard any subset of the stored quantum states. In the event that the two outer nodes of the repeater share an entangled state, this state is discarded and its fidelity is returned as a reward. A detailed description of the MDP can be found in App. A.

### C. Quantum key distribution - secret key rate

The secret key rate $R$ is a suitable measure to evaluate the performance of a quantum repeater, as it combines the relevant properties, namely quantum state fidelity and raw transmission rate, into one convenient figure of merit. Furthermore, long-range QKD is one of the main applications motivating the development of quantum repeaters. The secret key rate is defined as the amount of

secret bits distributed between two communicating parties, commonly denoted Alice and Bob, in bits per time and can be written as

$$R = Y \cdot r, \tag{2}$$

where $Y$ is the raw rate and $r$ the secret key fraction. The raw rate is the amount of raw (quantum) bits distributed between Alice and Bob per time and the secret key fraction is the potential amount of secret key that can be extracted per raw (quantum) bit via post-processing.

Throughout this work the Lo-Chau BB84 [2, 49] protocol is used to determine the secret key rates. The asymptotic secret key rate of the Lo-Chau BB84 protocol, see for example Ref. [50], reads as

$$R_{\mathrm{BB84}} = Y \cdot r_{\mathrm{BB84}}\left(e_1, e_2\right), \tag{3}$$

with the secret key fraction

$$r_{\mathrm{BB84}}\left(e_1, e_2\right) = 1 - h\left(e_1\right) - h\left(e_2\right), \tag{4}$$

where $h$ is the binary entropy function defined on the interval $[0, 1]$ as

$$h\left(p\right) = -p \log p - \left(1-p\right) \log\left(1-p\right), \tag{5}$$

where the logarithm is taken to the base two.

The quantum bit error rates $e_1$ and $e_2$ are the probabilities of non-coinciding measurements performed by the two parties on the shared data qubits in the respective coinciding basis. Commonly, the Pauli $X$ and Pauli $Z$ bases are chosen. In a realistic scenario, Alice and Bob have to estimate the bit error rates on a finite set of test data. In the proposal of Ref. [49] it was shown that the number of test bits can be chosen in the order of $\Omega(\log k)$, where $k$ is the length of the final key, while still achieving unconditional security. Thus, the fraction of test data can be chosen asymptotically close to zero.

The concrete derivation of the asymptotic secret key rate for our physical model, including only fiber losses and memory dephasing (see later Eq. (8)), can be found in App. B and it takes the simple form

$$R_{\mathrm{BB84}} = Y \cdot \left(1 - h\left(\frac{1}{2}\left(1 - \mathbb{E}\left[e^{-\frac{t}{\tau_c}}\right]\right)\right)\right), \tag{6}$$

where $\tau_c$ is the coherence time of the quantum memory for a bipartite quantum state and $t$ is the storage time of the quantum state.

### D. Memory dephasing - cut-off and swapping strategies

#### 1. *Overview of existing approaches*

One of the most significant error sources in practical implementations of memory-based quantum repeaters is the degradation of quantum states in the memories. The distribution of entanglement via a quantum repeater typically is a highly probabilistic process. Thus, the level of degradation caused by memory dephasing can become very large, with some nodes having to wait for others for so long that entanglement can hardly be preserved.

As already mentioned in the introduction, limiting the time during which the quantum states are stored is a common strategy to improve the fidelity of the distributed quantum states. This concept adapted to quantum repeaters based on imperfect memories was first introduced in Ref. [26]. In terms of resources, this is the least expensive approach to suppress memory errors compared to entanglement purification strategies relying on the distribution of additional entangled-state copies and extra classical communication rounds or, alternatively, strategies based on the implementation of more complicated quantum error correction codes. Again, the processes in a quantum repeater are highly probabilistic and computing the corresponding probability distributions for complex repeater schemes quickly becomes infeasible. In recent years, significant research has been devoted to this and results under various simplifying assumptions have been reported. In order to put our work and results in context we shall give a brief overview of some of the existing literature.

The distributed entangled quantum states in a two-segment repeater including cut-offs were thoroughly analyzed in the presence of various error models in Refs. [18, 24, 25, 47]. Reference [47] also considers more than two segments, however, adapted to specific distribution protocols. In Ref. [28] a closed, compact formula is derived to efficiently compute the raw rate of a quantum repeater with an arbitrary number of segments, including a memory cut-off, with the constraint of deterministic entanglement swapping operations and the (practically suboptimal) assumption that all swappings are performed at the end. In Ref. [27] an algorithm based on Markov chains and solving linear-equation systems is presented which exactly computes the average waiting time of quantum repeater chains with an, in principle, arbitrary number of segments including various swapping strategies. For up to four segments exact rate formulas are given. However, this algorithm is practically limited as its runtime in $\mathcal{O}(c^n)$ growth rather quickly with the number of segments $n$ and the cut-off $c$. In [22] swapping strategies in the presence of non-deterministic swapping operations are optimized for the best possible raw rate, but no cut-off was included.

In a more numerical approach, in Ref. [35] an efficient optimization over entanglement distribution schemes using dynamical programming is proposed under the assumption that there is no time-dependent decoherence in the quantum memories. The algorithm recursively solves larger repeater chains by dissecting them into smaller sub-chains which are optimized requiring the idealizing simplification that the sub-processes finish at the average time. In Ref. [32], the optimal memory buffer maxi-

mizing the rate of distillable entanglement of the average state at all nesting levels in a doubling repeater scheme is computed. In another work, in Ref. [29] algorithms to compute the probability distribution of the fidelity and waiting time for the first distributed entangled state in a quantum repeater and a numerical optimization over a cut-off is presented. Moreover, in Ref. [30], it was shown that the optimal policy maximizing the accumulated fidelity in a multiplexed repeater segment can, in the finite-horizon setting, be computed via a dynamical programming algorithm. In Ref. [33], a heuristic algorithm was developed which optimizes quantum repeater schemes in order to minimize the generation time of an entangled state between communicating parties for a fixed minimum success probability and fidelity. Their schemes also include memory buffers and entanglement distillation. To overcome the computational complexity the schemes were restricted to those that succeed at all levels near-deterministically. This is enforced by repeating all probabilistic processes a sufficient number of times to ensure a high probability of at least a single success. In Ref. [51], a genetic algorithm was presented and applied to NetSquid simulations [13] providing insights into the necessary hardware parameters for viable quantum repeaters. Recently, in Ref. [14], an exact rate analysis for quantum repeaters including experimental errors was presented. However, the cut-off was included primarily in a sequential repeater scheme and no optimization of the cut-off was made. Reference [31], as most relevant to the present work, was already discussed in the introduction.

As was stated in the introduction, our policies can dynamically adapt to the state of the repeater and decide for any individual quantum state if it should be discarded or not. Thus, we include an extended toolbox which offers the possibility of more sophisticated and, as we will present in Sec. IV E, better policies than previous approaches. In order to handle the complexity introduced by this generalization we use a DRL algorithm which is also novel in the optimization of quantum repeaters. We will distinguish the two policies by referring to the simpler policy as the cut-off policy and to the more sophistcated, better policy as learned policies.

#### 2. *Cut-off model and policies in this work*

In this paper, the cut-off policy is defined such that any quantum state whose accumulated storage time exceeds a chosen cut-off value $c$ will be discarded. Accumulated storage time here refers to the storage time that propagates additively through the swapping operations leading to the final quantum state shared between the most distant stations. If, for example, a state is stored for $n_1$ time steps and swapped with a state stored for $n_2$ time steps, the state after the swapping has an accumulated storage time of $n = n_1 + n_2$ time steps. Note that this differs from other common definitions of the cut-off, which usually only use the time the state is stored for in the current nesting level. However, as was explained in Sec. II B, our accumulated storage time does, in fact, correctly describe the propagated quantum state. This substantiates the choice of an accumulated cut-off, since it determines the quality of a quantum state more accurately than when only the storage time in one nesting level is taken into account. Also, note that a cut-off per nesting level was employed in Ref. [31] where it was found that for the presented example parameters a non-uniform cut-off, i.e. a different cut-off for each nesting level, does not yield a significant improvement of the secret key rate.

The cut-off policy as defined in our work, based on an accumulated storage time, actually ignores the fact that the information on which the decision to discard quantum states is based on might not be directly available at the node performing the operation at the corresponding time. Our basic assumption that the necessary information is indeed available is not obvious and may even seem unreasonable. Also, an argument that this serves as a bound is invalid, as classical communication is not an obstacle one could possibly circumvent. First, one should note that this does not physically contradict anything of how the quantum repeater is operated and functions. The only but clearly idealizing element of this is that decisions for the controlled part of the process, namely the discarding of the quantum states, are generally based on the entire state of the MDP. This means decisions of the policy performing an action in one node might be based on information that cannot be possibly accessible in this particular node. In the context of examining fully realistic quantum repeaters, all necessary classical communication must be taken into account. However, in this work, in order to allow for optimal comparability between unlearned and learned policies we choose these to be all-knowing to give the algorithm at any time complete information to discover new policies. Therefore, imposing the same conditions on the unlearned policies, simulated in the following section, as those imposed on the learned policies, described later, ensures that any advantage of the learned policies must be based on their better strategy and will not be based on any better assumptions.

Next, we want to briefly discuss swapping strategies. As a simplification, throughout this work, we assume the entanglement swapping to be error-free and deterministic. In the case that memory dephasing is the only error acting on the quantum states, the ideal swapping strategy is to perform the Bell measurement at any repeater node immediately when both adjacent segments have successfully distributed entanglement. This follows from the fact that a swapping operation reduces the number of bipartite quantum states that are simultaneously stored and thus the accumulated dephasing (for a rigorous and more mathematical proof of this statement, together with a systematic formalism to exactly calculate the secret key rate in such protocols that "swap as soon as possible", see Ref. [14]).

All schemes simulated in this work are what we call overlapping schemes. This means that, opposed to other

common analysis which often only considers the first distributed state, we consider the generation of many distributed states in an overlapping fashion. This further means that in the lower nesting levels states will be continued to be generated for future processing while other memories are still occupied for the next distributed state. In the four-segment case treated in this work this might not have a significant effect on the performance of the repeater, as the only case where this is actually beneficial is when an entangled state over three segments is shared. In this case, the segment in the middle of this connection can again start entanglement generation attempts.

## III. SIMULATING KEY DISTRIBUTION BASED ON QUANTUM REPEATERS

In this section, the cut-off policy is examined via simulations with respect to its effect on the secret key rates. The primary purpose of this is to provide benchmarks of achievable secret key rates to be surpassed, for the policies later considered with the RL algorithm in Sec. IV. Before presenting in detail the results of the simulation, we shall briefly describe the specific errors included in the simulation, extending the discussion on our physical model and errors from Sec. II A.

### A. Errors and imperfections

A realistic quantum repeater is a complex system with numerous parameters. The more realistic effects are included in a simulation, the more meaningful the results become regarding the assessment of a practical quantum repeater. On the other hand, excluding some sources of errors allows for a more focused view of the most important ones. Another way to look at this is that the results considering fewer errors serve as upper bounds for what is achievable.

The focus of this work is on strategies to counteract the dephasing of quantum memories. Therefore the only two imperfections that we choose to include are the dephasing of the quantum memories and the finite transmissivity of the quantum channel (while without the latter quantum repeaters would be pointless).

We model the degradation of the quantum states that are stored in the memories as a dephasing channel,

$$\mathcal{N}_Z(\rho) = (1-\nu)\,\rho + \nu Z\rho Z, \tag{7}$$

which we further specify as exponential decay at time $t$,

$$\mathcal{N}_Z(\rho, t) = \frac{1}{2}\left(1 + e^{-\frac{t}{2\tau_c}}\right)\rho + \frac{1}{2}\left(1 - e^{-\frac{t}{2\tau_c}}\right) Z\rho Z, \tag{8}$$

where $Z$ is a Pauli operator and $\tau_c$ is the coherence time of the quantum memory [24]. Note that the dephasing channel as defined in Eq. (8) is a single-qubit channel. In a quantum repeater, both qubits of a stored entangled bipartite state are subject to this error channel (individually, i.e. locally and independently) prior to entanglement swapping. This is also the reason for the factor 2 multiplied to the coherence time $\tau_c$, which was defined as the coherence time for a bipartite quantum state. The second imperfection included is the photon loss in the optical fibers, which was already described in Sec. II A.

### B. Results

This section presents the results of our numerical simulations. In the simulations, the secret key rates for a BB84 protocol are computed. A discussion of the required formulas can be found in App. B. Two-segment quantum repeaters have already been analyzed rather thoroughly [24, 25] and, being structurally simple, do not offer extra insights for our treatment. In this work, we choose to consider the class of four-segment repeaters. This is partially motivated by the existing literature, which often considers doubling the number of repeater segments, making four segments the logical next step beyond two. Here it is also worth noting that the more segments are included, the less feasible it can become in our model without entanglement purification and error correction that meaningful practical or even non-zero secret key rates can be obtained [14]. For repeaters of a larger scale such extra tools will have to be included.

For the attenuation length $L_{\text{att}}$ we choose 22 km which is today's experimental standard for the telecom wavelength of 1550nm [52, 53]. The speed of classical communication is assumed to be $2{\cdot}10^8$m/s, which corresponds to the typical value in optical fibers. The segment lengths are chosen within a commonly considered range of 20 - 70 km, and the coherence time of the memories is assumed to be within the scope of today's experimental possibilities [18].

#### 1. Uncertainties

In the simulations, since overall minimal uncertainties were achieved, 3-$\sigma$ confidence intervals ($\approx$ 99,73%) are displayed in all plots with uncertainty bars, instead of the more commonly chosen $\sigma$ intervals ($\approx$ 68,27%). However, note that the uncertainty bars are still small enough for many data points not to be easily visible, as they can be smaller than the displayed points of the data. One should not be misled by their width in the horizontal direction, since this axis corresponds to a discrete quantity and does not display an uncertainty. The style of displaying the uncertainties was chosen to increase their visibility. Their height corresponds to the 3-$\sigma$ confidence interval of the secret key rate.

It is also worth noting that the uncertainties get larger for a sparser entanglement distribution in time, thus for longer segment length and smaller cut-off parameters. The reason is that the sparser the received quantum

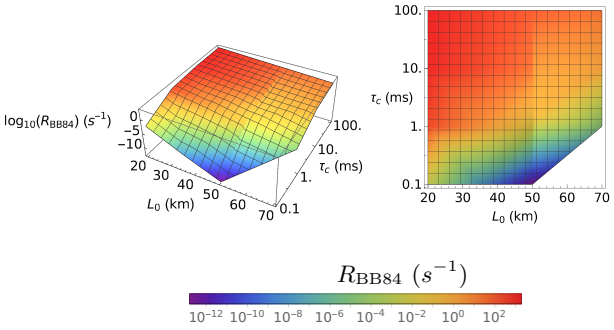


FIG. 2: Secret key rates of four-segment quantum repeaters without memory cut-off. The secret key rate is plotted in secret bits per second dependent on the segment length $L_0$ and the memory coherence time for a bipartite quantum state $\tau_c$. The secret key rate and the coherence time of the memories are displayed on a logarithmic scale.

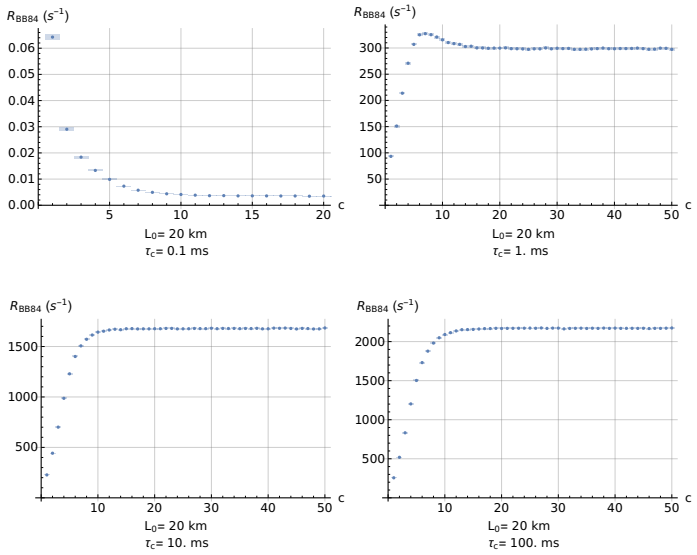


FIG. 3: Secret key rates of four-segment quantum repeaters in secret bits per second dependent on the memory cut-off parameter $c$ for segment length $L_0 = 20$ km and different coherence times of the memories for a bipartite quantum state $\tau_c$. Plotted with 3-$\sigma$ confidence intervals.

states are, the larger the required sample size is to obtain an equally good estimation of the average fidelity. Hence, the precision of some values with more significant uncertainties is limited by the same computational time needed to simulate larger sample sizes.

### 2. *Quantum repeater without cut-off*

In this subsection, the secret key rate of quantum repeaters that are not controlled by a policy including memory cut-off is examined. The results are presented in Fig. 2.

In order to deepen our understanding, we now further elaborate some implications of the parameters. One should note that since the classical communication time $\tau_0 = \frac{L_0}{c}$ scales linearly with the segment length $L_0$, a larger segment length $L_0$ does not only decrease the transmissivity of the channel, but it also decreases the effective coherence time per channel use. Furthermore, an increase in classical communication time results in lower raw rates per second for identical raw rates per channel use, thus reducing the efficiency of a channel use. However, these drawbacks are reasonable taking into account that the overall communication distance increases with the segment length, thus fundamentally lowering the transmission rate.

The results obtained in our simulations are what one would intuitively expect, as the secret key rate increases with smaller repeater segments and better coherence time of the memories. The secret key rate saturates for sufficiently large coherence times, because the fidelity of the distributed quantum states approaches unity, and it drops to zero for short coherence times when the fidelity approaches a value of half. With longer segment length the maximum secret key rate decreases. The larger the segment length, the higher are the requirements on the quantum memories to achieve reasonable secret key rates. The increase of the secret key rates with the quality of the memories is steeper for worse memories and gets less significant, as the secret key rate saturates towards perfect memories. The decrease of the secret key rate with increasing segment distance is significantly steeper for worse memories.

### 3. *Quantum repeater with cut-off*

In this subsection, the secret key rates of quantum repeaters that are controlled by the cut-off policy are examined. In Figs. 3 - 7 the relation of the secret key rate and the cut-off parameter is displayed for various parameter choices of the segment length $L_0$ and the coherence time of the memories $\tau_c$.

In the limit of large cut-off parameters, the policy is identical to the case without cut-off. This is apparent in all plots, as the secret key rate converges to the no cut-off rates for large cut-off parameters.

For all segment lengths, a similar behavior with respect to the coherence time of the memories is visible. In the regime of low coherence times (i.e., $\tau_c = 0.1$ ms), the secret key rates drop to values near zero, with the ideal cut-off parameter being the lowest possible value, which is one. This observation coincides with the results from Ref. [32]. With increasing coherence time, the optimal cut-off parameter shifts towards larger values, and its peak of the secret key rate for the optimal cut-off decreases relative to the secret key rate of the asymptotically converging quantum repeater without cut-off. This

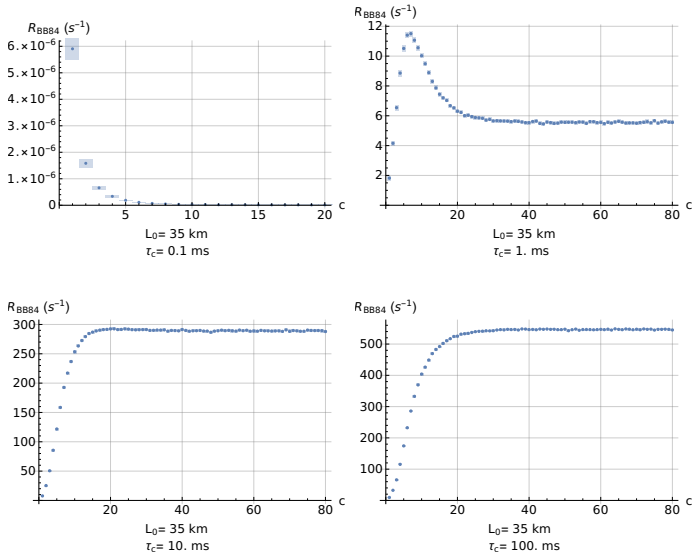


FIG. 4: Secret key rates of four-segment quantum repeaters in secret bits per second dependent on the memory cut-off parameter $c$ for segment length $L_0 = 35$ km and different coherence times of the memories for a bipartite quantum state $\tau_c$. Plotted with 3-$\sigma$ confidence intervals.

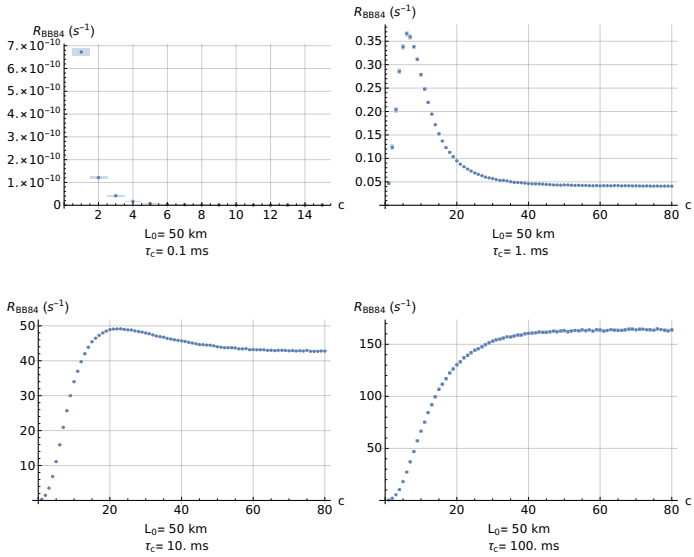


FIG. 5: Secret key rates of four-segment quantum repeaters in secret bits per second dependent on the memory cut-off parameter $c$ for segment length $L_0 = 50$ km and different coherence times of the memories for a bipartite quantum state $\tau_c$. Plotted with 3-$\sigma$ confidence intervals.

decrease causes the peak to vanish for larger coherence times. At this point, the coherence time reaches a quality value beyond which the cut-off policy does not offer any improvements on the secret key rate.

This behavior moves towards longer coherence times as the segment distance increases. In other words, the shapes of the simulated and plotted secret key rates are qualitatively identical, interpreting the coherence time in relation to the segment length. This indicates that, conceptually, the results when conditioned on the relation between the parameters should be qualitatively applicable to any parameter regime.

In Fig. 8, the ratio between the secret key rate of the best cut-off policy and that without cut-off is shown. This further illustrates the above observations. The advantage of the cut-off vanishes for better quantum memories and shorter segment lengths and it increases for worse quantum memories and larger segment lengths. These observations coincide with the results from Ref. [31].

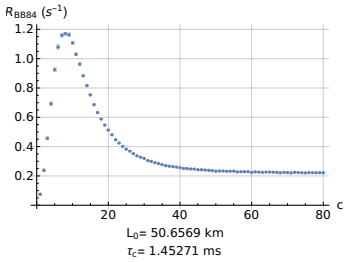


FIG. 6: Secret key rates of four-segment quantum repeaters in secret bits per second dependent on the memory cut-off parameter $c$ for segment length $L_0 = 50.6569$ km and coherence times for a bipartite quantum state $\tau_c = 1.45271$ ms. Plotted with 3-$\sigma$ confidence intervals. [a]

[a] These values correspond to the probability to generate initial entanglement in one segment $p = 0.1$ and the dephasing probability in one time step $\nu = 0.08$.

### C. Simulation - conclusion

In this section, the secret key rate of the BB84 protocol has been examined via numerical simulations for various parameters of a four-segment repeater setup. The two included imperfections were the dephasing of the quantum memories and the signal attenuation in the optical fibers.

The main observation that can be made with our results is that the worse the coherence time of the quantum memories is in relation to the segment length, the more significant is the advantage of the cut-off. This relation is steeper in the regime of bad memories in relation to the segment length and flattens for better memories until it converges to the point where the cut-off offers no improvement.

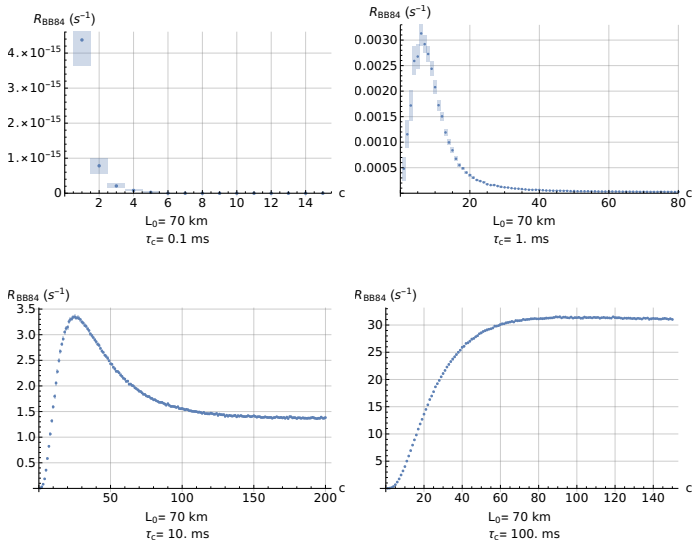


FIG. 7: Secret key rates of four-segment quantum repeaters in secret bits per second dependent on the memory cut-off parameter $c$ for segment length $L_0 = 70$ km and different coherence times of the memories for a bipartite quantum state $\tau_c$. Plotted with 3-$\sigma$ confidence intervals.

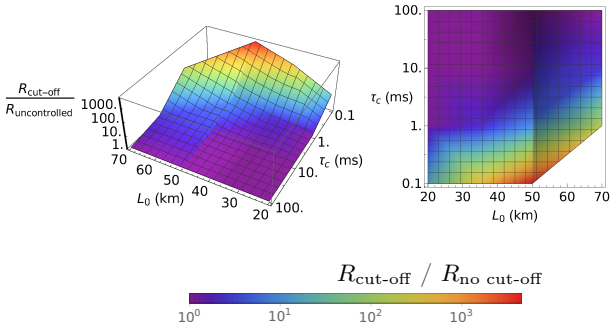


FIG. 8: Simulated ratio between the secret key rates of the scheme with the best memory cut-off policy and schemes entirely without memory cut-off. This ratio and the coherence time of the memories for a bipartite quantum state $\tau_c$ are plotted logarithmically.

## IV. DEEP REINFORCEMENT LEARNING APPLIED TO QUANTUM REPEATERS

In this section, we present our DRL approach to optimize secret key distribution employing quantum repeaters. The aim is a proof of concept, showing that DRL offers the possibility to provide sophisticated policies for memory treatment which outperform the more naive approaches.

First we will introduce our algorithm and its implementation. For the non-expert reader this includes a concise self-contained introduction to some DRL theory. After a brief discussion about the experimentation process the successful learning runs are presented. For conclusion, some observations about the learned policies are discussed.

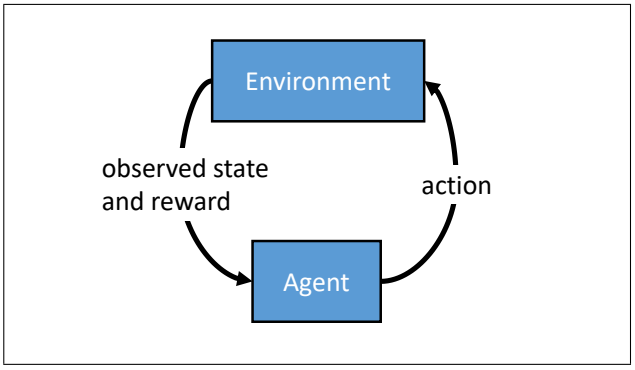


FIG. 9: Conceptual framework of an RL algorithm.

### A. DRL algorithm

In this subsection, we will discuss the algorithm that we have employed. RL is a class of machine learning algorithms which trains a so-called agent to optimize its behavior in an environment. This concept is illustrated in Fig. 9. Deep learning uses artificial neural networks to learn hierarchical representations in order to solve classification problems. In DRL these methods are combined by using an artificial neural network to model the agent [54]. For our DRL application the MDP described in Sec. II B serves as the model of the environment.

The rest of this subsection introduces some broadly known DRL theory with some intuitive explanations as well as the specific algorithm used in this work. For the interested reader Richard S. Sutton's and Andrew G. Barto's book [55] can serve as a comprehensive introduction to RL. A shorter, more practical introduction can be found on the website of Open Ai Spinning Up [56].

We denote the policy of the agent as the probabilistic policy parameterized by a vector $\theta$,

$$\pi_\theta(s, a) : S \times A \to [0, 1]. \tag{9}$$

The policy gives the probability of the agent taking the action $a \in A$ given the environment is in the state $s \in S$. In DRL the agent (i.e., the policy) is encoded in a neural network. In this case, the vector $\theta$ consists of the parameters (so-called weights) of the neural network.

Furthermore, we introduce the common definition of a trajectory as a sequence of states and actions up to a time $T$ called the (time) horizon,

$$\tau = (s_0, a_0, s_1, a_1 ...., s_T). \tag{10}$$

The central merit of the optimization is the accumulated discounted reward of a trajectory $\tau$ up to a time $T$,

$$R_0 = \left[ \sum_{t=0}^{T} \gamma^t r_t \right], \tag{11}$$

where $\gamma$ is the so-called discount parameter and $r_t$ the immediate reward at time $t$.

The usual objective of the agent is to optimize the policy $\pi_\theta$ to maximize the expected accumulated discounted reward over trajectories $\tau$,

$$\underset{\theta}{\text{maximize}} \underset{\tau\sim\pi_\theta}{\mathbb{E}} [R_0], \tag{12}$$

where $\underset{\tau\sim\pi_\theta}{\mathbb{E}}$ expresses that the expectation value is taken over all trajectories $\tau$ following the policy $\pi_\theta$. Note that the trajectory implies all returned immediate rewards obtained in the trajectory. Further, note that the dependency on $\gamma$ and $T$ is left out of the expectation value, because it will be fixed in the optimization. These parameters which are fixed in the run time of the algorithm are termed hyperparameters in order to distinguish them from those parameters that the algorithm optimizes.

Policy gradient algorithms optimize the policy in terms of the objective via a gradient ascent,

$$\theta_{k+1} = \theta_k + \alpha \nabla_\theta \underset{\tau\sim\pi_\theta}{\mathbb{E}} [R_0]\,|_{\theta=\theta_k}, \tag{13}$$

where $\alpha$ is a hyperparameter called the learning rate.

This already summarizes the concept of policy gradient optimization. The policy gradient $\nabla_\theta J(\pi_\theta)|_{\theta=\theta_k}$ is estimated on some training data, usually obtained from a simulation, and applied in an iterative fashion. One iteration is typically referred to as an epoch. As this gradient usually cannot be computed exactly and needs to be estimated, this introduces a finite variance into the stochastic gradient ascent. The fact that this gradient can be estimated is not obvious. The proof that this is in principle possible and the computations to do so are beyond what we want to cover in this section, and so we refer the interested reader to the aforementioned references [55, 56].

Unfortunately, simple DRL algorithms often suffer from slow and unstable convergence properties, i.e. they tend to converge very slowly towards a local optimum and to overshoot in parameter updates causing them to often drastically lose progress in the learning process. In recent years novel, improved DRL algorithms have been proposed to counteract these issues [36, 57, 58]. Our algorithm of choice is a proximal policy optimization (PPO) [58], which is a state of the art approach. PPO-clip is one of the proposed variants of Ref. [58] and is the one we use in this work. We will now explain the concept behind some of the improvements that were made in these algorithms.

A common approach is to generalize the objective of Eq. (12):

$$\underset{\theta}{\text{maximize}} \underset{\tau\sim\pi_\theta}{\mathbb{E}} [\Phi(\tau)], \tag{14}$$

where we call $\Psi_t$ the (generalized) objective function. In general, any objective function $\Psi(\tau)$ satisfying

$$\nabla_\theta \underset{\tau\sim\pi_\theta}{\mathbb{E}} [R_0]\,|_{\theta=\theta_k} = \nabla_\theta \underset{\tau\sim\pi_\theta}{\mathbb{E}} [\Phi(\tau)]\,|_{\theta=\theta_k}. \tag{15}$$

is a valid objective function, since it leaves the gradient in Eq. (12) invariant. Hence, one is free to choose this substitution, while still solving the same optimization problem.

In order to discuss how $\Phi(\tau)$ will be chosen we state the standard definitions of the on-policy value function $V^\pi(s_t)$, the on-policy action-value function $Q^\pi(s_t, a_t)$ and the advantage function $A^\pi(s_t, a_t)$.

We denote the accumulated reward after time $t$ as

$$R_t = \sum_{l=0}^{T-t} \gamma^l r_{t+l}. \tag{16}$$

Note that this is a generalization of Eq. (11), containing $R_0$ as a specific case. This is already an improvement, since due to the truncation of the reward in Eq. (16), the evaluation of an action by the advantage function is independent of the trajectory prior to the action. This makes intuitively sense, because the evaluation of an action should only follow from what happened in consequence of that action.

The on-policy value function $V^\pi(s_t)$ is the expected reward received after time $t$ given that the environment at time $t$ is in the state $s_t$,

$$V^\pi(s_t) = \underset{a_t, s_{t+1}, \ldots\sim\pi_\theta}{\mathbb{E}} [R_t]. \tag{17}$$

The on-policy action-value function $Q^\pi(s, a)$, is the expected reward received after time $t$ given the environment at time $t$ is in the state $s_t$ and the action $a_t$ is taken,

$$Q^\pi(s_t, a_t) = \underset{s_{t+1}, a_{t+1}, \ldots\sim\pi_\theta}{\mathbb{E}} [R_t]. \tag{18}$$

The difference between these functions defines the advantage function $A^\pi(s, a)$, which is a crucial merit in DRL,

$$A^\pi(s_t, a_t) = Q^\pi(s_t, a_t) - V^\pi(s_t). \tag{19}$$

Therefore, the advantage function evaluates how an action performs relative to the current policy. It can be shown that replacing the discounted reward $R_0$ in the objective in Eq. (12) by the advantage function leaves the gradient in Eq. (13) invariant [59] as demanded in Eq. (15), i.e.

$$\nabla_\theta \underset{\tau\sim\pi_\theta}{\mathbb{E}} [R_0]\,|_{\theta=\theta_k} = \nabla_\theta \underset{\tau\sim\pi_\theta}{\mathbb{E}} [A^\pi(s_t, a_t)]\,|_{\theta=\theta_k}. \tag{20}$$

To explain the appeal of this substitution we first note again that in most relevant applications the gradient in Eq. (13) cannot be determined exactly. The solution is to perform a stochastic gradient ascent, i.e. for each parameter update $\theta_k \to \theta_{k+1}$ the gradient is estimated on a finite sample size. The benefit of using the advantage function in the objective is then to decrease the variance of these estimations, which results in better convergence properties.

This can be intuitively understood by interpreting the advantage function as a rescaling of the reward relative

to the expected performance of the current policy. In the case where $Q^\pi(s_t, a_t)$, i.e. the performance merit conditioned on the action $a_t$, is better than the expectation value over all possible actions $V^\pi(s_t)$, the advantage function is positive, otherwise it is negative. It seems reasonable that we can think of the advantage function as a rescaled reward that becomes positive for any better-than-expected actions and negative for any worse-than-expected actions.

This is just one way in which DRL algorithms such as PPO can attempt to improve their convergence properties. More details can be found in the original PPO proposal [58]. We should also note that the advantage function is just one choice for the objective, as any function satisfying Eq. (15) would be a valid candidate. A detailed discussion on objective functions can be found in Ref. [59].

### B. Adaption of the DRL algorithm to non-additive rewards

We will now discuss how our model of QKD via quantum repeaters can be applied to the DRL approach presented in the previous subsection.

A first approach could be to assign the fidelity of an entangled quantum state distributed between the communicating parties to the immediate reward in a time step and return zero reward if no entangled state was distributed in this time step. This approach was, for example, chosen in Ref. [30] to optimize the distributed entanglement in a multiplexed point-to-point quantum communication channel.

In our case, we aim to optimize the secret key rate between the communicating parties. The distributed secret key, however, is not equivalent to the sum of the fidelities of the distributed entangled states. A simple solution consistent with the presented theory of DRL would be to consider a so-called episodic process where the reward of a trajectory is assigned to the last time step. Thus, one could assign the reward as

$$r_t = \begin{cases} 0 & t < T \\ R_{\mathrm{BB84}}(\tau) & t = T, \end{cases} \tag{21}$$

where $R_{\mathrm{BB84}}(\tau)$ is the secret key rate of the distributed quantum states in the trajectory $\tau$. In Eq. (6) it was shown how this can be calculated with the (accumulated) storage time of the quantum states (or, equivalently, their fidelity) of the quantum states which corresponds to the rewards obtained in the trajectory $\tau$. In App. A it is explained how the rewards in a trajectory are determined by the simulated MDP.

The essence of the problem with this approach is often referred to as the credit assignment problem [55, 60]. Intuitively, this form of reward includes little information about the causal connection between a specific action and its influence on the reward. Mathematically this results in an increasing variance of the estimation of the gradient with longer time horizons. Note that the key feature of the advantage function introduced in the preceding subsection is that it uses a time-grained evaluation of the process. Hence, an episodic reward severely hinders the means with which the convergence of the algorithm is improved. We were not able to achieve good convergence with this approach.

Therefore, in this work, we propose a novel approach where we incorporate a non-additive reward into the standard mathematical framework of RL. We redefine the optimization objective as

$$\underset{\theta}{\text{maximize}} \underset{\tau\sim\pi_\theta}{\mathbb{E}} \left[R_{\mathrm{BB84}}(\tau)\right]. \tag{22}$$

Therefore, the value functions now read

$$V^\pi(s_t) = \underset{a_t, s_{t+1}, \ldots \sim \pi_\theta}{\mathbb{E}} \left[R_{\mathrm{BB84}}(\tau_t)\right] \tag{23}$$

and

$$Q^\pi(s_t, a_t) = \underset{s_{t+1}, a_{t+1}, \ldots \sim \pi_\theta}{\mathbb{E}} \left[R_{\mathrm{BB84}}(\tau_t)\right], \tag{24}$$

where $\tau_t = (s_{t+1}, a_{t+1}, \ldots s_T)$ is the trajectory after time step $t$.

This new generalization replaces the accumulated reward after a time step by the secret key rate after that time step. Even though within the scope of this work we did not find a rigorous proof that this fulfills the condition of Eq. (15) and thus is an equivalent optimization, it did deliver well converging results empirically. Based on this approach we achieved reasonable convergence with good results which we will present in Sec. IV E. It is important to note that the lack of a rigorous proof for the optimization does not weaken the numerical results obtained, since the performance of a policy is evaluated independent of how the policy is found.

In our implementation, as is common practice, we used an actor-critic function $A_t$ to estimate the advantage function,

$$A_t = R_{\mathrm{BB84}}(\tau_t) - V^\pi(s_t), \tag{25}$$

where an artificial neural network is employed to approximate $V^\pi(s_t)$, which is trained on the same training data the agent is trained on. More details on the implications and properties of our advantage function and on the topic of discounting rewards can be found in the App. C. Below the Box Algorithm 1 shows the pseudo-code of our implementation. Finally, we would like to stress that generally this adaption could be applied to other optimization problems with non-additive rewards.

### C. Application of the algorithm to our physical model

In every time step the agent is given the decision to discard any of the bipartite quantum states stored in the

**Algorithm 1:** PPO Algorithm adapted from Ref. [56]

**while** *not converged* **do**

1. Collect set of trajectories $D_k = \{\tau_i\}$ by running policy $\pi_k = \pi(\theta_k)$ in the environment.
2. Compute reward functions $R_t = R_{\mathrm{BB84}}(\tau|_{t' \geq t})$ for every trajectory $\tau \in D_k$ and time step $t \in [0, T]$.
3. Compute $A_t = R_t - V_{\phi_k}(s_t)$.
4. Update the policy by maximizing the RL objective

$$\theta_{k+1} = \underset{\theta}{\operatorname{argmax}} \frac{1}{|D_k|T} \sum_{\tau \in D_k} \sum_{t=0}^{T} \min\left( \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_k}(a_t|s_t)} A_t, \operatorname{clip}\left( \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_k}(a_t|s_t)}, 1-\epsilon, 1+\epsilon \right) A_t \right) \tag{26}$$

   via Adam stochastic gradient ascent [61].

5. Fit value function by regression on mean-squared error,

$$\phi_{k+1} = \underset{\phi}{\operatorname{argmin}} \frac{1}{|D_k|T} \sum_{\tau \in D_k} \sum_{t=0}^{T} \left( V_\phi(s_t) - \Phi_t \right)^2 \tag{27}$$

   via Adam stochastic gradient descent [61].

quantum repeater based on the observation of the current state of the environment. It should be pointed out that this allows the agent to develop a vast range of complex policies. The agent is able to make decisions about each quantum state based on an arbitrary selection of the available information. [1] The decision about when to perform a swapping operation is not under control of the agent. As discussed in Sec. II D, the chosen swapping strategy is fixed to "swap as soon as possible".

For the present work, we decided to give the agent at any time the entire information about the current state of the environment. This means that the agent can make decisions on how to operate a quantum memory based on information that physically would not be immediately available at that time for a local operation on the memory. This assumption can be motivated by two primary reasons. The first is that the all-knowing agent can more easily find good solutions. When trying to implement algorithms that are sensitive to their tuned hyperparameters, it is usually a good approach to start with simpler problems and increase the complexity when the first results are obtained. The second reason is that one can possibly learn more about the behavior of the problem within the parameter space when as much information as possible is given into the analysis. This may also be considered as a first step towards methods where an agent is trained to mimic the behavior of the all-knowing agent as well as possible while having only the restricted information of a realistic scenario available. This two-step training of an agent could prove more efficient than trying to learn directly on the realistically restricted information. Lastly, it would be highly non-trivial to also include all classical communication in the simulation.

The benchmarks for the agents to surpass are the best naive strategies as determined in Sec. III B, which includes schemes both with and without cut-off.

*Remark:* It is not a priori known for which experimental setup the cut-off policy does not provide any benefit. Therefore, there is no reasonable measure for which the no-cut-off approach is as good as the cut-off approach. The cut-off policy can never become worse, as it emulates the no-cut-off approach for large cut-off parameter. Thus, we chose to take the maximum of all data points simulated in Sec. III for each experimental setup as the benchmark. This includes cut-off and no-cut-off strategies. For those cases where the cut-off does not offer an improvement, this means that the approximately optimal cut-off was run multiple times as the secret key rate converged for larger cut-offs, as discussed in Sec. III B. Since from these runs the maximum and not the average was taken, the benchmark is slightly biased towards better secret key rates for these cases. We expect this bias to be negligible for the comparison, since the uncertainties are small enough to prevent significantly overestimated rates.

### D. Implementation and experimentation

#### 1. Implementation

The learning algorithm of this work is a python implementation within the open-source DRL framework "OpenAI Spinning Up" [56] by "Open AI" which uses the "Open AI Gym" [62]. The Spinning Up framework provides some standard learning algorithms, a logger to store

[1] These policies can even emulate other repeater schemes. For example, a sequential repeater scheme like that in Refs. [14, 47]. This case is exactly emulated by discarding any quantum states that are generated in segments where the sequential scheme would not attempt any entanglement generation.

a variety of diagnostics during a learning run, as well as plotting tools to display them. Our implementation is a modified version of the PPO implementation provided by Spinning Up.

Internally the neural network and the sampling of actions is implemented using the "TensorFlow" [63] library.

The features that had to be modified and added to the implementation of the Spinning Up software will be explained in the following:

- The calculation of the secret key rate, raw rate, and average fidelity of a trajectory was implemented and included in the main loop of experience collection and storage.
- The neural network implemented in the framework did not support multidimensional action spaces. Therefore, the architecture of the neural network and the calculation of the action probabilities $\pi_\theta$ were generalized and adjusted.
- The PPO algorithm used a generalized advantage estimation [59], which is incompatible with the non-additive rewards of QKD. Hence, this was replaced by our generalized actor-critic function of Eq. (25).

The experience collection loop and the stochastic gradient ascent via Adam support multi-process parallelization via MPI (message passing interface) [64].

The final evaluation of the performance of a learned policy was done by loading its neural network into the simulation framework of Sec. III. Each agent was evaluated over 100 trajectories. The length of the trajectories for the simulation of each agent's policy can be found in App. D. Moreover, a program was implemented to store all the states seen by an agent and the count of actions it took separated for each state. This will be later used to analyze the behavior of the agents.

*Remark:* The implementation also stores the neural network representing the agent at intermediate epochs. Thus, one could analyze the agent at predetermined points in the learning process.

### 2. Hyperparameters

Below we list the tuned hyperparameters of the implemented algorithm. For details of those parameters which are not discussed in this section, we refer the reader to the original proposal of the algorithm [58] and the documentation of the original implementation [56]:

- Number of hidden layers, as well as the number of neurons for each hidden layer for the neural networks of the agent and value function estimation. Note that the numbers of neurons in the input and output layer are fixed by the state- and action-space of the environment, respectively.
- Activation functions for each layer of the two neural networks.
- Length of the simulated trajectories.
- Number of simulated trajectories per epoch.
- Learning rate for the policy.
- Learning rate for the value function $V^\pi(s_t)$.
- $\epsilon_{\text{clip}}$: Hyperparameter of the clipped objective of PPO in Eq. (26).
- Number of gradient steps in the policy update in one epoch.
- Number of gradient steps in the value function update in one epoch.
- Maximum Kullback-Leibler divergence which stops the gradient update steps of the policy early if exceeded.
- $\epsilon_{\text{Adam}}$: A parameter in the Adam optimization to improve numerical stability.

The agent as well as the value function estimation are represented by a four-layer (output and input layer, two hidden layers) artificial neural network. The hidden layers consist of 32 neurons each. For the activation function of the hidden layers of the agent the tangens hyperbolicus was chosen. The output layer of the agent uses a sigmoid activation in order to interpret the outputs as probabilities. These hyperparameters yielded good convergence properties empirically. The dimension of the input and output layers of the agent corresponds to the dimensions of state and action space of the MDP, respectively. This is 10 and 9 for a four-segment quantum repeater, respectively. The input layer of the value function estimation is also determined by the dimension of the MDP state space and the output dimension is one, thus yielding a scalar value of the value function. The choices for the rest of the hyperparameters are listed in Sec. E.

### 3. Experimentation

The process of the simulated quantum repeaters is highly probabilistic, resulting in high variances for the results of simulated trajectories of short length. This, together with the lack of discounting, leads to slow convergence of the algorithm.

In consequence, hyperparameter tuning turned out to be an extensive task in the setting of this work. The final tuning achieved visible learning progress in the order of an hour, which gives an impression of how time-consuming and challenging examining the vast hyperparameter space is where learning progress is significantly less.

With computation time being a major limiting factor and the algorithm being inherently well suited for parallelization, the use of a high-performance computation cluster could very well enable significantly more efficient

| Figure | $L_0$ (km) | $\tau_c$ (ms) | $\frac{R_{\text{agent}}}{R_{\text{no cut-off}}}$ | $\frac{R_{\text{agent}}}{R_{\text{benchamrk}}}$ |
|---|---|---|---|---|
| 10a | 20 | 0.1 | $18.0 \pm 0.2$ | $0.984 \pm 0.006$ |
| 10b | 20 | 10 | $1.001 \pm 0.002$ | $0.997 \pm 0.002$ |
| 10c | 35 | 10 | $1.014 \pm 0.004$ | $1.001 \pm 0.004$ |
| 10d | 50 | 100 | $0.996 \pm 0.003$ | $0.984 \pm 0.003$ |

TABLE I: Comparison between the learned policies and the benchmarks for those cases where the benchmark was approximately matched. $\tau_c$ is the coherence time of the memories for a bipartite quantum state.

experimentation. Unfortunately, the implementation on a cluster was not possible for this work.

The used TensorFlow implementation of the Adam optimizer was not always numerically stable. That caused the neural network to have "NaN" entries, forcing the learning algorithm to stop. This was counteracted by empirically adjusting the hyperparameter $\epsilon_{\text{Adam}}$, which helped but did not solve the problem entirely. As the functionality to resume aborted learning processes was not implemented within the scope of this work, the occurrence of this case was the reason that some of the learning processes presented in this section were not continued, even though they did not converge.

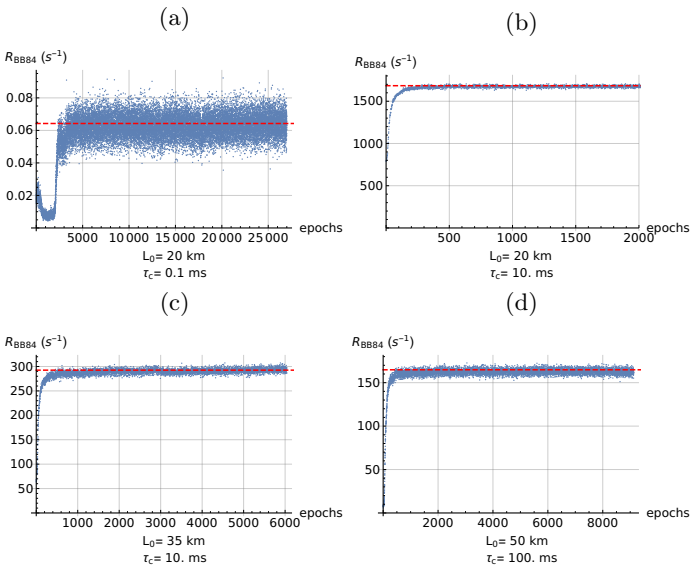


FIG. 10: Learning progress where the performance of the learned policies approximately matched the benchmark marked by the red dashed line. The average over the secret key rate of all trajectories for each epoch is plotted. An epoch is one iteration of the main loop of the algorithm. The runs 10a, 10b and 10d converged without breaking numerically. For these, a long stagnating part after the displayed epochs is not shown. $\tau_c$ is the coherence time of the memories for a bipartite quantum state.

### E. Results

The DRL implementation was applied to some of the parameter points of Sec. III B. The neural network of all stored agents of the learning runs as well as the gym environment simulating the quantum repeaters are publicly available in Ref. [65].

#### 1. *Performance and learning progress of the agents*

The learned policies can be classified into two categories. In the limiting cases where the naive cut-off offers no improvement on the secret key rate or where the optimal cut-off is one, the learning algorithm approximately reproduced these results. The learning process of these runs is plotted in Fig. 10. In other cases where the naive cut-off offered a significant improvement, the learning algorithm found a policy with even better performance. These cases are plotted in Fig. 11. The list of hyperparameter choices for each run, as well as some discussion on the runtime, can be found in App. E.

Tables I and II display the comparison of the learned policies with the benchmark strategies and a no-cut-off strategy.

Table I shows for the first category that the performance of the learned policies is approximately the same as the benchmark. The cases where the benchmarks were matched are exclusively those cases where the parameters correspond to limiting cases identified in Sec. III B, where the optimal strategy is either having no cut-off or the minimal possible cut-off. Therefore, it seems reasonable that in these regimes, there is no more room for improvement for the policies.

Table II presents the advantage of the learned policies that surpassed the benchmark. The narrow confidence intervals of the ratios indicate that the learned policies outperform the benchmarks. This is further shown via the quantity $\frac{r-1}{\Delta r}$, which is the difference between the measured ratio and the ratio a strategy with no benefit over the benchmark would yield in units of its uncertainty. This, therefore, prevents the conclusion that the result is a statistical fluctuation. The improvements over the benchmarks range from 2% to 38%. Furthermore, apparently the more improvement the naive cut-off offers over a quantum repeater without cut-off, the more potential lies in the learned policies for this quantum repeater, since this correlation is visible in the achieved and presented examples. It is worth mentioning that this is an analog result to what was stated in Sec. III C for the naive cut-off policy.

| Figure | $L_0$ (km) | $\tau_c$ (ms) | $R_{\text{agent}}\,(s^{-1})$ | $\frac{R_{\text{agent}}}{R_{\text{no cut-off}}}$ | $r = \frac{R_{\text{agent}}}{R_{\text{benchamrk}}}$ | $\frac{r-1}{\Delta r}$ |
|---|---|---|---|---|---|---|
| 11a | 20 | 1 | $350 \pm 8$ | $1.172 \pm 0.004$ | $1.068 \pm 0.003$ | 20.4453 |
| 11b | 35 | 1 | $15.1 \pm 0.7$ | $2.67 \pm 0.03$ | $1.31 \pm 0.01$ | 32.2178 |
| 11c | 50 | 1 | $0.51 \pm 0.01$ | $12.3 \pm 0.1$ | $1.382 \pm 0.009$ | 42.9056 |
| 11d | 50 | 10 | $50.3 \pm 0.8$ | $1.180 \pm 0.003$ | $1.024 \pm 0.002$ | 10.4957 |
| 11e | 52.6569 | 1.45271 | $1.43 \pm 0.04$ | $6.45 \pm 0.04$ | $1.218 \pm 0.005$ | 40.0445 |

TABLE II: Comparison between the learned policies and the benchmarks in the cases where the benchmark was surpassed. $\tau_c$ is the coherence time of the memories for a bipartite quantum state.

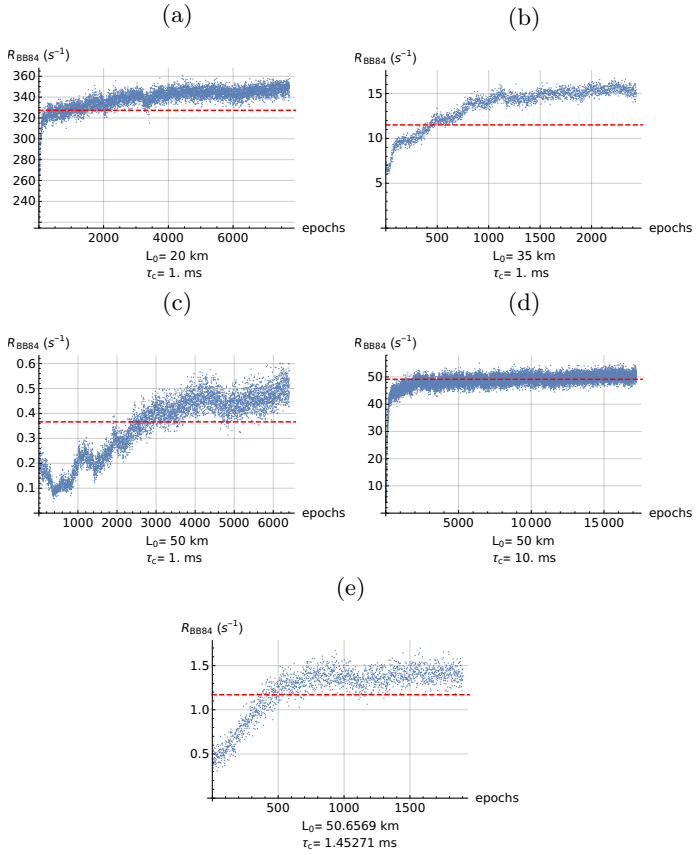


FIG. 11: Learning progress where the performance of the learned policies surpassed the benchmark marked by the red dashed line. The average over the secret key rate of all trajectories for each epoch is plotted. An epoch is one iteration of the main loop of the algorithm. $\tau_c$ is the coherence time of the memories for a bipartite quantum state.

### 2. *Discussion on the learned policies*

Extracting an in-depth understanding of a policy from the neural network of the agent is a non-trivial task and an open question of current research [66]. Nevertheless, some aspects gained by looking at the policies of the presented results shall be discussed in this subsection.

The DRL algorithm optimizes probabilistic policies, where the probability can be interpreted as how certain an agent is about the optimality of an action for a given state. In order to output probabilities, a sigmoid function is applied to the output layer of the neural network representing the policy. The sigmoid function converges asymptotically to one and zero for large absolute output weights. First, this means that in order to get an approximately guaranteed probability, comparably large parameter changes are necessary. Secondly, a real guaranteed probability can, in theory, never be achieved. This causes even well converged policies to take actions that they evaluate to be non-ideal occasionally. This is useful in the learning process, as the agent occasionally reviews these actions, but slightly reduces the final performance of the agent. A suggestion to improve this is to map probabilities above a high threshold to one and analogously for low probabilities to zero in the final policies. Nonetheless, this has not been done in this work.

The actions of those policies that matched the benchmark were confirmed to reproduce their respective benchmark strategy approximately. The agents that surpassed the benchmark are more challenging to comprehend. A complete understanding of the policies could not be achieved in this work. Nevertheless, three observed patterns that give insights about how the policies achieved better performance are given in the following:

- The policies keep quantum states with higher storage times shared between nodes that are farther apart compared to quantum states shared between closer nodes, which are discarded much earlier.

- The policy's decision to keep or discard appears to be influenced by the existence and quality of other quantum states in the repeater. If, for example, in the proximity of a quantum state, another quantum state exists, the agent is more likely to keep the quantum state.

- The policy discards a quantum state comparably early if a single bipartite quantum state is shared

between the two nodes surrounding the first quantum state.

These strategies might seem intuitive or even obvious and could probably be thought of without a machine learning approach. The advantage of the machine learning approach, however, is the scaling of these strategies. Even though the concept might be clear, finding exact numbers remains a task on its own, for which DRL is used here.

On the other hand, some decisions of the agents were obviously unreasonable. This includes decisions of well converged probabilities. It shows that the full potential of the DRL approach has not been achieved yet in this work and that there is still room for improvement in the algorithm and policies. The following examples are two decisions that an agent made with high probability that fall into this category:

- In the case mentioned above, where a policy would discard a quantum state $A$ if it is the only one left in the repeater but then keeps it because of another quantum state $B$ in the proximity, some unreasonable actions occurred. In some cases, it kept the state $A$ but discarded the state $B$, even though the decision to keep state $A$ was dependent on the existence of state $B$.

- In a scenario where two bipartite states with exactly symmetric positions in the repeater where stored, the agent decided to discard the one with the lower and to keep the one with the higher storage time.

Another observation worth mentioning is that the policies tended to be asymmetric in the sense that two symmetric segments were often treated differently. This might happen because of non-ideal convergence, but one should note that nothing contradicts the possibility that an asymmetric policy might be optimal, despite the structure of a four-segment repeater being symmetric.

Furthermore, the policies were less converging in their probabilities than their convergence in performance might suggest. This could be explained by the fact that two different actions do not necessarily cause different secret key rates. Some actions might be equivalent or very close in performance. Thus, the agent will not, or will only very slowly, learn to choose one action over the other.

### F. Future work

One suggestion for further improvements is based on the fact that states of the environment have significantly different frequencies in a trajectory. Therefore, the agent has superfluous experience about prevalent states of the MDP while having very sparse information about others. This creates an imbalance and thus sample inefficiency. A suggestion to counteract this is to create an environment that uses a random initial state of the trajectory in ways that flatten the distribution of the states the agent sees.

The next step towards analyzing more realistic repeater chains would be to include classical communication by having multiple cooperating agents where each agent has access only to the information available at one corresponding repeater node.

### G. Deep reinforcement learning - conclusion

A proof of concept that DRL can be exploited for QKD via quantum repeaters to find sophisticated policies that improve the secret key rate over "naive" approaches has been achieved. The results of the learned policies ranged from a reproduction of performance up to a 38% improvement over the naive simulation approach of Sec. III. Furthermore, we have identified shortcomings of the policies and suggested enhancements to the algorithm to further optimize policies, showing potential room for improvement. Additionally, utilizing a high-performance computation cluster would probably go a long way in enhancing the effectiveness with which solutions can be found. This provides motivation to further pursue DRL algorithms in control problems of multi-segment quantum repeaters.

The improvements of the learned policies as presented here might seem small for making a strong impact on long-distance QKD applications. This, however, could be misleading, as the examined repeaters only had four segments. Practical applications most likely will use considerably more segments. This increases the number of possible policies and, in particular, the asymmetry between the repeater nodes. Therefore, we would intuitively expect the improvements to increase for more segments, as the parameter space grows in complexity. Additionally, the complexity of quantum repeaters generally grows with the number of segments, making other approaches less feasible, while DRL is especially suited for problems that can be simulated but no longer analytically analyzed. This emphasizes the meaning of a proof of concept. It has been conceptually shown that DRL can lead to improvements, while the full potential of the approach remains to be seen.

## V. CONCLUSION

In this work, a Markov decision process (MDP) to model the generation of spatially distributed, entangled quantum states via a memory-based quantum repeater was developed. In principle, the model can describe quantum repeaters with arbitrary numbers of segments and include arbitrary qubit Pauli and erasure channels as error sources. Moreover, entanglement swapping and discarding of quantum states are actions available to operate the quantum repeater.

Based on this MDP, a simulation was implemented. Among the above-mentioned general error channels, the particular error sources included in the simulation are the random, time-dependent dephasing of quantum states in the memories and distance-dependent photon losses in the optical quantum channels (as well as constant losses or inefficiencies at the repeater stations and interfaces). The entanglement swapping is assumed to be error-free and deterministic, and all schemes in this paper "swap as soon as possible" (which is the optimal strategy in the presence of memory dephasing). The simulation was used to analyze the secret key rate of the BB84 quantum key distribution (QKD) protocol via four-segment quantum repeaters in a broad parameter space of the segment length and the quantum memories' coherence time. Moreover, the simulation was used to analyze the behavior of the secret key rate under the variation of a controlled limited storage time of the quantum states - the so-called memory cut-off - in the same parameter space. It was found that there exists a parameter regime in the limit where the memory coherence times are so large or small in relation to the segment length that the ideal control is to discard no or all intermediate quantum states, respectively.

Furthermore, a deep reinforcement learning (DRL) algorithm was implemented to examine the possibilities of finding sophisticated solutions for the control of the quantum memories improving the secret key rate of quantum repeaters. The algorithm of choice was a proximal policy optimization (PPO) [58], which was applied to the aforementioned simulated quantum repeaters. First, the results of the limit parameter regime were successfully reproduced. More specifically, and most importantly, the algorithm found policies outperforming the benchmarks provided by the previously employed, "naive" simulation (i.e., a standard simulation without the help of a learning agent), serving as a proof of concept that DRL can indeed offer a valid approach to optimize the memory storage times. This proof of concept is a first step in laying the groundwork to develop and apply DRL algorithms to realistic and practical long-distance quantum key distribution.

The DRL algorithm introduced and employed here, though adapted to the special problem of computing secret key rates in quantum repeaters subject to memory dephasing, is not yet an optimal algorithm. First, it suffers from rather slow convergence properties due to the sparse entanglement distribution per time step for realistic parameters and the lack of an additive reward model. Hence, a next step would be to increase computational power to improve the efficiency with which policies can be found. A further suggestion is to modify the experience loop of the algorithm in a way that flattens the distribution of environment states the agent interacts with to improve the balance of the gradient steps for the different states.

Examining the achieved policies, it is clear that these are not perfect. This suggests that there still is further potential in the approach, even for the examined repeater constellations. In the future, one could expand the simulation to more realistic scenarios, for example, by including more error sources and classical communication, therefore finding solutions for practical settings. In conclusion, the full potential of optimizing policies for quantum repeaters using DRL remains to be seen.

Long-distance key distribution based on quantum repeaters and DRL are both rapidly moving and developing fields, with their full future impact being unforeseeable. This work contributes new insights by combining the two fields and showing viability, as well as problems and limitations of the approach.

## ACKNOWLEDGMENTS

We thank the BMBF in Germany for support via PhotonQ, Q.Link.X/QR.X and the BMBF/EU for support via QuantERA/ShoQC. We also thank Frank Schmidt and Evgeny Shchukin for helpful discussions.

## Appendix A: Markov decision process modeling a quantum repeater

Here we will formulate the model by explaining each element of the tuple $(S, A, P_a, R_a)$.

### 1. States $S$

Any state of the Markov model is fully described by the spatially separated bipartite quantum states stored in the quantum memories. Thus, in a general treatment, the full density matrices of all quantum states would be necessary to determine the state of the MDP. By restricting the errors to be Pauli channels as defined in Eq. (1), it is possible to define the states in a much more convenient form. In this restricted scenario, it follows directly from the fact that Pauli channels and entanglement swapping commute, that treating the precise development of all quantum states is equivalent to treating a noiseless quantum repeater and applying any contributing errors to the final quantum state. That means it is sufficient to keep track of any error occurring and propagating the counts additive through the swapping. Therefore, all possible states of the MDP can be encoded in a triangular matrix $S$, with rows and columns corresponding to the repeater nodes. The entries $s_{ij}$ are either the symbol $\xi$, indicating that the two repeater nodes $i$ and $j$ do not share an entangled state or a vector with an integer component for each error, storing its accumulated count on the bipartite quantum state.

Generally, the number of possible states is infinite. In realistic scenarios, the memory strategy restricts the transitions between the states in a way that only a finite

amount of states remains accessible with non-vanishing probability.

### 2. Actions $A$

The set of possible actions $A$ can be separated into two independent sets of actions, $A = A_s \bigotimes A_d$.

- The first set corresponds to entanglement swapping. It is encoded in a vector $A_s$ with a boolean component for each node. If a component is true, the entanglement swapping operation is performed at that node. Otherwise, no action will be performed at that node.

- The second set corresponds to the discarding of states. It is encoded in a triangular matrix $A_d$ with a boolean entry for each pair of two different nodes. If an entry is true, the bipartite quantum state of this pair is discarded. Otherwise, no operation will be performed for this quantum state.

It is important to note that in the treatment of the MDP, these actions are instantaneous. The reason for this is that the MDP only models the development of the quantum states, which is independent of classical communication. That does include entanglement swapping. Even though entanglement swapping requires classical communication, the time when classical communication is performed is irrelevant for the development of the quantum states and can thus be excluded from the treatment.

### 3. Transitional properties $P_a$ and $R_a$

To find a complete, analytic expression for $P_a$ and $R_a$ is an infeasible task. Fortunately, this is not necessary to simulate the process. It is sufficient to find mathematical operations realizing the simulated processes on the states. The time step of the process is defined by a round of classical communication, which is the time it takes for a node to send classical information to an adjacent node. This time step is the time a quantum state will be stored when initially distributed in one segment before the corresponding nodes communicated the successful distribution. In the following, it will be described how each process that takes place in one time step is simulated.

1. Initial entanglement generation in each segment:
   Within one time step in each segment, an attempt to distribute initial entanglement is performed. This is simulated via sampling in each segment respective to an entanglement generation rate $p$.

2. Natural, uncontrolled development of the quantum states:
   The count of any error that occurs on stored states is increased by one in the tuples in the state matrix $S$. The for this work relevant example is the dephasing channel.

3. Performing the actions:

   (a) Entanglement swapping at each station $j$ where $(A_s)_j$ is true :
   For any tuple $\{i,k\}$, where $s_{i,j} \neq \xi$ and $s_{j,k} \neq \xi$:

      i. $s_{i,k} \leftarrow s_{i,j} + s_{j,k}$, where "+" is the element wise addition of the two vectors. This follows from the fact that Pauli channels and entanglement swapping commute with each other.

      ii. $s_{i,j} \leftarrow \xi$
      $s_{j,k} \leftarrow \xi$

      iii. Increase the count of any error that occurs due to entanglement swapping by one in the vector $s_{i,k}$.

   (b) State discarding:
   For each tuple $\{i,j\}$ where $(A_d)_{i,j}$ is true, set $s_{i,j} \leftarrow \xi$.

4. Return the immediate reward $r = F(s_{0,n})$ and set $s_{0,n} \leftarrow \xi$. Where $F(s)$ is the fidelity between the quantum state determined by the tuple $s$ and the quantum state, a noiseless repeater would have distributed. This could be kept more general, by instead returning the vector $s_{0,n}$ and thus the entire quantum state, but this was not necessary for this work.

### 4. Example step of the MDP

In Fig. 12 one example step of the MDP is illustrated. In this example, the only error is the dephasing of the quantum memories, which accumulates on a state for every time step it is stored. Therefore the entries $s_{j,k}$ are the cumulated storage time of the quantum state stored in the nodes $j$ and $k$.

The state of the MDP prior to the step is shown in Fig. 12(a). The state after the uncontrolled development, which consists of generating initial entanglement and accumulating the storage time is shown in Fig. 12(b). The controlled swapping operation is depicted in Fig. 12(c) with the resulting state in Fig. 12(d). The discarding action is displayed in Fig. 12(e) with the resulting state in Fig. 12(f).

## Appendix B: Secret key rate of entanglement-based BB84 in the presence of dephasing

In the following, we will calculate the secret key rate for a BB84 protocol where the quantum states are subject to a dephasing channel.

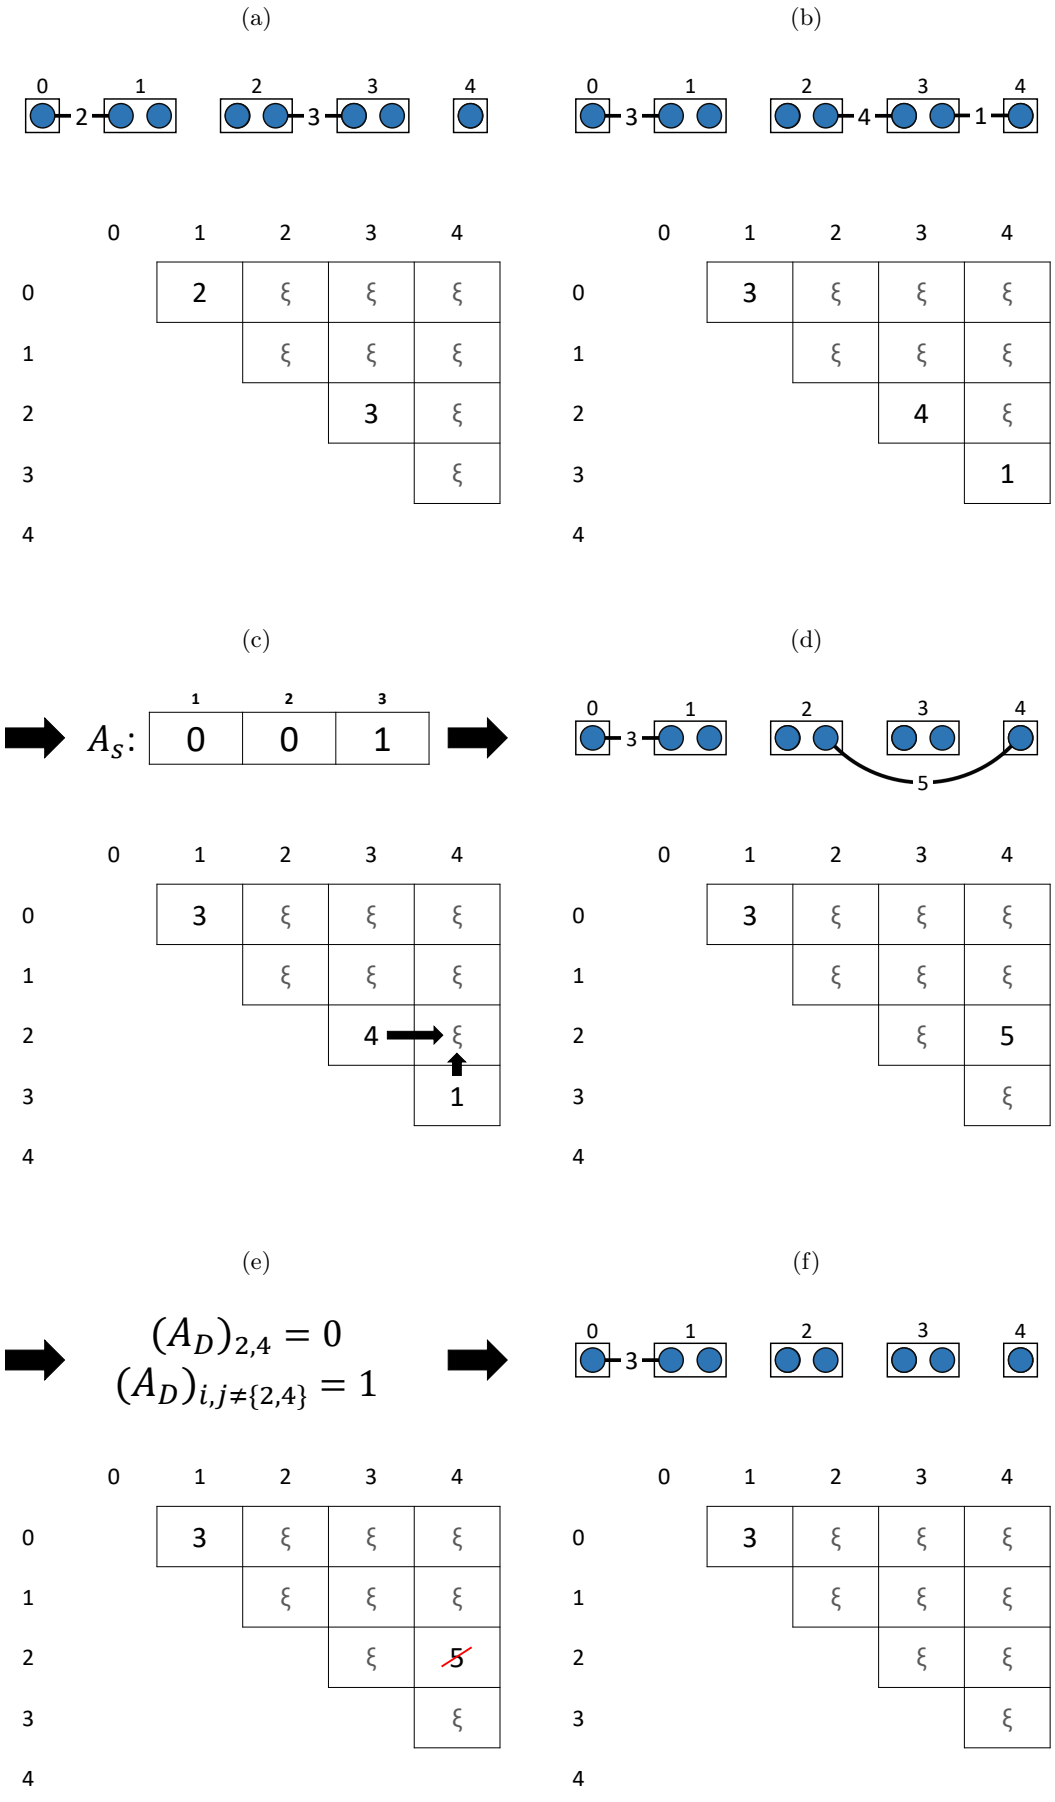


FIG. 12: Illustration of one example step of the Markov decision process.

Without loss of generality, we will choose the initially distributed, undisturbed entangled states as $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. The form of a degraded state follows directly from Eq. (8) and reads

$$\rho = (1-\nu)\,|\Phi^+\rangle\langle\Phi^+| + \nu\,|\Phi^-\rangle\langle\Phi^-|\,, \tag{B1}$$

where

$$0 < \nu < \frac{1}{2}. \tag{B2}$$

The bit error rates in the $X$ and $Z$ bases are defined as

$$e_Z = \mathbb{E}\left[1 - \langle 00|\,\rho\,|00\rangle + \langle 11|\,\rho\,|11\rangle\right] \tag{B3}$$

and

$$e_X = \mathbb{E}\left[1 - \langle ++|\,\rho\,|++\rangle + \langle --|\,\rho\,|--\rangle\right]. \tag{B4}$$

A straightforward calculation yields

$$e_Z = 0 \tag{B5}$$

and

$$e_X = \mathbb{E}\left[\nu\right]. \tag{B6}$$

Inserting the bit error rates into Eq. (3) and Eq. (4) gives

$$R_{\mathrm{BB84}} = Y \cdot \left(1 - h\left(\mathbb{E}\left[\nu\right]\right)\right). \tag{B7}$$

Alternatively, by inserting the exponential dephasing channel defined in Eq. (8), Eq. (B7) takes the form

$$R_{\mathrm{BB84}} = Y \cdot \left(1 - h\left(\frac{1}{2}\left(1 - \mathbb{E}\left[e^{-\frac{t}{\tau_c}}\right]\right)\right)\right), \tag{B8}$$

where $\tau_c$ is the coherence time of the quantum memory for a bipartite quantum state and $t$ the storage time of the quantum state.[2]

## Appendix C: Discounting rewards

### a. *Finite and infinite horizons and the relevance of discounting rewards*

The in Sec. IV A introduced discounted reward function of Eq. (16) is a fundamental objective function in RL:

$$R_t = \sum_{l=0}^{T-t} \gamma^l r_{t+l}. \tag{C1}$$

A discounting parameter $\nu$ smaller one introduces the discounting into the accumulated reward. This servers to intuitive purposes. First, it ensures that the objective function remains finite for "infinite-horizon" trajectories (T $\to \infty$). Secondly, and maybe most importantly for most practical applications, it makes the agent weight sooner rewards more than rewards in the farther future. This improves the evaluation of an action as it is bias towards its short-term consequences rather than on rewards in the distant future, which is increasingly independent from the action at the time $t$. In this way, discounting reduces noise in the objective function.

As was stated before, the non-additive rewards in this work are not suited for discounting in this way. Even more detrimental , non-negative rewards can actually decrease the secret key rate of a trajectory. The lack of discounting necessitates the use of a finite-horizon objective function, which makes good gradient estimation difficult. This can be explained by understanding that, when the trajectory is too long, the objective function becomes very noisy, since for earlier actions, the long trajectory introduces rewards that are mostly independent from early actions. On the other hand, if the trajectory is too short, delayed rewards of an action might not be obtained, cause the trajectory ended, causing a wrong evaluation of the action.

In consequence, good convergence is expected to be significantly harder to achieved compared to a problem setting where discounting can be used.

### b. *A proposal for generalized discounting*

In this work an attempt was made to propose discounted objective function of an arbitrary (not necessarily additive) reward function $\Psi(\tau)$ via:

$$\Phi_t = \sum_{l=0}^{T-t} \gamma^l \Psi(\tau_t). \tag{C2}$$

Applied to our optimization this reads,

$$\Phi_t = \sum_{l=0}^{T-t} \gamma^l R_{\mathrm{BB84}}(\tau_t). \tag{C3}$$

This seemed like a reasonable approach intuitively but yielded no useful results during the minimal testing we did. Therefore, this method was not used further. Even though, unfortunately, this could not be achieved here, with sufficient hyperparameter tuning this method could possibly improve the convergence properties of the optimization significantly, as of the reasons discussed above.

## Appendix D: Trajectory lengths of the simulations

In this appendix the lengths of the trajectories used in the simulations and the experience collection of the agents is given in Tab. III.

[2] In practical finite-size applications one can use confidence intervals on the estimation of the expectation value to determine an error rate which is larger than the true expectation value with approximately guaranteed probability.

| $L_0$ (km) | $\tau_c$ (ms) | $T$, no cut-off | $T$, cut-off | $T$, agent |
|---|---|---|---|---|
| 20 | 0.1 | $10^4$ | $10^4$ | $10^5$ |
| 20 | 1 | $10^4$ | $10^4$ | $10^5$ |
| 20 | 10 | $10^4$ | $10^4$ | $10^5$ |
| 20 | 100 | $10^4$ | $10^4$ | |
| 35 | 0.1 | $10^4$ | $10^4$ | |
| 35 | 1 | $10^4$ | $10^4$ | $10^5$ |
| 35 | 10 | $10^4$ | $10^4$ | $10^5$ |
| 35 | 100 | $10^4$ | $10^4$ | |
| 50 | 0.1 | $10^6$ | $10^6$ | |
| 50 | 1 | $10^5$ | $10^4$ | $10^5$ |
| 50 | 10 | $10^5$ | $10^4$ | $10^5$ |
| 50 | 100 | $10^4$ | $10^4$ | $10^5$ |
| 50.6569 | 1.45271 | $10^5$ | $10^5$ | $10^5$ |
| 70 | 0.1 | $10^5$ | $10^6$ | |
| 70 | 1 | $10^5$ | $10^5$ | |
| 70 | 10 | $10^5$ | $10^5$ | |
| 70 | 100 | $10^5$ | $10^5$ | |

TABLE III: Length of the trajectories of the simulations.

### Appendix E: Hyperparamters of the DRL runs

The DRL program was run with four processes on a desktop computer with a four-core cpu.[3] An epoch with four trajectories with each 8000 simulated time steps took about 13 seconds on average. Therefore, a thousand epochs take around 3.6 hours. In the operated hyperparameter space, the computation time was roughly linear with the number of simulated time steps within the epoch. This serves as an orientation for the required computation time of each run.

The architecture of the neural networks representing the policy and the value function was chosen identical for every learning run as it seemed sufficient in complexity and degrees of freedom for the task, and smaller architectures did not improve computation time significantly, empirically. The architecture consisted of 2 hidden layers with 32 neurons each. The activation function of the hidden layers was chosen to be the tangens hyperbolicus. The neural network representing the policy applies a sigmoid as the activation on the output layer so that the final output can be interpreted as probabilities.

Abbreviations for the hyperparameters:

- L: Length of the simulated trajectories.
- N: Number of simulated trajectories per epoch.
- $\alpha_\pi$: Learning rate for the policy.
- $\alpha_V$: Learning rate for the value function.
- $\epsilon_{\text{clip}}$: Epsilon of the clipped objective of PPO in Eq. (26).
- $n_\pi$: Number of gradient steps in the policy update in one epoch.
- $n_V$: Number of gradient steps in the value function update in one epoch.
- KL: Maximum KL-divergence which stops the gradient update steps of the policy early if exceeded.
- $\epsilon_{\text{Adam}}$: The parameter in the Adam optimization to improve numerical stability.

[3] Intel(R) Core(TM) i7-3770 CPU @ 3.40GHz

| figure | $L_0$ (km) | $\tau_c$ (ms) | L | N | $\alpha_\pi$ | $\alpha_V$ | $\epsilon_{\text{clip}}$ | $n_\pi$ | $n_V$ | KL | $\epsilon_{\text{Adam}}$ |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 10a | 20 | 0.1 | 800 | 4 | $1 \cdot 10^{-4}$ | $1 \cdot 10^{-3}$ | 0.2 | 120 | 80 | 0.015 | 0.01 |
| 11a | 20 | 1 | 8000 | 4 | $4 \cdot 10^{-4}$ | $1 \cdot 10^{-3}$ | 0.2 | 120 | 80 | 0.015 | 0.01 |
| 10b | 20 | 10 | 8000 | 4 | $1 \cdot 10^{-3}$ | $1 \cdot 10^{-3}$ | 0.2 | 120 | 80 | 0.015 | 0.01 |
| 11b | 35 | 1 | 8000 | 4 | $1 \cdot 10^{-3}$ | $1 \cdot 10^{-3}$ | 0.2 | 120 | 80 | 0.015 | 0.01 |
| 10c | 35 | 10 | 8000 | 4 | $4 \cdot 10^{-4}$ | $1 \cdot 10^{-3}$ | 0.2 | 120 | 80 | 0.015 | 0.01 |
| 11c | 50 | 1 | 8000 | 4 | $8 \cdot 10^{-4}$ | $1 \cdot 10^{-3}$ | 0.2 | 1000 | 80 | 0.015 | 0.1 |
| 11d | 50 | 10 | 8000 | 4 | $4 \cdot 10^{-3}$ | $1 \cdot 10^{-3}$ | 0.2 | 1000 | 80 | 0.015 | 1 |
| 10d | 50 | 100 | 8000 | 4 | $4 \cdot 10^{-3}$ | $1 \cdot 10^{-3}$ | 0.2 | 1000 | 80 | 0.015 | 1 |
| 11e | 50.6569 | 1.45271 | 8000 | 4 | $4 \cdot 10^{-4}$ | $1 \cdot 10^{-3}$ | 0.2 | 120 | 80 | 0.015 | $1 \cdot 10^{-8}$ |

TABLE IV: Hyperparameters of the learning runs presented in Sec. IV E


[1] P. W. Shor and J. Preskill, Phys. Rev. Lett. **85**, 441 (2000).

[2] C. H. Bennett and G. Brassard, Theoretical Computer Science **560**, 7 (2014), theoretical Aspects of Quantum Cryptography – celebrating 30 years of BB84.

[3] H.-L. Yin, T.-Y. Chen, Z.-W. Yu, H. Liu, L.-X. You, Y.-H. Zhou, S.-J. Chen, Y. Mao, M.-Q. Huang, W.-J. Zhang, H. Chen, M. J. Li, D. Nolan, F. Zhou, X. Jiang, Z. Wang, Q. Zhang, X.-B. Wang, and J.-W. Pan, Phys. Rev. Lett. **117**, 190501 (2016).

[4] A. Boaron, G. Boso, D. Rusca, C. Vulliez, C. Autebert, M. Caloz, M. Perrenoud, G. Gras, F. Bussières, M.-J. Li, D. Nolan, A. Martin, and H. Zbinden, Phys. Rev. Lett. **121**, 190502 (2018).

[5] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai, G.-B. Li, Q.-M. Lu, Y.-H. Gong, Y. Xu, S.-L. Li, F.-Z. Li, Y.-Y. Yin, Z.-Q. Jiang, M. Li, J.-J. Jia, G. Ren, D. He, Y.-L. Zhou, X.-X. Zhang, N. Wang, X. Chang, Z.-C. Zhu, N.-L. Liu, Y.-A. Chen, C.-Y. Lu, R. Shu, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan, Science **356**, 1140 (2017), https://www.science.org/doi/pdf/10.1126/science.aan3211.

[6] G. Vallone, D. Bacco, D. Dequal, S. Gaiarin, V. Luceri, G. Bianco, and P. Villoresi, Phys. Rev. Lett. **115**, 040502 (2015).

[7] M. Lucamarini, Z. L. Yuan, J. F. Dynes, and A. J. Shields, Nature **557**, 400 (2018).

[8] S. Wang, Z.-Q. Yin, D.-Y. He, W. Chen, R.-Q. Wang, P. Ye, Y. Zhou, G.-J. Fan-Yuan, F.-X. Wang, Y.-G. Zhu, P. V. Morozov, A. V. Divochiy, Z. Zhou, G.-C. Guo, and Z.-F. Han, Nature Photonics **16**, 154 (2022).

[9] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, Scientific Reports **6**, 20463 (2016).

[10] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, Phys. Rev. Lett. **81**, 5932 (1998).

[11] S. Slussarenko and G. J. Pryde, Applied Physics Reviews **6**, 041303 (2019), https://doi.org/10.1063/1.5115814.

[12] S. Khatri, C. T. Matyas, A. U. Siddiqui, and J. P. Dowling, Phys. Rev. Research **1**, 023032 (2019).

[13] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpędek, M. Skrzypczyk, L. Wubben, W. de Jong, D. Podareanu, A. Torres-Knoop, D. Elkouss, and S. Wehner, Communications Physics **4**, 164 (2021).

[14] L. Kamin, E. Shchukin, F. Schmidt, and P. van Loock, Exact rate analysis for quantum repeaters with imperfect memories and entanglement swapping as soon as possible (2022), arXiv:2203.10318 [quant-ph].

[15] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, Nature **414**, 413 (2001).

[16] P. van Loock, T. D. Ladd, K. Sanaka, F. Yamaguchi, K. Nemoto, W. J. Munro, and Y. Yamamoto, Phys. Rev. Lett. **96**, 240501 (2006).

[17] L. Childress, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, Phys. Rev. Lett. **96**, 070504 (2006).

[18] P. van Loock, W. Alt, C. Becher, O. Benson, H. Boche, C. Deppe, J. Eschner, S. Höfling, D. Meschede, P. Michler, F. Schmidt, and H. Weinfurter, Advanced Quantum Technologies **3**, 1900141 (2020).

[19] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, Nature **558**, 268 (2018).

[20] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, and M. D. Lukin, Nature **580**, 60 (2020).

[21] S. Langenfeld, P. Thomas, O. Morin, and G. Rempe, Phys. Rev. Lett. **126**, 230506 (2021).

[22] E. Shchukin and P. van Loock, Phys. Rev. Lett. **128**, 150502 (2022).

[23] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin, Rev. Mod. Phys. **83**, 33 (2011).

[24] F. Rozpędek, K. Goodenough, J. Ribeiro, N. Kalb, V. C. Vivoli, A. Reiserer, R. Hanson, S. Wehner, and D. Elkouss, Quantum Science and Technology **3**, 034002 (2018).

[25] F. Rozpędek, R. Yehia, K. Goodenough, M. Ruf, P. C. Humphreys, R. Hanson, S. Wehner, and D. Elkouss, Phys. Rev. A **99**, 052330 (2019).

[26] O. A. Collins, S. D. Jenkins, A. Kuzmich, and T. A. B. Kennedy, Phys. Rev. Lett. **98**, 060502 (2007).

[27] E. Shchukin, F. Schmidt, and P. van Loock, Phys. Rev. A **100**, 032322 (2019).

[28] L. Praxmeyer, Reposition time in probabilistic imperfect memories (2013), arXiv:1309.3407 [quant-ph].

[29] S. Brand, T. Coopmans, and D. Elkouss, IEEE Journal on Selected Areas in Communications **38**, 619 (2020).

[30] S. Khatri, Policies for elementary link generation in quantum networks (2020), arXiv:2007.03193 [quant-ph].

[31] B. Li, T. Coopmans, and D. Elkouss, IEEE Transactions on Quantum Engineering **2**, 1 (2021).

[32] S. Santra, L. Jiang, and V. S. Malinovsky, Quantum Science and Technology **4**, 025010 (2019).

[33] K. Goodenough, D. Elkouss, and S. Wehner, Phys. Rev. A **103**, 032610 (2021).

[34] S. E. Vinay and P. Kok, Phys. Rev. A **99**, 042313 (2019).

[35] L. Jiang, J. M. Taylor, N. Khaneja, and M. D. Lukin, Proceedings of the National Academy of Sciences **104**, 17291 (2007), https://www.pnas.org/content/104/44/17291.full.pdf.

[36] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, IEEE Signal Processing Magazine **34**, 26 (2017).

[37] Y. Duan, X. Chen, R. Houthooft, J. Schulman, and P. Abbeel, Benchmarking deep reinforcement learning for continuous control (2016), arXiv:1604.06778 [cs.LG].

[38] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, Nature **518**, 529 EP (2015).

[39] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, Trust region policy optimization (2015), arXiv:1502.05477 [cs.LG].

[40] N. Heess, G. Wayne, D. Silver, T. Lillicrap, Y. Tassa, and T. Erez, Learning continuous control policies by stochastic value gradients (2015), arXiv:1510.09142 [cs.LG].

[41] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, PRX Quantum **1**, 010301 (2020).

[42] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, Phys. Rev. X **8**, 031084 (2018).

[43] H. P. Nautrup, N. Delfosse, V. Dunjko, H. J. Briegel, and N. Friis, Quantum **3**, 215 (2019).

[44] S. D. Reiß, *Applying Deep Reinforcement Learning to Key Distribution Based on Quantum Repeaters*, Master's thesis, Johannes Gutenberg University, Mainz, Germany (Nov. 2019).

[45] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, Phys. Rev. A **53**, 2046 (1996).

[46] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Phys. Rev. Lett. **76**, 722 (1996).

[47] F. Schmidt and P. van Loock, Phys. Rev. A **102**, 042614 (2020).

[48] G. Bowen and S. Bose, Phys. Rev. Lett. **87**, 267901 (2001).

[49] H.-K. Lo, H. F. Chau, and M. Ardehali, Journal of Cryptology **18**, 133 (2005).

[50] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. Pereira, M. Razavi, J. S. Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden, Advances in quantum cryptography (2019), arXiv:1906.01645 [quant-ph].

[51] F. F. da Silva, A. Torres-Knoop, T. Coopmans, D. Maier, and S. Wehner, Quantum Science and Technology **6**, 035007 (2021).

[52] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, Rev. Mod. Phys. **81**, 1301 (2009).

[53] G. P. Agrawal, Fiber-optic communication systems, in *"Fiber-optic communication systems, 2nd ed., by Govind P. Agrawal. New York, NY: Wiley, 1997, "A Wiley-Interscience publication", ISBN 0471175404."* (1997) p. 55.

[54] S. S. Mousavi, M. Schukat, and E. Howley, in *Proceedings of SAI Intelligent Systems Conference (IntelliSys) 2016*, edited by Y. Bi, S. Kapoor, and R. Bhatia (Springer International Publishing, Cham, 2018) pp. 426–440.

[55] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (The MIT Press, 2018).

[56] J. Achiam, Spinning Up in Deep Reinforcement Learning (2018).

[57] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, Trust region policy optimization (2017), arXiv:1502.05477 [cs.LG].

[58] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms (2017), arXiv:1707.06347 [cs.LG].

[59] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, High-dimensional continuous control using generalized advantage estimation (2015), arXiv:1506.02438 [cs.LG].

[60] M. Minsky, Proceedings of the IRE **49**, 8 (1961).

[61] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization (2014), arXiv:1412.6980 [cs.LG].

[62] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, Openai gym (2016).

[63] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from tensorflow.org.

[64] C. The MPI Forum, in *Proceedings of the 1993 ACM/IEEE Conference on Supercomputing*, Supercomputing '93 (Association for Computing Machinery, New York, NY, USA, 1993) p. 878–883.

[65] S. D. Reiß, Git repository supplementary to master's thesis, https://github.com/SimonReiss/Master-Thesis.

[66] A. Dawid, P. Huembeli, M. Tomza, M. Lewenstein, and A. Dauphin, Phase detection with neural networks: Interpreting the black box (2020), arXiv:2004.04711 [quant-ph].