

Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber

Horta Ferreira da Silva, F.; Avis, G.; Slater, J.A.; Wehner, S.D.C.

DOI

[10.1088/2058-9565/ad7499](https://doi.org/10.1088/2058-9565/ad7499)

Publication date

2024

Document Version

Final published version

Published in

Quantum Science and Technology

Citation (APA)

Horta Ferreira da Silva, F., Avis, G., Slater, J. A., & Wehner, S. D. C. (2024). Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber. *Quantum Science and Technology*, 9(4), Article 045041. <https://doi.org/10.1088/2058-9565/ad7499>

Important note

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

Copyright

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

Takedown policy

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

PAPER • OPEN ACCESS

Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber

To cite this article: Francisco Ferreira da Silva *et al* 2024 *Quantum Sci. Technol.* **9** 045041

View the [article online](#) for updates and enhancements.

You may also like

- [The use of fractionated fly ash of thermal power plants as binder for production of briquettes of coke breeze and dust](#)
E Yu Temnikova, A R Bogomolov and A A Lapin
- [SRP Meeting: Radiation Emergency Preparedness, London, 24 June 1998](#)
- [A Framework for Requirements Prioritization Process in Agile Software Development](#)
Khaled AbdElazim, Ramadan Moawad and Essam Elfakharany

Quantum Science and Technology



PAPER

Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber

OPEN ACCESS

RECEIVED
23 October 2023

REVISED
29 July 2024

ACCEPTED FOR PUBLICATION
28 August 2024

PUBLISHED
10 September 2024

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Francisco Ferreira da Silva^{1,2,3,4,*} , Guus Avis^{1,2,3,4} , Joshua A Slater¹  and Stephanie Wehner^{1,2,3,*} 

¹ QuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands

² Quantum Computer Science, EEMCS, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands

³ Kavli Institute of Nanoscience, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands

⁴ These authors contributed equally.

* Authors to whom any correspondence should be addressed.

E-mail: f.hortafferiradasilva@tudelft.nl and s.d.c.wehner@tudelft.nl

Keywords: quantum repeaters, quantum networks, entanglement distribution, quantum key distribution, blind quantum computing

Abstract

We perform a numerical study of the distribution of entanglement on a real-world fiber grid connecting the German cities of Bonn and Berlin. The connection is realized using a chain of processing-node quantum repeaters spanning roughly 900 kilometers. Their placement is constrained by the fiber grid we consider, resulting in asymmetric links. We investigate how minimal hardware requirements depend on the target application, as well as on the number of repeaters in the chain. We find that requirements for blind quantum computing are markedly different than those for quantum key distribution, with the required coherence time being around two and a half times larger for the former. Further, we observe a trade-off regarding how target secret-key rates are achieved when using different numbers of repeaters: comparatively low-quality entangled states generated at a high rate are preferred for higher numbers of repeaters, whereas comparatively high-quality states generated at a lower rate are favored for lower numbers of repeaters. To obtain our results we employ an extensive simulation framework implemented using NetSquid, a discrete-event simulator for quantum networks. These are combined with an optimization methodology based on genetic algorithms to determine minimal hardware requirements.

1. Introduction

The Quantum Internet promises to enable various novel applications that are provably impossible using the classical internet alone. These include the secure distribution of secret keys [1, 2], distributed quantum computation [3] and secret sharing [4]. The requirements on hardware for a functioning quantum internet are application-dependent [5], but the generation and distribution of high-quality entanglement is necessary for all applications beyond quantum key distribution (QKD).

Recent years have seen major advances in quantum networking, with state-of-the-art experiments demonstrating entanglement generation at metropolitan distances [6] and a three-node quantum network in a lab [7]. However, a large-scale quantum network requires long-distance entanglement generation, which has thus far not been realized. This is due to the fact that the probability of not losing photons, the typical travelling carriers of quantum information, decays exponentially with the distance traversed in optical fiber. Classically, signals are amplified to overcome loss. The same solution is impossible in the quantum case due to the no-cloning theorem [8–10], which renders copying arbitrary quantum states impossible.

Quantum repeaters [11, 12] have been proposed as a possible solution for dealing with photon loss by introducing intermediate nodes between the two distant locations over which entanglement is to be established. Multiple platforms have been considered for implementing quantum repeaters (see, e.g. [13, 14]), including ensemble-based systems [15, 16], trapped ions [17–19], neutral atoms [18, 20] and color



Figure 1. Map of Germany overlaid with a depiction of the fiber path connecting the German cities of Bonn and Berlin that we investigate, provided by Deutsche Telekom (DT). The white circles represent locations where DT plans to install trusted nodes and where, when building a repeater chain, processing nodes or heralding stations could be placed. These locations are connected to each other through fiber drawn in black. The maximum number of repeaters that can be placed between Bonn and Berlin in this fiber network is seven. We consider all possible repeater placements, assuming that the heralding stations are placed as symmetrically as possible (there are 986 such placements). The distance between Bonn and Berlin is 917.1 km via fiber, and approximately 480 km as the crow flies. The reason for such a large difference between the two values is that other major German cities, such as Hannover and Dortmund, are connected through the fiber link as well.

centers in diamond [21]. All but the first are examples of processing nodes. These are capable of not only storing and transmitting information, but also of performing quantum gates.

Despite recent progress, a scalable quantum repeater has yet to be demonstrated. Part of the challenge is that hardware requirements are not fully known. There is a large body of work investigating such requirements (see [22] and references therein). However, to the best of our knowledge, there has yet to be a study of how hardware requirements for chains of multiple repeaters, with their placement constrained by a real-world fiber grid, depend on (i) the number of repeaters used and (ii) the application for which the entanglement will be used.

It is likely that real-world deployment of quantum networks will make use of existing fiber infrastructure [23]. Previous work has shown that accounting for this fact significantly affects the hardware requirements for a single processing-node repeater setup [22]. This emphasizes the importance of taking the constraints imposed by existing fiber grids into consideration when estimating repeater hardware requirements. Moreover, existing fiber grids are already being used to deploy trusted-node networks throughout the world [24]. These networks allow for the distribution of keys, albeit without end-to-end security [25]. A natural development in the deployment of quantum networks is to upgrade trusted-node networks by replacing them with quantum-repeater networks [5]. This provides further motivation for considering real-world fiber grids when studying hardware requirements.

Given that future quantum networks will be used for different applications, it is also pertinent to ask whether the required hardware quality depends on the application to be executed. QKD is one of the best-known applications for quantum networks [1, 2]. In its entanglement-based version [26], two parties that share entangled pairs can perform measurements in two different bases and compare outcomes in order to distill a provably-secret key that can then be used for cryptographic protocols. Verifiable blind quantum computation (VBQC) is a promising application that allows a client to execute a computation on a powerful remote server securely and privately [27, 28]. QKD is the most often considered application when evaluating hardware requirements, but it is not clear whether its requirements are representative of those for other quantum-networking applications, particularly as QKD does not require multiple live entangled states at the same time (unlike, for example, VBQC).

In this work, we investigate minimal hardware requirements for connecting two end nodes which are separated by roughly 900 km of real-world optical fiber using a chain of processing-node quantum repeaters. We study how these requirements change with the number of repeaters used and the imposed performance target. In particular, we compare minimal hardware requirements for performing QKD at different rates (1, 10 and 100 Hz) and for performing VBQC.

1.1. Setup

In this section, we introduce the setup we consider, as well as all the hardware parameters for which we will determine requirements. We consider the quantum-network path depicted in figure 1, with two end nodes situated in Bonn and Berlin separated by 917.1 km of optical fiber corresponding to 214.7 dB of attenuation (at a telecom wavelength of 1550 nm). There are a total of sixteen locations between the end nodes where equipment can be placed, namely repeater nodes and heralding stations. Note that imposing this restriction implies that the links will be asymmetric in both length and attenuation. Throughout this paper we assume that such a heralding station must be placed between every pair of neighboring network nodes (i.e. end

nodes or repeater nodes), as these are required when entanglement is generated between those nodes through the interference and measurement of entangled photons [7, 22, 29–35]. This data has been provided to us by Deutsche Telekom (DT), Germany’s largest telecommunications provider, which plans to install trusted nodes in the locations depicted in figure 1.

We assume neighboring nodes perform heralded entanglement generation [36, 37]. That is, entanglement consists of a series of attempts, and at the end of each attempt the partaking nodes learn whether an entangled state was successfully created or not. Examples of protocols for heralded entanglement generation are the double-click protocol [30, 32–35, 38], where photons are interfered and measured at a heralding station and success is declared in case two detectors click, the single-click protocol [7, 29, 31, 39, 40], where photons are also interfered but success is only declared in case one detector clicks, and direct transmission of an entangled photon from one node to the next where it is stored in heralded quantum memory [41–43]. Here, we employ a simplified, more abstract model for heralded entanglement generation, which retains features of both the single-click and double-click models. We do this so that the protocol and its interplay with other components of the repeater chain can be readily understood and our modelling is not overly platform specific. First of all, we assume that each node can perform heralded entanglement generation with two neighbours in parallel, which is not currently possible for all quantum-repeater platforms [21, 22]. Second, we model the elementary-link states ρ that are created upon the completion of a successful attempt as depolarized Bell states, i.e.

$$\rho = W|\phi^+\rangle\langle\phi^+| + \frac{1-W}{4}\mathbb{1}, \quad (1)$$

where W is related to the fidelity F to the ideal Bell state $|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ as $F = (1 + 3W)/4$ and $\mathbb{1}$ is the four-dimensional identity matrix. We note that real entangled states generated in quantum-repeater chains are often not depolarized states [22, 44]. Yet, as depolarized Bell states represent a worst-case type of noise [45], using a depolarizing model ensures that we will not find hardware requirements that are artificially low due to this simplification. Third, we take the time t_{attempt} required to perform one attempt to be given by

$$t_{\text{attempt}} = \frac{L}{c}, \quad (2)$$

where L is the fiber distance between the two nodes and $c = 200\,000 \text{ km s}^{-1}$ is the speed of light in fiber. That is, it corresponds to the communication time associated with sending photons to a heralding station that is exactly in the center between two nodes and then receiving a message with the measurement outcome. This is equivalent to the time required to directly transmit a photon from one node to the next. In reality it may be longer, as the attempt time could be further limited by the rate of the photon source, local operations or the synchronization of emission times [7, 46]. Finally, we take the success probability p_{el} of each attempt to be

$$p_{\text{el}} = p_{\text{det}} \times 10^{-\frac{\alpha_{\text{att}}}{10}L}. \quad (3)$$

Here, p_{det} is the probability that an emitted photon that is led through fiber to a detector is detected, given that it is not lost while travelling in fiber. This parameter combines multiple sources of loss, such as the detector’s efficiency, the probability of emitting the photon in the right mode and the probability of successfully sending the photon into the fiber, but not the fiber’s attenuation losses. α_{att} is the fiber’s attenuation coefficient (in dB km^{-1}). Therefore, the success probability corresponds to the success probability of directly transmitting a photon between the nodes and measuring it there. We note that for the double-click protocol the scaling with length would be the same, although the prefactor would be different (p_{det}^2 instead of p_{det} , as two photons must be detected). For the single-click protocol the scaling would be more gentle (roughly replacing L by $L/2$ in the exponential), and while the prefactor would be linear in p_{det} , there would also be a factor that depends on the device settings (specifically on the bright-state parameters chosen at both nodes, which tune a trade-off between success probability and state fidelity [7, 22, 47]). Additionally, we allow also for the possibility of multiplexed heralded entanglement generation [48–50]. This essentially consists of performing multiple attempts of generating the same elementary-link state in parallel. Multiplexing can be done across multiple degrees of freedom, such as frequency, time or space. We remain agnostic regarding how the multiplexing is performed, including it in our model as one parameter corresponding to the number of multiplexing modes used, n . The probability of successfully generating an elementary link assuming the use of multiplexing is then the probability that at least one of the multiplexing modes succeeds:

$$p_{\text{multiple modes}} = 1 - (1 - p_{\text{el}})^n. \quad (4)$$

The nodes implement a swap-asap protocol [51, 52]. That is, as soon as a node holds two entangled states, one shared with each of its neighbours, it performs an entanglement swap in order to create an entangled state spanning a larger distance. We assume this swap is realized deterministically, since we are modelling processing nodes that can implement a swap using quantum gates and measurements on their processors. It may however introduce noise, which we model as depolarizing. We quantify how well the swap can be performed using the swap-quality parameter s_q . The d -dimensional depolarizing noise channel of parameter p acts on a state ρ as follows,

$$\rho \rightarrow p\rho + (1-p)\frac{\mathbb{1}}{d}. \quad (5)$$

This means that, with probability p , ρ is left unchanged, and with probability $1-p$ it is mapped to the maximally-mixed state, i.e. all information is lost. Then, we model entanglement swapping as a two-qubit depolarizing channel (i.e. $d=4$) with parameter $p=s_q$ followed by a perfect entanglement-swapping operation (i.e. a measurement in the Bell basis [53]). We assume that the gates and measurements applied by the end nodes when executing QKD and VBQC are noiseless and instantaneous. States stored in memory undergo decoherence, which we model as exponential depolarizing noise, i.e.

$$\rho \rightarrow e^{-t/T}\rho + (1-e^{-t/T})\frac{\mathbb{1}}{d}, \quad (6)$$

where t is the time for which the state ρ has been held in memory and T is the memory's coherence time. To combat the effects of memory decoherence, entangled states are discarded after a local cut-off time. The cut-off time is defined as follows: a timer starts once a state is created in memory through the successful generation of an elementary link. If the timer reaches the local cut-off time, the state is discarded. That is, the qubit holding the state is reset. Additionally, the node sends a classical message along the chain so that the qubit with which the first qubit was entangled can also be reset. As a result, a number of elementary links in the chain must be regenerated (with the exact number depending on how far away the entangled qubit was).

1.2. Applications

Having discussed our modelling of the entanglement generation process between Bonn and Berlin, we turn to the applications that will make use of the entanglement, QKD and VBQC. We investigate the BB84 QKD protocol [2] (in its entanglement-based form [26]) between the end nodes situated in Bonn and Berlin. We record the entanglement generation rate and estimate the quantum bit error rate (QBER) that would have been obtained when measuring the generated state in order to estimate the achievable asymptotic secret-key rate (SKR) as per the following equation [54]:

$$\text{SKR} = E_R \cdot \max(0, (1 - 2H(Q))), \quad (7)$$

where E_R is the entanglement-generation rate, $H(p) = -p\log_2(p) - (1-p)\log_2(1-p)$ is the binary entropy function and Q is the QBER. We note that all the noise sources we consider are depolarizing, hence the entangled states generated will be of the form of the state shown in equation (5). Therefore, the QBER is the same irrespective of the measurement basis. The end nodes do not wait until end-to-end entanglement is established before measuring their qubits. Instead, they measure them as soon as they have established entanglement with their nearest neighbours, as this minimizes the amount of time states spend in memory, resulting in laxer hardware requirements.

We note that while it is possible to perform QKD by measuring entangled states, it is also possible to perform QKD without entanglement distribution (e.g. using measurement-device-independent QKD [55] or twin-field QKD [56]). However, such schemes cannot be used for any other quantum-network applications, such as quantum teleportation or VBQC, and hence we do not consider them any further.

We also investigate a two-qubit version of the VBQC protocol introduced in [28]. In such protocols, a client wishes to delegate a computation to a powerful remote server in a secure and verifiable fashion [27]. In particular, we investigate the repeated execution of test rounds of the protocol, which consist of the server performing a controlled-Z gate followed by a measurement. In these rounds the client knows the computation's expected outcome, and can therefore compare them to the observed outcomes. Under the assumption of an honest server, wrong outcomes are a result of noise. We call this the *BQC test protocol*. The fraction of successful BQC test protocol rounds is therefore a metric for the quality of the entanglement used for transmitting qubits. We define the success rate as the number of rounds of the protocol that can be executed with a successful result per time unit. More concretely, if p_s is the success probability of a test round and R_{rounds} is the rate at which rounds can be executed, the BQC-test-protocol success rate is given by:

$$R_{\text{BQC}} = R_{\text{rounds}} \cdot p_s. \quad (8)$$

While the BQC test protocol is in and of itself not an interesting application of a quantum network, it can be considered a benchmark for how well the network is suited to VBQC and possibly other multi-qubit applications. The fact that, in contrast with QKD, it requires the distribution of multiple entangled states and the storage of qubits between rounds makes it a more meaningful benchmark for quantum-network applications that require multiple live qubits contemporaneously. Further details on the BQC test protocol can be found in appendix D.

The two applications we have just introduced define our performance targets.

1.3. Minimal hardware requirements

We wish to find the *minimal hardware requirements* that are needed to realize different target SKRs and BQC-test-protocol success rates. These correspond to the minimal improvements over state-of-the-art hardware parameters that enable meeting the targets. We phrase the problem of finding minimal hardware requirements as a constrained optimization problem. Namely, we wish to minimize the hardware improvement while ensuring that the constraint induced by the performance target is met. This constraint is relaxed through a process known as scalarization [57, 58], resulting in a single-objective optimization problem, in which the quantity to be minimized is the sum of the cost associated to the hardware improvement and a penalty term for the rate target. The resulting cost function is given by:

$$C = w_1 \left(1 + (R_{\text{target}} - R_{\text{real}})\right)^2 \cdot \Theta(R_{\text{target}} - R_{\text{real}}) + w_2 H_C(x_1, \dots, x_N), \quad (9)$$

where H_C is the hardware improvement cost associated to parameter set $\{x_1, \dots, x_N\}$, w_i are the weights assigned to the objectives, Θ is the Heaviside step function, R_{target} is the rate target and R_{real} is the rate of application execution achieved by the parameter set. We note that R_{real} and R_{target} can be either a SKR or a BQC-test-protocol success rate. H_C maps sets of hardware parameters to a number, the cost, which represents how large of an improvement over the state-of-the-art they represent. In order to compute this cost consistently across different parameters, we use no-imperfection probabilities as done in [22]. By no-imperfection probability, we mean the probability that there is no error or loss associated to a given parameter. For example, the no-error probability associated to a photon detection probability p_{det} (defined in section 1.1) of 0.1 is 0.1. For the no-error probabilities associated to the other hardware parameters, see appendix C.1. We say that a parameter is improved by a factor of k if its no-imperfection probability becomes $\sqrt[k]{p_{ni}}$, with p_{ni} being the state-of-the-art no-imperfection probability. For example, improving the no-imperfection probability of 0.1 associated to $p_{\text{det}} = 0.1$ by a factor of 5, we get a no-imperfection probability of ≈ 0.63 , corresponding to $p_{\text{det}} \approx 0.63$. The hardware cost associated to a set of parameters is given by the sum of the improvement factors of the parameters. The weights w_i are chosen such that the term of the overall cost function corresponding to meeting the rate target is always larger than the one corresponding to the hardware cost, ensuring that even though we have relaxed the constraints by scalarizing, we are still effectively requiring that the minimal hardware requirements are such that the performance target is met. To ensure this, we picked $w_1, w_2 \gg w_3$, such that $w_1 (1 + (R_{\text{target}} - R_{\text{real}})^2) \Theta(R_{\text{target}} - R_{\text{real}}) \gg w_2 H_C(x_1, \dots, x_N)$. Specifically, we set $w_1 = 1 \times 10^{100}$ and $w_2 = 1$. No particular heuristic was used to select these numbers.

We note that the hardware cost is meant only to represent a measure of the hardness of improving the hardware to a certain level, and not any form of monetary cost. At present quantum repeater systems are research setups, with commercial solutions only starting to emerge. Therefore, assigning any specific commercial cost numbers would be too speculative at this point, and would require an in-depth study outside the scope of this project.

1.4. State-of-the-art parameters

Computing minimal hardware requirements as described in section 1.3 is done with respect to a baseline over which we are improving. In this work, this baseline consists of parameters measured for color centers in diamond, as they are physical systems using which various quantum-networking primitives have been demonstrated. These include long-lived quantum memories [59], remote entanglement generation [32, 33], quantum teleportation [38], entanglement distillation [39], entanglement swapping [60], a three-node network [7] and memory-enhanced quantum communication [41]. The result of such extensive experimentation is that these devices have been thoroughly characterized, rendering finding hardware parameters to use as baseline easy. We do not impose that all parameters must have been demonstrated in the same experiment or even with the same color center. The parameters we consider are shown in table 1. Details on how these parameters were determined can be found in appendix A. Note that even though these parameter values are based on quantum-networking experiments using color centers, our modelling remains

Table 1. State-of-the-art color-center parameters. We note that not all of these parameter values have been realized in a single experiment. We have number of modes as 1 without reference because to the best of our knowledge multiplexed entanglement generation has not been demonstrated using color centers. These parameters are introduced in section 1.1.

Parameter	Value
Coherence time	1 s [61]
Number of multiplexing modes	1
Fidelity of elementary links	0.83 [33]
Photon detection probability	0.255 [41]
Swap quality	0.83 [39, 62]

hardware-agnostic. In fact, one could consider parameter values derived from trapped-ion or neutral-atom experiments while still using the same modelling and simulations.

1.5. Determining minimal hardware requirements

In order to determine minimal hardware requirements, we need to (i) be able to evaluate how a given set of hardware parameters performs and (ii) optimize over the parameter space to find the parameters that minimize the requirements while still performing adequately (i.e. the parameters that minimize the cost function defined in equation (9)).

We evaluate the performance of hardware parameters using general processing-node repeater-chain simulations developed in NetSquid, a discrete-event based quantum-network simulator [51]. The simulations are general in the sense that they can be used to investigate swap-asap repeater chains of arbitrary size and spacing (i.e. nodes and heralding stations need not be equidistant). They take into account time-dependent noise, classical control communication and the constraints imposed by a real-world fiber network. The code for executing such simulations has been made publicly available at [63] and is largely based on the simulations first introduced in [22]. Our code that utilizes these simulations to produce the results here presented can be found at [64] (and the corresponding data at [65]).

Given that we can evaluate the performance of any parameter set on the Bonn-Berlin path, we perform parameter optimization using a genetic algorithm [66] to minimize the cost function defined in section 1.3 using the high-performance computing cluster Snellius. For further details, see appendix C.

2. Impact of number of repeaters on hardware requirements

In this section we answer the question of how hardware requirements are affected by the number of repeaters deployed in a quantum network. Specifically, we investigate the minimal hardware requirements for performing BB84 between the German cities of Bonn and Berlin at a key rate of 10 Hz. We assume the cities are connected by the network path shown in figure 1. We determine what these minimal requirements are in two cases: (i) optimizing over the number of repeaters and (ii) restricting the number of repeaters to specific values. In both cases we optimize over the placement of the repeaters.

2.1. Absolute minimal number of multiplexing modes

Before determining minimal requirements, we aim to answer the question of what are the absolute minimal number of multiplexing modes required to perform QKD between the German cities of Bonn and Berlin at rates of 1, 10 and 100 Hz. By *absolute minimal* number of multiplexing modes, we mean the minimum number of multiplexing modes that is required if the only source of imperfection in the setup is fiber attenuation. This provides a lower bound on the number of multiplexing modes in the minimal hardware requirements, as the introduction of other hardware imperfections can only lead to more stringent demands on the number of modes. We emphasize that for the purposes of answering this question we are temporarily setting aside the real-world fiber path introduced in figure 1. Instead, we are going to consider a *symmetrized* version of that path. By this we mean a path with the same total length and attenuation, but in which nodes and heralding stations are placed equidistantly, and where the attenuation is evenly distributed throughout the path, i.e. all elementary links have the same attenuation. The reason for doing so is that the minimal number of modes for this path is a lower bound for the same quantity on any other path with the same total length and attenuation. To see this, we note that it has been shown in [67] that repeater chains of the type studied here perform best when all nodes are positioned as symmetrically as possible. This implies that such a chain will have less stringent hardware requirements to attain a given performance target in comparison to chains which are subject to real-world restrictions such as the ones imposed by the fiber path shown in figure 1, and, therefore, also less stringent requirements on the number of multiplexing modes.

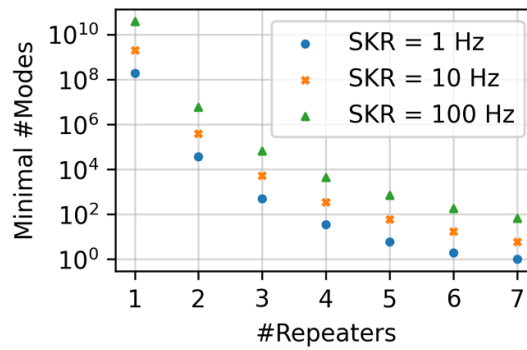


Figure 2. Minimal number of multiplexing modes required to achieve 1, 10 and 100 Hz of SKR over 917.1 km of fiber with a total of 214.7 dB of attenuation, corresponding to a symmetrized version of the path between Bonn and Berlin that we investigate. That is, for N repeaters, the symmetrized path has $N + 1$ elementary links, each of length $917.1/(N + 1)$ km and of attenuation $214.7/(N + 1)$ dB. We assume that there are no hardware imperfections, and that repeaters are uniformly spaced.

Determining the absolute minimal number of multiplexing modes serves two purposes. First, it allows us to limit the search space of the optimization we run for finding minimal hardware requirements. Second, it gives us a general idea of how many repeaters might be required to achieve the target with reasonable hardware demands. For example, if for a specific number of repeaters hundreds of thousands of multiplexing modes are required to meet the target without any noise sources, that may indicate that using that number of repeaters is not practically feasible.

In figure 2 we show the absolute minimal number of modes required to distribute secret key at rates of 1, 10 and 100 Hz using BB84 in the symmetrized Bonn—Berlin path. We find that more multiplexing modes are required for higher rate targets, and that this number grows superexponentially as the number of repeaters decreases, so as to counteract the effects of photon loss in fiber. Further, we find that achieving a SKR of 10 Hz with fewer than 3 repeaters requires hundreds of thousands of multiplexing modes even in the absence of any sources of noise. As the hardware cost (defined in section 1.3) associated with so many multiplexing modes far outweighs typical values for the minimal total hardware cost found for three or more repeaters we limit the rest of our investigation to configurations with three or more repeaters.

2.2. Minimal hardware requirements for QKD

We now turn our attention to the minimal hardware requirements for performing QKD at a rate of 10 Hz using the BB84 protocol. In particular, we investigate them along the path connecting Bonn and Berlin depicted in figure 1. As figure 2 illustrates, the number of repeaters used can have a considerable impact on the hardware requirements. Further, it is expected that the same is true for the placement of repeaters and heralding stations [22, 67]. With this in mind, we ask two questions: (i) what are the minimal hardware requirements when allowing for the placement of up to the largest number of repeaters that fits in the fiber path (seven) and (ii) what are the minimal hardware requirements when restricting the maximum number of repeaters to five. We expect that this will lead to different parameter regimes, illustrating two possible directions towards achieving the target performance.

In figure 3 we show the directions along which hardware must be improved for distributing secret key at rates of 10 Hz using BB84 in the network path connecting Bonn and Berlin. The corresponding minimal hardware requirements can be found in table 2. In each case we find that the hardware requirements are minimized when the number of repeaters used is maximized. That is, for seven repeaters in case (i) and five repeaters in case (ii). Hardware requirements are more stringent in case fewer repeaters are used. In particular, the overall photon detection probability excluding attenuation in fiber must be improved to a much larger degree (0.79 vs 0.36) if only five repeaters are used. This is needed to overcome the increased attenuation losses associated with the longer elementary links. The coherence time required when using five repeaters is also larger than the time required when using seven repeaters (4.2 s vs 3.1 s). This can be explained by the fact that keeping the entanglement-generation rate high is more costly in case of five repeaters. Therefore keeping the QBER small to extract as many secret bits as possible from each entangled state is more valuable. Furthermore, since the entanglement-generation rate is smaller for five repeaters, qubits are stored for longer times before they can be swapped and hence a larger coherence time is required to achieve the same QBER. We study this interplay further in section 2.3. Finally, we notice that while the requirements on most hardware parameters are more stringent for five repeaters as compared to seven repeaters, this is not the case for the requirement on the swap quality. In fact, the requirement on the swap quality is even slightly looser for five repeaters (0.996 vs 0.997). This is explained by the fact that when there

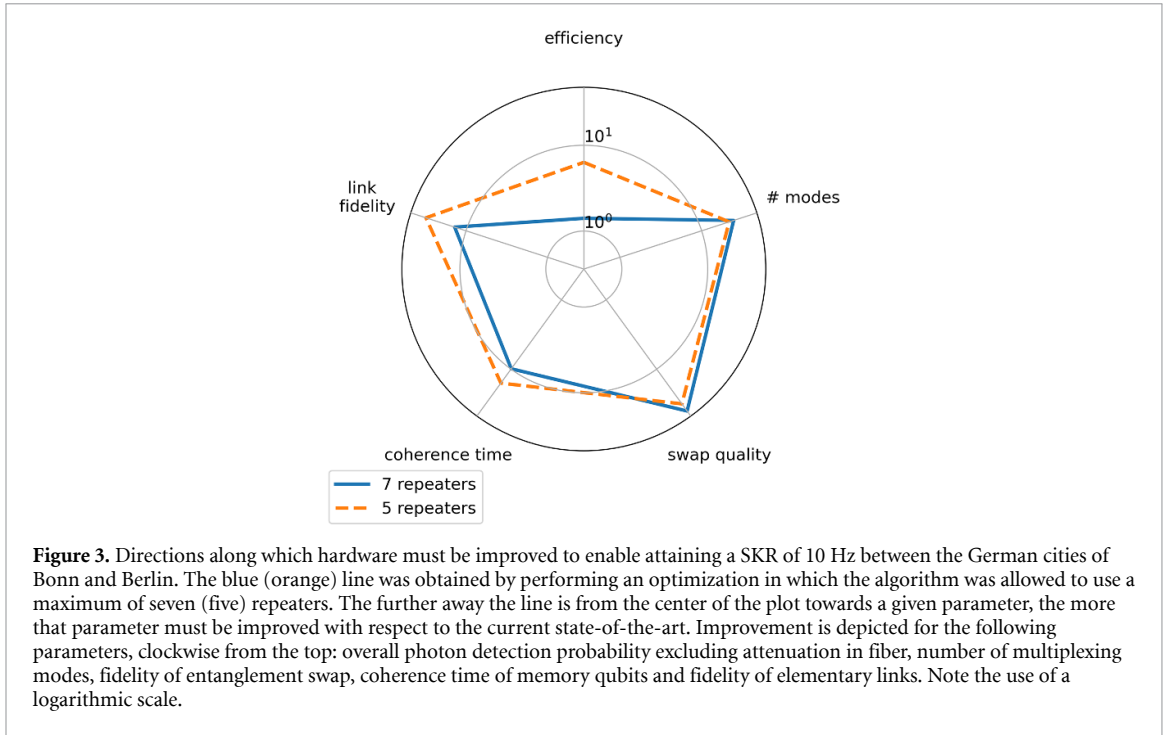


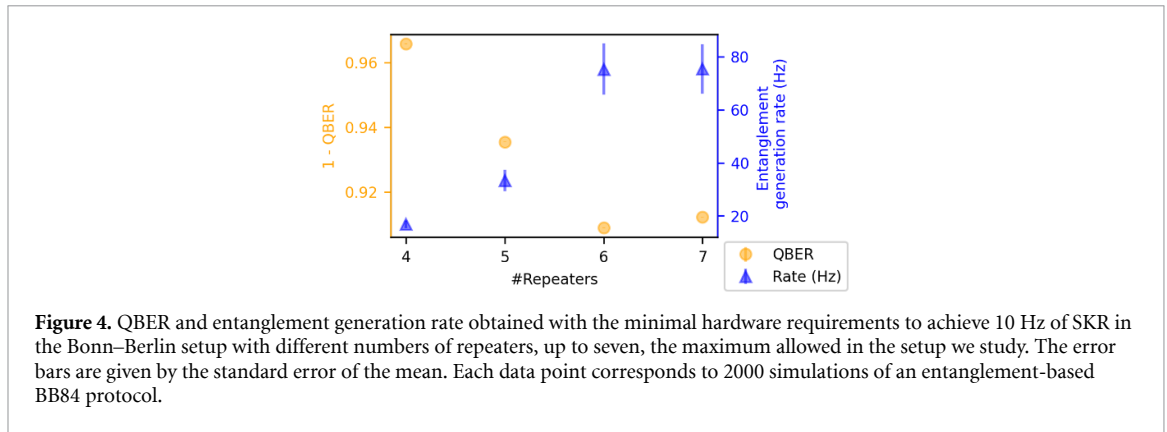
Table 2. Minimal hardware requirements to achieve 1, 10 and 100 Hz of secret-key rate and 10 Hz of blind quantum computing test protocol success rate between the German cities of Bonn and Berlin. The photon detection probability p_{det} is the probability of a photon being detected given that it is not lost in fiber. It combines multiple sources of loss, such as the detector's efficiency, the probability of emitting the photon in the right mode and the probability of successfully sending the photon into the fiber. More details can be found in section 1.1 (where all parameters are introduced) and appendix A.

Application	QKD			BQC	
Rate (Hz)	1	10	100	10	
Number of repeaters	7	Max 5	Max 7	7	
Coherence time (s)	1.81	4.23	3.14	10.1	7.99
Number of multiplexing modes	175	544	592	799	172
Fidelity of elementary links	0.989	0.995	0.987	0.996	0.845
Photon detection probability p_{det}	0.604	0.785	0.360	0.804	0.552
Swap quality	0.996	0.996	0.997	0.997	0.881

are more repeaters, there are more entanglement swaps associated with every end-to-end entangled state. Therefore, when there are more repeaters the final error rate is more sensitive to noise in the swaps, creating a larger incentive to improve the associated parameter in the seven-repeater case as compared to the five-repeater case.

2.3. SKR: QBER and entanglement generation rate

A specific value for the SKR can be obtained through many different pairs of values for the entanglement-generation rate and the QBER, as follows from equation (7). This opens up a trade-off between the entanglement generation rate and the QBER, as briefly discussed in section 2.2. Here, we investigate this trade-off more deeply by repeating our process for determining minimal hardware requirements to achieve an SKR of 10 Hz while keeping the number of repeaters a fixed parameter. We did this for 4, 5, 6 and 7 repeaters. For each case, we still optimize over all possible placements of the repeaters in the fiber grid. In figure 4 we show the QBER and entanglement-generation rate achieved with the minimal hardware requirements for the best setup found by our optimization procedure for varying number of repeaters. We observe two different regimes. For 4 and 5 repeaters, which we name the ‘few-repeater’ regime, we find a low QBER ($\sim 5\%$) and an entanglement-generation rate of 20–30 Hz. On the other hand, for 6 and 7 repeaters, i.e. the ‘many-repeater’ regime, we find a comparatively higher QBER ($\sim 9\%$) and an entanglement generation rate of almost 80 Hz. In other words, in the many-repeater regime, distributing many entangled pairs of comparatively lower quality requires less hardware improvement. On the other hand, in the few-repeater regime it seems to be more feasible to distribute fewer pairs of comparatively higher quality. As the number of repeaters used decreases, it becomes harder to overcome the effect of fiber attenuation, which makes improving the quality of the entangled states delivered a more attractive option for increasing the SKR.



We finalize by remarking that, perhaps surprisingly, the variance in the time it takes to distribute one entangled state appears to grow as the number of repeaters in the chain increases (as shown by the increasing error bar on the rate in figure 4). While interesting, further investigation is beyond the scope of this work.

We show the repeater placement corresponding to the minimal hardware requirements found when optimizing over the number of repeaters and their placement in appendix B.

3. Impact of target on hardware requirements

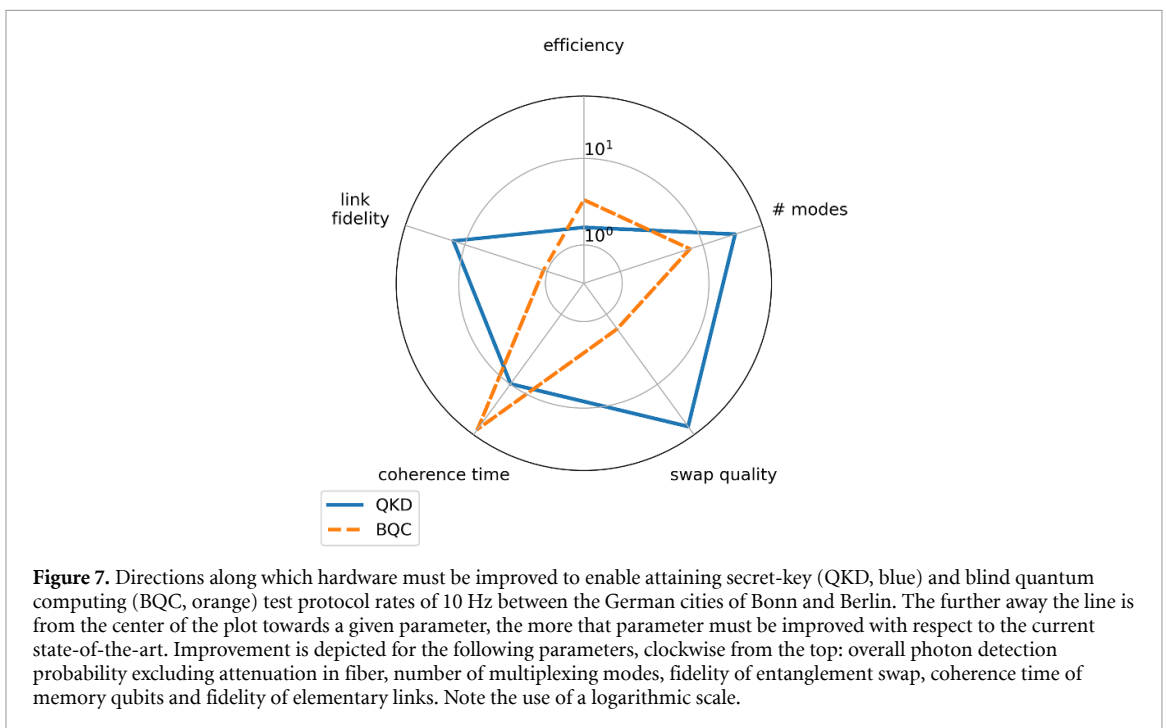
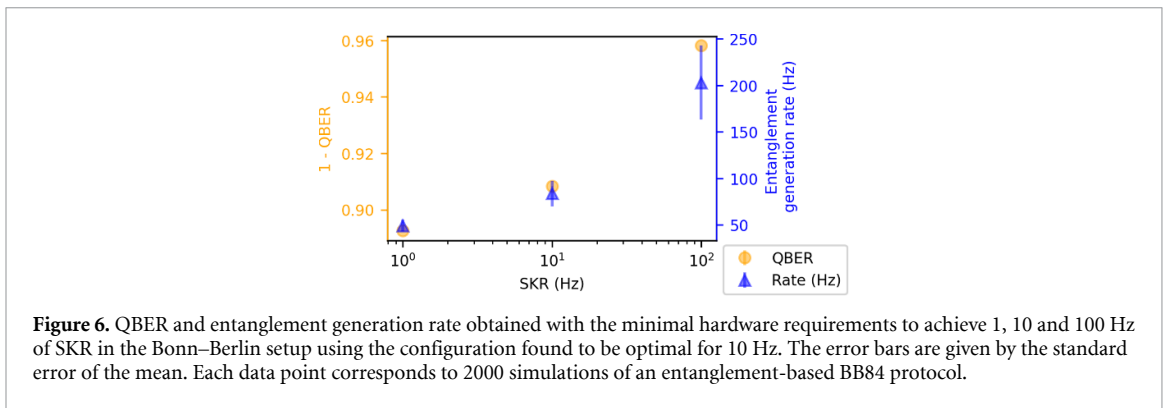
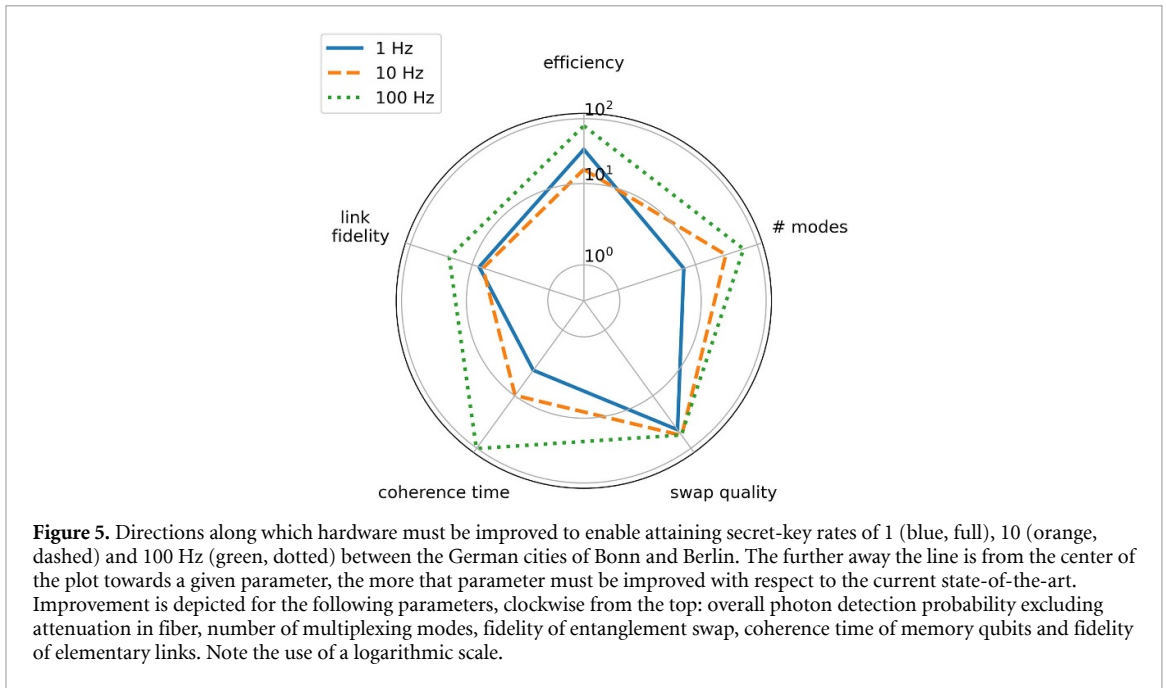
We now turn our attention to the impact of the performance target on the hardware requirements. We approach this from two angles: (i) the impact of varying the SKR target and (ii) the impact of holding the required rate constant but changing the target quantum-network application. It is clear that, given the same repeater chain, increasing the target rate will lead to more stringent requirements. However, it is not *a priori* obvious if the relative importance of the different parameters will change as the target rate is increased. It is further not obvious how changing the target application impacts the hardware requirements. These are questions of practical relevance: given that one wishes to build a repeater chain capable of distributing entanglement to perform QKD at a rate of 100 Hz, it seems crucial to know whether building a repeater chain for performing QKD at a rate of 1 Hz is a step in the right direction. In other words, this investigation can shed light on whether the process of improving hardware for quantum-repeater chains should be approached incrementally. The same question holds for the different target applications. It is likely that quantum repeaters will initially be used for QKD as they begin to replace their trusted-node predecessors, and only progressively start to be used for applications that require multiple live qubits. We would then like to know whether the hardware improvements necessary to perform QKD using quantum repeaters are similar to the ones for multi-qubit applications.

3.1. Requirements for different SKR targets

In figure 5 we show the directions along which hardware must be improved for distributing secret key at rates of 1, 10 and 100 Hz using BB84 in the network path connecting Bonn and Berlin. The corresponding minimal hardware requirements can be found in table 2. The hardware requirements become more stringent as the SKR target grows. Further, the coherence time requires significantly less improvement in the 1 Hz case when compared to 10 and 100 Hz. This comes as something of a surprise, given that we expect qubits to spend less time in memory for higher SKR values, as these should correspond to higher entanglement-generation rates (and hence lower waiting times). In order to further investigate why this happens, we show in figure 6 the QBER and entanglement generation rate achieved with the minimal hardware requirements for the best setup found by our optimization procedure for different SKR targets. We find that both the entanglement generation rate and $1 - \text{QBER}$ increase with the target SKR. We conjecture that the increase in coherence time observed for higher SKR targets is due to the necessary entanglement generation rate being very high. In fact, it is so high that it requires a huge number of multiplexing modes, which in turn imply a very high cost. This makes it comparatively less costly to extract more key from each entangled state than to generate states faster. Finally, we remark that for each of the targets the optimization solutions exceed the repeaterless bound for QKD by a large margin, as shown in appendix E.

3.2. Requirements for secret-key and BQC success rates

In figure 7 we show the directions along which hardware must be improved for performing QKD and BQC at a rate of 10 Hz in the network path connecting Bonn and Berlin. The corresponding minimal hardware requirements can be found in table 2. It is plain to see that the two applications require improvements in



distinct parameters. In particular, we emphasize the much larger coherence time required for BQC, corresponding to roughly a factor of 2.5 difference (7.99 vs 3.14 seconds). This can be explained by the fact that BQC, unlike QKD, requires two entangled pairs to be alive at the same time, implying that one entangled pair must be stored at the end nodes while the second one is generated. Further, the fact that the minimal coherence time required for BQC is high means that comparatively less noise will be caused by decoherence. This, in turn, means that in order to achieve the same state quality, the swap quality and the elementary link fidelity need not be as good.

We have also observed that there is a significant difference in the entanglement generation rates achieved by the parameter sets corresponding to the improvements shown in figure 7. The minimal hardware requirements for QKD achieve an entanglement generation rate of almost 80 Hz, whereas the ones for the BQC-test-protocol result in an entanglement generation rate of around 20 Hz. In the same vein as what was discussed in section 2.3, this is a result of the SKR and the BQC test protocol success rate being composite quantities, depending not only on the rate at which entangled states are delivered, but also on the quality of these states. We believe that the difference observed in entanglement generation rate between the two applications is due to the fact that there is a threshold state quality to obtain non-zero secret-key ($\sim 11\%$ QBER or equivalently ~ 0.84 fidelity, both under the assumption of depolarizing noise). Such a threshold does not exist for the BQC test protocol. This fundamental difference means that the state quality requirements are more stringent in the QKD case, making improving the entanglement generation rate a more attractive possibility. We do however note that even though the BQC test protocol does not impose a threshold on state quality, the complete VBQC protocol proposed in [28] does.

4. Conclusion

We have determined minimal hardware requirements for generating entanglement between two nodes separated by roughly 900 km of real-world optical fiber using a chain of processing-node quantum repeaters. We investigated both how such requirements depend on how many repeaters are employed and on the quantum-network application for which the entanglement is used. Notably, we have found that the hardware requirements for performing QKD and a simplified form of BQC are qualitatively different, with BQC requiring a coherence time which is roughly a factor of 2.5 larger for the same target rate in the setup we investigated. We further observed that given that most metrics one is interested in when evaluating quantum-network performance depend on both the rate at which entanglement is generated and its quality, there is room for trade-offs: for example, we found that when employing a large number of repeaters to achieve a given SKR in the setup we studied it is easier to generate many entangled pairs of comparatively lower quality, with the opposite being true if fewer repeaters are used.

The BQC requirements we determined were obtained for a simplified form of the protocol, which is useful as a benchmark for quantum-network performance but is not an interesting application in and of itself. It would be interesting to learn how the results presented would change if instead a complete VBQC protocol such as the one introduced in [28] were studied. Doing so would require optimizing also over protocol parameters, by adaptively choosing the proportion of test and computation rounds depending on the computation being performed, the accepted tolerance and the network's noise parameters. We speculate that considering the full-on VBQC protocol would push the hardware parameters to enable the network to generate high-fidelity entangle states. This is because the protocol imposes an upper bound on the fraction of test rounds that can fail (this depends on the type of computation that is performed, but is 25% in the simplest case [28]).

We have not considered the possibility of employing entanglement distillation [11, 68]. The effect this would have on the hardware requirements is not obvious, since it would likely make some parameters, such as the two-qubit gate fidelity or the coherence time, more stringent, and others, such as the fidelity of elementary links, less stringent. An interesting direction of future study would be to investigate in what scenarios distillation could lead to a faster realization of quantum-repeater-enabled entanglement distribution. The reason we have here shied away from considering entanglement distillation is that repeaters that can rely only on qubit storage and entanglement swapping are closer to the state-of-the-art [7, 35, 69], hence making them more appropriate for a study on requirements for upgrading trusted nodes. Furthermore, finding good distillation protocols in repeater chains with noisy gates, limited memories, and asymmetric node placement is highly non-trivial, and an area of active study [70–72].

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.4121/22193539.v1>.

Acknowledgments

We thank Bethany Davies for critical reading of the manuscript. We thank Deutsche Telekom for sharing data regarding their fiber network. This work was supported by the QIA-Phase 1 project which has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101102140. G A was supported by NWO Zwaartekracht QSC 024.003.037.

Code availability

The code that was used to perform the simulations and generate the plots in this paper has been made available at <https://gitlab.com/softwarequtech/simulation-code-for-requirements-for-upgrading-trusted-nodes-to-a-repeater-chain-over-900-km-of-optical-fiber> [64].

Appendix A. Baseline parameters

Here we discuss how we determined the baseline hardware parameters shown in table 1. We did so by following two steps: (i) finding state-of-the-art color-center hardware parameters in the literature and (ii) converting these to the hardware model we employ. In table 3 we show the relevant state-of-the-art color center parameters we have identified and provide their respective references. We now discuss how these are converted to the parameters shown in table 1. The elementary-link fidelity and number of modes can be used directly without conversion. Color-center memories have both an electron qubit (also known as communication qubits due to their optical interface) and possibly multiple carbon qubits (also known as memory qubits due to being long-lived). We assume a 'best-of-both-worlds' situation, in which the qubits in our model are both endowed with an optical interface that allows them to generate entanglement and a long (1s baseline) memory lifetime. This simplification allows us to treat all qubits in the nodes equally. As explained in section 1.1 we combine all photon-related inefficiencies, with the exception of fiber attenuation, into one parameter, p_{det} . This is done as follows:

$$p_{\text{det}} = p_{\text{photon interface}} \cdot p_{\text{conv}}, \quad (\text{A1})$$

where $p_{\text{photon interface}}$ is the photonic interface efficiency and p_{conv} is the frequency-conversion efficiency. This results in the 0.255 number reported in table 1. We note that the experiment reported in [41] does not consist of entanglement generation through a heralding station, as we assume in this paper. We have made a best guess of how the parameters reported there would translate to a scheme where entangled photons are interfered and measured at a heralding station. An entanglement swap in a color center (this concrete example was demonstrated using a nitrogen-vacancy center) consists of single-qubit gates on both carbon and electron, two-qubit gates and measurement and initialization of the electron (see figure 17 in supplementary note 5 of [51] for an image of the circuit). We make the simplifying assumption that all errors are depolarizing. First, we convert each of the initialization and gate fidelities in table 3 to depolarizing parameters (in accordance with equation (5)), and then multiply the depolarizing parameters corresponding to all the operations in the circuit together to obtain the swap quality (which parametrizes a depolarizing channel as detailed in section 1.1), i.e.

$$s_q = (1 - p_{\text{carbon}})^2 \cdot (1 - p_{\text{electron-carbon}}) \cdot (1 - p_{\text{electron}})^2 \cdot (1 - p_{\text{electron init}}) \cdot (1 - p_{\text{electron meas}})^2 \cdot (1 - p_{\text{retrieve}}), \quad (\text{A2})$$

where p_{carbon} is the depolarizing parameter of the carbon single-qubit gate, $p_{\text{electron-carbon}}$ of the two-qubit gate, p_{electron} of the electron single-qubit gate, $p_{\text{electron init}}$ of the electron initialization, $p_{\text{electron meas}}$ of the electron measurement and p_{retrieve} of the retrieve operation (maps the carbon state to the electron, see figure 17(b) in supplementary note 5 of [51]).

Table 3. State-of-the-art color center parameters. We have number of modes as 1 without reference because to the best of our knowledge multiplexed entanglement generation has not been demonstrated using color centers.

Parameter	State-of-the-art value
Number of modes	1
Carbon coherence time	1 s [61]
Elementary-link fidelity	0.83 [33]
Electron initialization fidelity	0.995 [7]
Carbon initialization fidelity	0.99 [61]
Electron-carbon two-qubit gate fidelity	0.97 [39]
Electron single-qubit gate fidelity	0.995 [7]
Carbon single-qubit gate fidelity	0.999 [62]
Electron readout fidelity	0.93(0) 0.995(1) [60]
Photonic interface efficiency	0.855 [41]
Frequency conversion efficiency	0.3 [73]



Figure 8. Map of Germany overlaid with a depiction of the fiber path connecting the German cities of Bonn and Berlin that we investigated. The white circles represent end nodes, in Bonn and Berlin, and repeater nodes elsewhere. This placement corresponds to the best found by our optimization method, in the sense that it allowed for minimization of hardware requirements for a target secret-key rate of 10 Hz.

Table 4. Length and attenuation of elementary links depicted in figure 8.

Link	Length (km)	Attenuation (dB)
Bonn–Wuppertal	138.9	32.8
Wuppertal–Münster	133.2	31.4
Münster–Warmsen	126.2	29.6
Warmsen–Hannover	97.2	22.7
Hannover–Liebenburg	122.0	28.4
Liebenburg–Magdeburg	115.5	26.9
Magdeburg–Havel	103.9	24.3
Havel–Berlin	80.2	18.6

Appendix B. Repeater placement chosen by optimization method

As described in section 2.2, we determined minimal hardware requirements for performing QKD at a rate of 10 Hz over the network path depicted in figure 1. In doing so, we optimized over the number of repeaters used and their placement. We then used the placement our optimization method found to perform best for determining minimal hardware requirements for other performance targets, as described in section 3. In figure 8 we show this placement. In table 4 we show the lengths and attenuations of the elementary links defined by the repeater placement.

In order to optimize over the number of repeaters and their placement, we have first generated all the 986 possible ways repeaters can be placed in the network (such that there is still space for the required heralding stations between repeaters and end nodes). To each configuration we assigned a number r corresponding to the number of repeaters in the network. Then, for each configuration we computed the chain asymmetry parameter defined as [67]

$$\mathcal{A}_{\text{chain}} = \frac{1}{r} \sum_{i=1}^r \frac{|L_{\text{left},i} - L_{\text{right},i}|}{L_{\text{left},i} + L_{\text{right},i}}, \quad (\text{B1})$$

where $L_{\text{left},i}$ ($L_{\text{right},i}$) is the distance between repeater node i and its left- (right-) hand neighboring node. Next, we ordered all the configurations with the same value of r by their values of $\mathcal{A}_{\text{chain}}$, and label their position in this ordering as n . This number is then an identifier for how asymmetric (as quantified by the chain asymmetry parameter) a configuration is relative to the other configurations with the same number of repeaters. $n = 0$ corresponds to the most symmetric setup and $n = m_r - 1$ corresponds to the most asymmetric setup, where m_r is the number of configurations with r repeaters. All configurations are then stored in a table by their values for r and n .

Then, when we optimize over the hardware parameters, we also optimize over two additional parameters. These are r (the number of repeaters) and a , which is a number between zero and one. Given a pair (r, a) , the configuration that is used is chosen as follows. First, the number a is mapped to a value of n using

$$n = \text{round}(a(m_r - 1)), \quad (\text{B2})$$

(where round denotes rounding to the closest integer) i.e. it uses $n = 0$ for $a = 0$ and $n = m_r - 1$ for $a = 1$. Second, the unique configuration defined by the values of r and n is taken from the table and used in the simulation. The reason why we optimize over a instead of over n directly is that a quantifies how asymmetric the chosen configuration is in a way that is independent of r (the range is always between 0 and 1, instead of between 0 and $m_r - 1$). This makes it easier to vary r and a independently compared to r and n .

Appendix C. Optimization method

In this appendix we provide more details regarding our optimization methodology. This methodology is based on genetic algorithms, which come in several different flavors. Our particular implementation is heavily based on the one introduced in [66] and used in [22, 74, 75] to which we refer the interested reader. There are two things that we do discuss in this appendix. First, as the parameter set we use here is different than in [22, 66, 74, 75], we explain in section C.1 how we define the probability of no imperfection for each of these, as required by the definition of the hardware cost function H_c in section 1.3. Second, we have employed a simple local optimization performed on the best solution found by the genetic algorithm. This local optimization method has been used and described in [74], but has not yet appeared in a peer-reviewed publication. We therefore explain it below in section C.2. Additionally, we also give the details of the machine used to perform the actual optimizations in section C.3. Finally we would like to remark that the code for our implementation, together with the tools required for integration with NetSquid simulations, is publicly accessible at [76].

C.1. No-imperfection probabilities

We show in table 5 the probability of no-imperfection for all parameters considered in our hardware models.

We start by defining the quantity $p_{\text{surv baseline}}$ that appears in this table more rigorously. It is computed as follows,

$$p_{\text{surv baseline}} = 10^{-\overline{\alpha_{\text{att}}}/10}, \quad (\text{C1})$$

with $\overline{\alpha_{\text{att}}}$ given by,

$$\overline{\alpha_{\text{att}}} = \frac{2}{N} \sum_{i=1}^N \alpha_{\text{att},i} L_i. \quad (\text{C2})$$

Here, L_i is the length of fiber segment i in the fiber path under consideration, $\alpha_{\text{att},i}$ is the attenuation coefficient of fiber segment i (i.e. the amount of attenuation per unit length), and N is the total number of fiber segments in the path. For the fiber path considered in this paper (i.e. the one depicted in figure 1), $N = 17$. An elementary link between two neighboring nodes must consist of at least two fiber segments to allow for the installation of a heralding station. $\overline{\alpha_{\text{att}}}$ can then be thought of as the total amount of attenuation on an elementary link made up of two times the average fiber segment. This means $p_{\text{surv baseline}}$ is the probability of a photon surviving traveling through this average elementary link. The reason for constructing this quantity is that it provides a baseline for the photon survival probability in fiber, which can then be improved upon by increasing the number of multiplexing modes, thereby enabling us to associate a cost function.

The coherence time T represents a timescale for depolarization, with the probability of the state becoming maximally mixed over a period of time t being given by $1 - e^{-t/T}$, with the respective probability of no-imperfection then being $e^{-t/T}$. In this case, improving T by a factor of k is equivalent to multiplying it by k .

Table 5. Probabilities of no-imperfection for hardware parameters we optimized over in this work. $p_{\text{surv baseline}}$ (defined in equation (C1)) is the probability of one photon (i.e. no multiplexing) surviving traveling an elementary link made up out of two times the average fiber segment in the fiber path we study (shown in figure 1).

Parameter	Probability of no-imperfection
Photon detection probability p_{det}	p_{det}
Coherence time T	$e^{-1/T}$
Swap quality s_q	s_q
Elementary link fidelity F_{el}	F_{el}
Number of multiplexing modes N	$1 - (1 - p_{\text{surv baseline}})^N$

For the swap quality, s_q is the probability that the two-qubit state before the Bell-state measurement is not replaced with a maximally mixed state, and therefore s_q is the corresponding probability of no imperfection. Finally, for the elementary-link fidelity we take the fidelity itself to be the probability of no imperfection.

C.2. Local optimization

Genetic algorithms are derivative-free optimization algorithms that are particularly useful when applied to functions whose cost landscape is largely unknown but is assumed to have many local minima [77]. Through a balancing act of exploration (i.e. investigation of many different areas of parameter space) and exploitation (i.e. investigation of local optima) they often manage to avoid being trapped in local optima as gradient-based methods are wont to. Nevertheless, use of a genetic algorithm does not guarantee that one can find the global optimum. Further, one can not even be sure that one has maximally exploited the best optimum found. For this reason, we complement the exploration performed by the genetic algorithm with a deterministic local optimization method which we apply to the best parameter set found by the genetic algorithm. The algorithm used is a variation of an iterative local search algorithm [78]. It consists of iteratively making small changes on a parameter and evaluating the cost associated to the resulting parameter set. In case it has decreased, it is kept and we again make a small change on the same parameter. If the cost increases, we discard the change and move on to another parameter. This process is repeated for all parameters being optimized over. We must however emphasize that this still does not guarantee that the global optimum will be found.

More details on this method can be found in chapter 4.2 of [74].

C.3. Performing the optimization

Each optimization run was executed on a thin node of the Snellius supercomputer [79]. Each of these nodes is endowed with 2 AMD Rome 7H12 CPUs (2.6 GHz), for a total of 128 cores and a total of 256 GiB of memory.

Appendix D. BQC test protocol

In this appendix we describe the BQC test protocol that is used as a performance metric in this paper. This protocol consists of repeated execution of test rounds as required by the VBQC protocol presented in [28]. In each round of the VBQC protocol, a server is tasked by a client to execute a quantum computation on qubits transmitted by the client and then send the classical result of that computation back to the client. In test rounds, the client has prepared the transmitted qubits in such a way that it knows the correct outcome of the computation.

Therefore, executing test rounds allows the client to verify whether the server is honest. However, noise in the quantum hardware can also lead to failed test rounds. The more often test rounds fail due to noise, the harder it is for the client to verify the server's honesty.

The BQC test protocol that we consider is not itself a VBQC protocol. In fact, its only purpose is to benchmark how suited a quantum network could be to perform BQC protocols (and perhaps other applications that require multiple live qubits simultaneously). The performance metric that we consider for this protocol is the success rate, defined as the average number of successful test rounds that can be executed per time unit (i.e. the product of the rate R and success probability p_s , as in equation (8)). We specifically consider an entanglement-based two-qubit version of the protocol. In that case, the protocol is as follows:

1. The client chooses d and r uniformly at random from $\{0, 1\}$ and θ from $\{j\pi/4\}_{0 \leq j \leq 7}$, and then defines two quantum states, $|\text{dummy}\rangle = |i\rangle$ and $|\text{trap}\rangle = |+\theta\rangle$, where $|\pm\phi\rangle \equiv \frac{1}{\sqrt{2}}(|0\rangle \pm e^{i\phi}|1\rangle)$. Additionally, it uniformly at random designates $|\phi_1\rangle$ to be $|\text{dummy}\rangle$ or $|\text{trap}\rangle$. $|\phi_2\rangle$ is designated to be the option that was not chosen.

2. When an entangled state shared between the client and server becomes available, the client uses quantum teleportation to transmit the state $|\phi_1\rangle$ to the server. The server stores the received state in quantum memory.
3. When a second entangled state becomes available, the client uses quantum teleportation to transmit the state $|\phi_2\rangle$ to the server.
4. The server performs a CZ gate between its two qubits.
5. The server measures the qubit that was used to receive the state $|\text{trap}\rangle$ in the basis $\{|+\theta+r\pi\rangle, |-\theta+r\pi\rangle\}$ and transmits the result back to the client.
6. The client declares the test round a success if it receives measurement result matching its expectation (i.e. if the outcome is equal to $d \oplus r$, where \oplus is addition modulo two), and a failure otherwise.
7. The client and server go back to Step 1 to start the next test round.

Alternatively, remote state preparation [80] could be used to prepare the required states at the server, which may be easier to execute on real hardware than quantum teleportation. In fact, it is proven in [22] that using remote state preparation for the VBQC protocol in [28] is equivalent to using quantum teleportation in case the client and server implement gates noiselessly. Therefore the success rate will, under these assumptions, be the same whether quantum teleportation or remote state preparation is used.

We here assume that classical communication between the client and the server happens instantaneously and that both the client and server are able to perform gates and measurements noiselessly and instantly. However we do not assume they are able to store qubits indefinitely; the first teleported state undergoes depolarizing noise as described in equation (6), where the coherence time T is the same as the coherence time of the repeater nodes (i.e. it is varied by the optimizations performed in this paper). Under these assumptions, R_{rounds} is simply half the rate at which entanglement can be distributed when entanglement is being generated continuously, as one test round can be performed for every two entangled states that are produced. In order to calculate the success probability, we use the following result from [22]:

$$1 - p_s = e^{-\frac{\Delta t}{T}} [F_{\text{dummy}} (1 - F_{\text{trap}}) + F_{\text{trap}} (1 - F_{\text{dummy}})] + \frac{1}{2} (1 - e^{-\frac{\Delta t}{T}}). \quad (\text{D1})$$

Here, Δt is the time between the transmission of the first qubit and the second qubit. For the fidelities F_{trap} and F_{dummy} , let the density matrices for the state $|\text{dummy}\rangle$ after transmission to the server be ρ_{dummy} and ρ_{trap} for $|\text{trap}\rangle$. Then $F_{\text{dummy}} = \langle \text{dummy} | \rho_{\text{dummy}} | \text{dummy} \rangle$ and $F_{\text{trap}} = \langle \text{trap} | \rho_{\text{trap}} | \text{trap} \rangle$.

We then determine the success rate as follows. First, we simulate continuous entanglement generation between the end nodes of a repeater chain. Each time an end-to-end entangled state is generated it is removed from the simulation and stored as raw data, together with the time at which it was generated. Then, after the simulation has finished, we process the raw data to determine what the success rate would have been if the entangled states had been consumed by the BQC test protocol. To this end, we divide the data into single test rounds, each consisting of two entangled states that were generated in succession. We assign each test round a duration t , which is the amount of time between the start of the round and the end of the round (i.e. when the second state was generated), and a storage time Δt , which is the time between when the first entangled state and the second entangled state were generated. We furthermore calculate the p_s of that round using equation (D1), where we average over the two possible choices in the protocol for how $|\phi_1\rangle$ and $|\phi_2\rangle$ are designated (i.e. whether the first entangled state is used to transmit the dummy and the second to transmit the trap or vice versa). Then we calculate the rate as

$$R = \frac{1}{\langle t \rangle}, \quad (\text{D2})$$

where $\langle t \rangle$ is the average value of t over all the test rounds. Finally, we use R and the average value of p_s to calculate the success rate according to equation (8). The processing code that realizes this calculation has been made publicly available at [81].

Appendix E. Comparison to repeaterless bound

Each quantum channel has an associated secret-key capacity, which is the asymptotically optimal number of secret-key bits that can be shared between two parties per use of the channel. In quantum-repeater studies, SKRs are often compared to what is known as the repeaterless bound or the PLOB bound [82]. The repeaterless bound is the secret-key capacity corresponding to a pure-loss channel with a loss parameter equal to the total transmission loss associated with sending a photon directly from end node to end node. For an optical-fiber path, as considered in this paper, this loss parameter corresponds to the total attenuation loss

of the fiber connecting the two end nodes. Observing key rates beating the repeaterless bound can then be considered proof that the repeaters genuinely improve performance. In this appendix, we compare the performance achieved for our optimization solutions to the repeaterless bound and observe that it is always exceeded by at least about 14 orders of magnitude.

First, we calculate the repeaterless bound for the network path under consideration. The total attenuation losses on the fiber path add up to 214.6 dB (see figure 1). This corresponds to a transmissivity (which is one minus the loss probability) of $\eta \approx 3.47 \times 10^{-22}$. The repeaterless bound K is then calculated as [82]

$$K = -\log_2(1 - \eta) \approx 5.0 \times 10^{-22} \text{ secret bits per channel use.} \quad (\text{E1})$$

Now, we determine a lower bound on the number of secret bits shared per channel use when using the QKD-based optimization solutions. First we upper bound the rate of channel uses, which we denote C . In order to be conservative, we define the total number of channel uses in a repeater chain as the sum over the number of channel uses of all the segments of the repeater chain, as is done in, e.g. [83]. As the maximum number of repeaters that can be placed in the fiber path is seven (see figure 1), which corresponds to the chain being made up out of eight elementary links, we can bound the rate of channel uses as

$$C \leq 8C_{\text{link}}, \quad (\text{E2})$$

where C_{link} is an upper bound on the rate of channel uses of the elementary links in the chain. We will now proceed to determine such a bound.

In our model, attempts at entanglement generation in an elementary link take exactly the amount of time that a light signal would need to travel through that elementary link, i.e. L/c (see equation (2)). This means that an elementary link can perform attempts at a rate of at most c/L . In practice the average rate will be smaller, as entanglement generation is halted after success and until the necessary qubits in the repeater are freed up again. The number of channel uses corresponding to each attempt is equal to the number of multiplexing modes, multiplied by a factor of two in case the double-click protocol is used (as each photon is then encoded in two modes). While our model is agnostic about the exact entanglement-generation protocol used, we will here be pessimistic and assume the double-click protocol is used. The largest number of multiplexing modes over all optimization results is 799 (see table 2). Moreover, the shortest possible elementary link possible between two adjacent repeaters on the fiber path (which consists of two adjacent fiber segments, such that a heralding station can still be placed in between) is 74.4 km. Therefore, we upper-bound the rate of channel uses on any of the elementary links as

$$C_{\text{link}} = \frac{c}{74.4 \text{ km}} \times 799 \times 2 \approx 4.3 \times 10^6 \text{ Hz.} \quad (\text{E3})$$

This gives $C \lesssim 3.4 \times 10^7$ Hz.

Now, the amount of secret bits produced by the repeater chain per channel uses R_c can be obtained by dividing the SKR R by the rate of channel uses C . As $R \geq 1$ Hz (1 Hz being the smallest target rate that we consider),

$$R_c = \frac{R}{C} \gtrsim 2.9 \times 10^{-8} \text{ secret bits per channel use.} \quad (\text{E4})$$

As we see, our pessimistic lower bound on R_c is about 14 orders of magnitude better than the repeaterless bound K . Therefore, it can be unambiguously concluded that the optimization solutions found in this paper ‘beat the bound’.

ORCID iDs

Francisco Ferreira da Silva  <https://orcid.org/0000-0003-3642-4350>

Guus Avis  <https://orcid.org/0000-0001-7035-8065>

Joshua A Slater  <https://orcid.org/0000-0002-7297-6163>

Stephanie Wehner  <https://orcid.org/0000-0002-8433-0730>

References

- [1] Ekert A K 1991 Quantum cryptography based on Bell’s theorem *Phys. Rev. Lett.* **67** 661
- [2] Bennett C H and Brassard G 2014 Quantum cryptography: public key distribution and coin tossing *Theor. Comput. Sci.* **560** 7–11
- [3] Buhrman H and Röhrig H 2003 Distributed quantum computing *Int. Symp. on Mathematical Foundations of Computer Science* (Springer) pp 1–20
- [4] Hillery M, Bužek V and Berthiaume A 1999 Quantum secret sharing *Phys. Rev. A* **59** 1829

- [5] Wehner S, Elkouss D and Hanson R 2018 Quantum internet: a vision for the road ahead *Science* **362** eaam9288
- [6] Lago-Rivera D, Grandi S, Rakonjac J V, Seri A and de Riedmatten H 2021 Telecom-heralded entanglement between multimode solid-state quantum memories *Nature* **594** 37–40
- [7] Pompili M et al 2021 Realization of a multinode quantum network of remote solid-state qubits *Science* **372** 259–64
- [8] Park J L 1970 The concept of transition in quantum mechanics *Found. Phys.* **1** 23–33
- [9] Wootters W K and Zurek W H 1982 A single quantum cannot be cloned *Nature* **299** 802–3
- [10] Dieks D 1982 Communication by EPR devices *Phys. Lett. A* **92** 271–2
- [11] Briegel H-J, Dür W, Cirac J I and Zoller P 1998 Quantum repeaters: the role of imperfect local operations in quantum communication *Phys. Rev. Lett.* **81** 5932
- [12] Dür W, Briegel H-J, Cirac J I and Zoller P 1999 Quantum repeaters based on entanglement purification *Phys. Rev. A* **59** 169
- [13] Munro W J, Azuma K, Tamaki K and Nemoto K 2015 Inside quantum repeaters *IEEE J. Sel. Top. Quantum Electron.* **21** 78–90
- [14] Muralidharan S, Li L, Kim J, Lütkenhaus N, Lukin M D and Jiang L 2016 Optimal architectures for long distance quantum communication *Sci. Rep.* **6** 20463
- [15] Sangouard N, Simon C, De Riedmatten H and Gisin N 2011 Quantum repeaters based on atomic ensembles and linear optics *Rev. Mod. Phys.* **83** 33
- [16] Duan L-M, Lukin M, Cirac I and Zoller P 2001 Long-distance quantum communication with atomic ensembles and linear optics *Nature* **414** 413–8
- [17] Monroe C et al 2013 Quantum networks with atoms and photons *J. Phys.: Conf. Ser.* **467** 012008
- [18] Reiserer A and Rempe G 2015 Cavity-based quantum networks with single atoms and optical photons *Rev. Mod. Phys.* **87** 1379–418
- [19] Krutyanskiy V, Canteri M, Meraner M, Bate J, Krucmarzsky V, Schupp J, Sangouard N and Lanyon B P 2023 Telecom-wavelength quantum repeater node based on a trapped-ion processor *Phys. Rev. Lett.* **130** 213601
- [20] Langenfeld S, Thomas P, Morin O and Rempe G 2021 Quantum repeater node demonstrating unconditionally secure key distribution *Phys. Rev. Lett.* **126** 230506
- [21] Ruf M, Wan N H, Choi H, Englund D and Hanson R 2021 Quantum networks based on color centers in diamond *J. Appl. Phys.* **130** 070901
- [22] Avis G, Ferreira da Silva F, Coopmans T, Dahlberg A, Jirovská H, Maier D, Rabbie J, Torres-Knoop A and Wehner S 2023 Requirements for a processing-node quantum repeater on a real-world fiber grid *npj Quantum Inf.* **9** 100
- [23] Rabbie J, Chakraborty K, Avis G and Wehner S 2022 Designing quantum networks using preexisting infrastructure *npj Quantum Inf.* **8** 5
- [24] Stanley M, Gui Y, Unnikrishnan D, Hall S and Fatadin I 2022 Recent progress in quantum key distribution network deployments and standards *J. Phys.: Conf. Ser.* **2416** 012001
- [25] Salvail L et al 2010 Security of trusted repeater quantum key distribution networks *J. Comput. Secur.* **18** 61–87
- [26] Bennett C H, Brassard G and Mermin N D 1992 Quantum cryptography without bell's theorem *Phys. Rev. Lett.* **68** 557
- [27] Broadbent A, Fitzsimons J and Kashefi E 2009 universal blind quantum computation 2009 50th Annual IEEE Symp. on Foundations of Computer Science (IEEE) pp 517–26
- [28] Leichte D, Music L, Kashefi E and Ollivier H 2021 Verifying BQP computations on noisy devices with minimal overhead *PRX Quantum* **2** 040302
- [29] Cabrillo C, Cirac J I, Garcia-Fernandez P and Zoller P 1999 Creation of entangled states of distant atoms by interference *Phys. Rev. A* **59** 1025
- [30] Barrett S D and Kok P 2005 Efficient high-fidelity quantum computation using matter qubits and linear optics *Phys. Rev. A* **71** 060310
- [31] Humphreys P C, Kalb N, Morits J P, Schouten R N, Vermeulen R F, Twitchen D J, Markham M and Hanson R 2018 Deterministic delivery of remote entanglement on a quantum network *Nature* **558** 268–73
- [32] Bernien H et al 2013 Herald ed entanglement between solid-state qubits separated by three metres *Nature* **497** 86–90
- [33] Hensen B et al 2015 Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres *Nature* **526** 682–6
- [34] Stephenson L J, Nadlinger D P, Nichol B C, An S, Drmota P, Ballance T G, Thirumalai K, Goodwin J F, Lucas D M and Ballance C J 2020 High-rate, high-fidelity entanglement of qubits across an elementary quantum network *Phys. Rev. Lett.* **124** 110501
- [35] Krutyanskiy V et al 2023 Entanglement of trapped-ion qubits separated by 230 meters *Phys. Rev. Lett.* **130** 050803
- [36] Azuma K, Economou S E, Elkouss D, Hilaire P, Jiang L, Lo H-K and Tzitrin I 2023 Quantum repeaters: from quantum networks to the quantum internet *Rev. Mod. Phys.* **95** 045006
- [37] Northup T E and Blatt R 2014 Quantum information transfer using photons *Nat. Photon.* **8** 356–63
- [38] Pfaff W et al 2014 Unconditional quantum teleportation between distant solid-state quantum bits *Science* **345** 532–5
- [39] Kalb N, Reiserer A A, Humphreys P C, Bakermans J J, Kamberling S J, Nickerson N H, Benjamin S C, Twitchen D J, Markham M and Hanson R 2017 Entanglement distillation between solid-state quantum network nodes *Science* **356** 928–32
- [40] Slodička L, Hétet G, Röck N, Schindler P, Hennrich M and Blatt R 2013 Atom-atom entanglement by single-photon detection *Phys. Rev. Lett.* **110** 083603
- [41] Bhaskar M K et al 2020 Experimental demonstration of memory-enhanced quantum communication *Nature* **580** 60–64
- [42] Langenfeld S, Welte S, Hartung L, Daiss S, Thomas P, Morin O, Distant E and Rempe G 2021 Quantum teleportation between remote qubit memories with only a single photon as a resource *Phys. Rev. Lett.* **126** 130502
- [43] Lin G W, Zou X B, Lin X M and Guo G C 2009 Herald ed quantum memory for single-photon polarization qubits *Europhys. Lett.* **86** 30006
- [44] Hermans S, Pompili M, dos Santos Martins L, Montblanch A R-P, Beukers H, Baier S, Borregaard J and Hanson R 2023 Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study *New J. Phys.* **25** 013011
- [45] Horodecki M, Horodecki P and Horodecki R 1999 General teleportation channel, singlet fraction and quasidistillation *Phys. Rev. A* **60** 1888–98
- [46] Pompili M et al 2022 Experimental demonstration of entanglement delivery using a quantum network stack *npj Quantum Inf.* **8** 1–10
- [47] Childress L, Taylor J M, Sørensen A S and Lukin M D 2005 Fault-tolerant quantum repeaters with minimal physical resources and implementations based on single-photon emitters *Phys. Rev. A* **72** 052330
- [48] Askarani M F, Chakraborty K and Do Amaral G C 2021 Entanglement distribution in multi-platform buffered-router-assisted frequency-multiplexed automated repeater chains *New J. Phys.* **23** 063078
- [49] van Dam S B, Humphreys P C, Rozpedek F, Wehner S and Hanson R 2017 Multiplexed entanglement generation over quantum networks using multi-qubit nodes *Quantum Sci. Technol.* **2** 034002

- [50] Sinclair N *et al* 2014 Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control *Phys. Rev. Lett.* **113** 053603
- [51] Coopmans T *et al* 2021 Netsquid, a network simulator for quantum information using discrete events *Commun. Phys.* **4** 1–15
- [52] Inesta A G, Vardoyan G, Scavuzzo L and Wehner S 2022 Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs *npj Quantum Inf.* **9** 46
- [53] Bennett C H, Brassard G, Crépeau C, Jozsa R, Peres A and Wootters W K 1993 Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels *Phys. Rev. Lett.* **70** 1895
- [54] Shor P W and Preskill J 2000 Simple proof of security of the bb84 quantum key distribution protocol *Phys. Rev. Lett.* **85** 441
- [55] Lo H-K, Curty M and Qi B 2012 Measurement-device-independent quantum key distribution *Phys. Rev. Lett.* **108** 130503
- [56] Minder M, Pittaluga M, Roberts G L, Lucamarini M, Dynes J F, Yuan Z and Shields A J 2019 Experimental quantum key distribution beyond the repeaterless secret key capacity *Nat. Photon.* **13** 334–8
- [57] Pascoletti A and Serafini P 1984 Scalarizing vector optimization problems *J. Optim. Theory Appl.* **42** 499–524
- [58] Schaffer J D 1985 Some experiments in machine learning using vector evaluated genetic algorithms *Technical Report* (Vanderbilt University)
- [59] Bradley C *et al* 2022 Robust quantum-network memory based on spin qubits in isotopically engineered diamond *npj Quantum Inf.* **8** 122
- [60] Hermans S, Pompili M, Beukers H, Baier S, Borregaard J and Hanson R 2022 Qubit teleportation between non-neighbouring nodes in a quantum network *Nature* **605** 663–8
- [61] Bradley C, Randall J, Abobeih M, Berrevoets R, Degen M, Bakker M, Markham M, Twitche D and Taminiau T 2019 A ten-qubit solid-state spin register with quantum memory up to one minute *Phys. Rev. X* **9** 031045
- [62] Taminiau T H, Cramer J, van der Sar T, Dobrovitski V V and Hanson R 2014 universal control and error correction in multi-qubit spin registers in diamond *Nat. Nanotechnol.* **9** 171–6
- [63] Avis G and Ferreira da Silva F 2023 NetSquid-QRepChain (available at: <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-qrepchain>)
- [64] Avis G and Ferreira da Silva F 2023 Simulation code for Requirements for upgrading trusted nodes to a repeater chain over 900 km of optical fiber *GitLab* (available at: <https://gitlab.com/softwarequtech/simulation-code-for-requirements-for-upgrading-trusted-nodes-to-a-repeater-chain-over-900-km-of-optical-fiber>)
- [65] Ferreira da Silva F, Avis G and Wehner S 2023 Replication data for: requirements for upgrading trusted-nodes to a repeater chain over 900 km of optical fiber (arXiv:2303.03234)
- [66] da Silva F F, Torres-Knoop A, Coopmans T, Maier D and Wehner S 2021 Optimizing entanglement generation and distribution using genetic algorithms *Quantum Sci. Technol.* **6** 035007
- [67] Avis G, Kneigiens R, Sørensen A S and Wehner S 2024 Asymmetric node placement in fiber-based quantum networks *Phys. Rev. A* **109** 052627
- [68] Deutsch D, Ekert A, Jozsa R, Macchiavello C, Popescu S and Sanpera A 1996 Quantum privacy amplification and the security of quantum cryptography over noisy channels *Phys. Rev. Lett.* **77** 2818
- [69] Ritter S, Nölleke C, Hahn C, Reiserer A, Neuzner A, Uphoff M, Mücke M, Figueroa E, Bochmann J and Rempe G 2012 An elementary quantum network of single atoms in optical cavities *Nature* **484** 195–200
- [70] Krastanov S, Albert V V and Jiang L 2019 Optimized entanglement purification *Quantum* **3** 123
- [71] Goodenough K, Elkouss D and Wehner S 2021 Optimizing repeater schemes for the quantum internet *Phys. Rev. A* **103** 032610
- [72] Haldar S, Barge P J, Cheng X, Chang K-C, Kirby B T, Khatri S, Wong C W and Lee H 2024 Policies for multiplexed quantum repeaters: theory and practical performance analysis (arXiv:2401.13168)
- [73] Zaske S *et al* 2012 Visible-to-telecom quantum frequency conversion of light from a single quantum emitter *Phys. Rev. Lett.* **109** 147404
- [74] Labay Mora A 2021 Genetic algorithm-based optimisation of entanglement distribution to minimise hardware cost *Master's Thesis* Delft University of Technology
- [75] Labay Mora A, Ferreira da Silva F and Wehner S 2023 Reducing hardware requirements for entanglement distribution via joint hardware-protocol optimization *Quantum Sci. Technol.* **9** 045001
- [76] Torres-Knoop A, Coopmans T, Maier D and Silva F 2020 smart-stopos (available at: <https://gitlab.com/aritoka/smart-stopos>)
- [77] Goldberg D E 1989 *Genetic Algorithms in Search, Optimization and Machine Learning* 1st edn (Addison-Wesley Longman Publishing Co., Inc.)
- [78] Luke S 2013 *Essentials of Metaheuristics* vol 2 (Lulu Raleigh)
- [79] Snellius (available at: www.surf.nl/en/dutch-national-supercomputer-snellius)
- [80] Bennett C H, DiVincenzo D P, Shor P W, Smolin J A, Terhal B M and Wootters W K 2001 Remote state preparation *Phys. Rev. Lett.* **87** 077902
- [81] NetSquid-SimulationTools (available at: <https://gitlab.com/softwarequtech/netsquid-snippets/netsquid-simulationtools>)
- [82] Pirandola S, Laurenza R, Ottaviani C and Banchi L 2017 Fundamental limits of repeaterless quantum communications *Nat. Commun.* **8** 1–15
- [83] Rozpedek F, Yehia R, Goodenough K, Ruf M, Humphreys P C, Hanson R, Wehner S and Elkouss D 2019 Near-term quantum-repeater experiments with nitrogen-vacancy centers: overcoming the limitations of direct transmission *Phys. Rev. A* **99** 052330