



Comparing one- and two-way quantum repeater architectures



Prateek Mantri , Kenneth Goodenough & Don Towsley

Quantum repeaters are an essential building block for realizing long-distance quantum communications. However, due to the fragile nature of quantum information, these repeaters suffer from loss and operational errors. Prior works have classified repeaters into three broad categories based on their use of probabilistic or near-deterministic methods to mitigate these errors. Besides differences in classical communication times, these approaches also vary in technological complexity, with near-deterministic methods requiring more advanced hardware. Recent increases in memory availability and advances in multiplexed entanglement generation motivate a fresh comparison of one-way and two-way repeater architectures. In this work, we present a two-way repeater protocol that combines multiplexing with application-aware distillation, designed for a setting where sufficient high-quality memory resources are available—reflecting architectural assumptions expected in large-scale network deployments. We introduce a recursive formulation to track the full probability distribution of Bell pairs in multiplexed two-way repeater architectures, enabling the performance analysis of multiplexed repeater schemes which use probabilistic n -to- k distillation. Using this framework, we compare the proposed two-way protocol with one-way schemes in parameter regimes previously believed to favour the latter, and find that the two-way architecture consistently outperforms one-way protocols while requiring lower technological and resource overheads.

Quantum communication is poised to enable transformative applications in quantum sensing^{1–4}, distributed quantum computing^{5–6}, secure communications^{7–8}, and quantum secret sharing^{9–10}, among others. As quantum processors grow in sophistication, a scalable network for quantum information transfer between spatially separated nodes is becoming increasingly critical.

However, transmitting quantum information over long distances poses significant challenges, primarily due to losses that grow exponentially with distance in optical fibers. Unlike classical communication, quantum networks cannot employ classical “receive and re-transmit” strategies because of the no-cloning theorem, which prohibits the duplication of unknown quantum states. Prior studies have established fundamental limits on direct quantum information transmission^{11,12}, underscoring the need for innovative solutions such as quantum repeaters.

Quantum repeaters are specialized devices designed to extend the range of quantum communications by dividing long segments into shorter segments, thereby mitigating losses through specific quantum gates and measurement operations. These devices significantly enhance the viability of long-distance quantum communications by employing shorter, manageable links to create extended connections. However, the implementation of these operations is fraught with errors, which can restrict the effective

distance for practical quantum communication. To address both loss and operational errors, researchers have proposed various probabilistic (heralded generation and distillation operations) and near-deterministic (quantum error correction) approaches. These strategies have led to the development of three generations of quantum repeaters, each with distinct characteristics and technological requirements^{13,14}.

First-generation architectures use probabilistic entanglement generation to mitigate loss errors and heralded entanglement purification (often probabilistic) to mitigate operational errors. Second-generation architectures also use probabilistic entanglement generation but incorporate near-deterministic quantum error correction to address operational errors. Third-generation architectures solely rely on near-deterministic quantum error correction schemes to correct both operation and loss errors. The probabilistic solutions, while more feasible with current technology, necessitate heralding and consequently suffer from increased temporal costs associated with classical communication. In contrast, quantum error correction-based approaches typically employ one-way signaling and have the potential to achieve higher secret-key rates. In this manuscript, we focus on first-generation (referred to in this manuscript as the two-way schemes) and third-generation (referred to as the one-way schemes) quantum

Robert and Donna Manning College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA, USA.

e-mail: pmantri@cs.umass.edu

repeater architectures. As a special case, we also include a particular variant of second-generation networks that aims to use multiplexing to establish at most a single elementary link between the segments. However, this approach does not employ error correction to protect against operational noise.

To be viable, one-way schemes typically require almost perfect operations, along with complex encoding and decoding circuitry with a large number of measurement and gate operations. Moreover, to be able to correct for fiber loss errors, one-way schemes also require repeaters to be more closely spaced compared to two-way schemes^{15,16}. This trade-off between complexity and performance motivates a re-examination of both approaches across relevant parameter regimes. In particular, it is important to assess whether the purported gains of one-way schemes—despite their high technological and resource demands—can instead be realized using the comparatively simpler two-way architectures.

This motivates a thorough comparison between different quantum repeater approaches. In their pivotal work, Muralidharan et al.¹⁴ performed a foundational comparison of the three generations of quantum repeaters, identifying coupling efficiency (η_c), gate errors (ϵ_G), and gate times (t_G) as critical parameters for evaluation. They delineated the specific parameter regimes in which each generation is expected to excel. However, as quantum technologies advance, revisiting these comparisons with updated models and technologies is essential to ensure accurate assessments and practical guidance for implementation.

One such example of advances in quantum technology has been in the area of long-lived memories^{17–19}—a crucial requirement for the viability of two-way schemes over long distances. Unlike one-way repeaters, which may only need such memories in the case of slow gate operations, two-way protocols necessarily depend on memories that outlive the round-trip classical communication time for heralded feedback. Most prior analyses of two-way repeater schemes have focused on scenarios where memory availability is highly constrained, making memories the most significant cost factor^{14,20}. This focus has shaped the strategy of using multiplexing to maximize the success probability of at least one successful link per segment^{14,21–23}. This kind of scheme has also been referred to as the second generation without encoding or “2G-NC”¹⁴.

Some studies have explored more aggressive forms of multiplexing that aim for multiple simultaneous successes per segment^{20,24–28}. However, most of these studies analyze throughput in the context of linear-chain quantum networks with probabilistic swapping realized through optical circuits, but without any distillation capabilities. In such networks, a nested swapping schedule provides limited benefit. In the cases where nested distillation schemes have been considered in the context of multiple-success multiplexing, deterministic distillation

protocols are often used to simplify calculations^{20,25}. Some protocols also employ blind operations—where the repeater proceeds without waiting for classical information, typically at the cost of an exponentially decreasing success probability—when considering a nested swapping schedule, simplifying the computation of the number of Bell pairs delivered and the time required^{20,25}. To the best of our knowledge, no prior work has performed a detailed comparison between one-way and multiplexed two-way protocols (MTPs) under such considerations.

In this manuscript, we investigate the performance of multiplexed two-way (first-generation) and one-way (third-generation) quantum repeater architectures in parameter regimes where one-way schemes have previously been considered advantageous. We consider a nested swapping protocol tailored for a setting in which a sufficient number of high-quality quantum memories are available—reflecting the architectural assumptions expected in practical large-scale network deployments. This protocol supports flexible distillation scheduling optimized for different service metrics, such as secret-key rate or fidelity thresholds. We aim to provide an even-handed performance comparison between these architectures under conditions that reflect large-scale network implementations. To support this, we introduce a recursive numerical framework that captures the full probability distribution of the number of available Bell pairs at each stage of the protocol. This framework enables a fine-grained comparison of architectures using performance metrics relevant to real-world deployment.

Methods

Two-way repeater architecture

In this manuscript, we consider a linear network with each repeater station being equipped with a large number of optically active memories or emitters. These memories emit photons, which are then sent to a station located at the midpoint of the link connecting the two repeater stations. At this midpoint station, photons from two different repeaters are entangled and measured together, effectively creating a Bell pair link shared between the repeaters (see Fig. 1).

Multiplexing for exactly one success across the network (2G-NC). Multiplexing is a well-known technique in telecommunications and computer networks, where multiple information channels are combined over a shared medium. Multiplexed quantum repeaters can be used to overcome the probabilistic loss errors associated with signal decay in optical fibers. This is achieved by attempting multiple Bell pair generation attempts in parallel through either spatial, time-bin, or frequency multiplexing. Due to poor rates associated with entanglement generation sources, various proposals have been made over time for the use of multiplexing to improve rates in quantum repeaters^{29–31}. The most basic of these proposals involves parallelizing operations and sending multiple

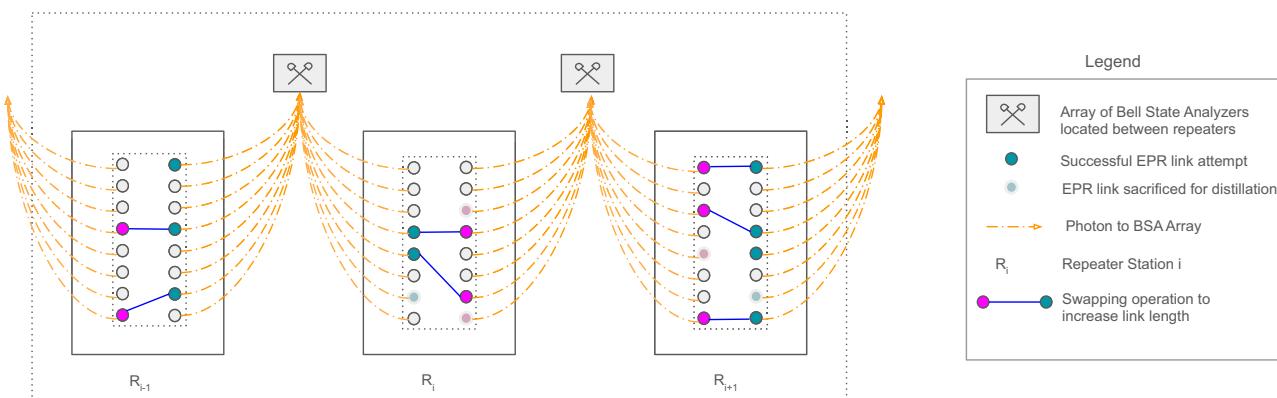


Fig. 1 | A schematic of the multiplexed two-way repeater scheme. Each repeater (denoted R_i , etc.) has multiple emitter memories located on both sides. Each of these memories emits photons, creating a spatially or spectrally multiplexed burst. At the

midpoint between any two repeaters, an array of Bell state analyzers exists to entangle photons coming from both sides. Only a fraction of photons from either side survive the journey and reach the midpoint station.

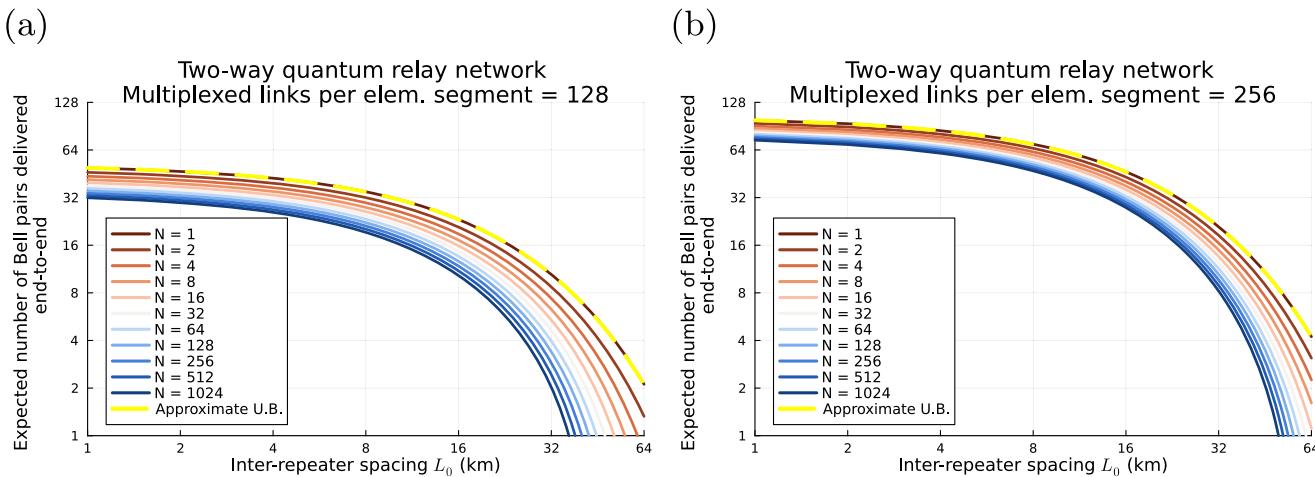


Fig. 2 | Expected number of surviving Bell pairs per shot versus inter-repeater distance in a quantum relay network. **a** Results for 128 multiplexed elementary links; **b** Results for 256 multiplexed elementary links. The y-axis denotes the expectation of the minimum number of Bell pairs generated across all segments. With deterministic swapping operations, the minimum number of Bell pairs across

segments also translates into the number of Bell pairs delivered end-to-end. The x-axis denotes the inter-repeater spacing, and N denotes the number of segments. The yellow dashed line shows an upper bound used by prior analyses to approximate the number of Bell pairs delivered.

photons over an optical fiber (using time-division, spectral, or spatial multiplexing), performing Bell State Measurements (BSMs), effectively creating entanglement between multiple matter qubits. To simplify analysis, these techniques often focus on using multiplexing to maximize the success of at least one Bell pair, with only one Bell pair kept between repeaters in the event of multiple successes. Moreover, this technique requires spatial or temporal multiplexing, which is often realized using lossy optical switches, adding further loss²². However, a recent proposal by Chen et al.³¹ uses spectral multiplexing and parallel entanglement creation to achieve high transmission rates without the added losses usually incurred from spatial or temporal multiplexing methods.

Another line of thought has explored multiplexing for quantum networks in multiple degrees of freedom of a single photon^{30,32}. This scheme uses a single photonic pulse to entangle multiple pairs of remote memories, minimizing the need for extensive spatial channels and precise temporal coordination. These proposals simplify the infrastructure needs while enhancing the rate at which entangled pairs can be generated in a quantum network. Furthermore, these techniques have also been extended to one-way schemes^{33,34}.

In Muralidharan et al. 2016¹⁴, the authors categorize the use of multiplexing in a two-way protocol as a “second generation without encoding” (2G-NC) scheme. This approach primarily aims to improve the probability of generating exactly one Bell pair between neighboring repeater stations (see Eq. (2) in subsection “Multiplexing with exactly one success—2G-NC”). It has been shown to be effective in regimes with moderate to low gate errors and low coupling efficiency. To benchmark the MTP considered in the manuscript, we have conducted a comparative analysis using the 2G-NC protocol (see Results). Consistent with the formulation in ref. 14, we do not incorporate entanglement distillation into the 2G-NC scheme (see subsection “Distillation”).

Multiplexing for more than one success. In this manuscript, we primarily focus on multiplexing schemes that allow for the creation of multiple Bell pairs across segments. The emphasis lies in using multiplexed channels to generate multiple elementary links and not just to boost the success probability of single elementary link generation. This effectively reduces the inefficiencies associated with the basic multiplexing scheme. In our setup, we consider generating these multiplexed links across segments at the same time or with an insignificant time delay. By insignificant time delay, we mean that the time difference between entangled photons arriving at the midpoint station is significantly smaller than the elementary link propagation delay and memory decoherence times, as may be achieved through spatial, spectral, or

time-division multiplexing. This corresponds to an idealized synchronization of photon arrival times at the midpoint station. While this assumption simplifies the analysis and is common in theoretical models^{14,35,36}, achieving such synchronization experimentally remains challenging. However, recent progress in photonic memory and emission timing control^{37,38} indicates that synchronized operation is a realistic target for near-future implementations.

This enhancement allows us to balance the qubit resources required in one-way repeaters vis-à-vis two-way repeaters and to more accurately estimate end-to-end performance by tracking the full probability distribution of successful links. While multiplexing can enhance delivery rates, the number of end-to-end Bell pairs that can be delivered decreases as the number of segments increases in a linear relay network. Figure 2 shows the expected number of end-to-end Bell pairs that can be delivered in a single shot for a linear quantum relay network with deterministic swapping operations, for varying numbers of segments. The yellow dashed line denotes an approximate number of end-to-end Bell pairs by the quantity $M \cdot \pi_0$, where M is the multiplexed channels and π_0 is the elementary link success probability. This quantity has been used by some prior analyses²⁵ as an upper bound for the expected Bell pairs that a quantum relay network can deliver. As evident from Fig. 2, keeping track of the probability distribution allows us to provide a more precise expectation of output Bell pairs than the models considered in prior works.

Another important consideration is that end-to-end links created using a relay approach will potentially suffer a decay in fidelity owing to swapping operations—we address this in subsection “Link propagation”. To deliver as many high-fidelity end-to-end Bell pairs as possible, one may need to consider distillation operations (see subsection “Distillation”). However, distillation schemes like DEJMPS are inherently probabilistic with success rates dependent upon the fidelity of the input Bell pairs. This makes an exact analysis difficult. To calculate and optimize repeater schemes, it is important to determine the probability distribution of the number of successfully distilled pairs at each step. We model the number of available Bell pairs as a random variable, with a distribution affected by non-deterministic operations (see Eqs. (1) and (10) in subsection “Recursive formulation of the probability distribution”).

In subsection “Recursive formulation of the probability distribution,” we present a recursive formulation that builds upon the idea of tracking probability distributions and incorporates additional elements such as distillation (see subsection “Distillation”) and nested swapping (see subsection “Link propagation”). This formulation provides a more comprehensive framework for modeling the dynamics of Bell pair generation, distillation,

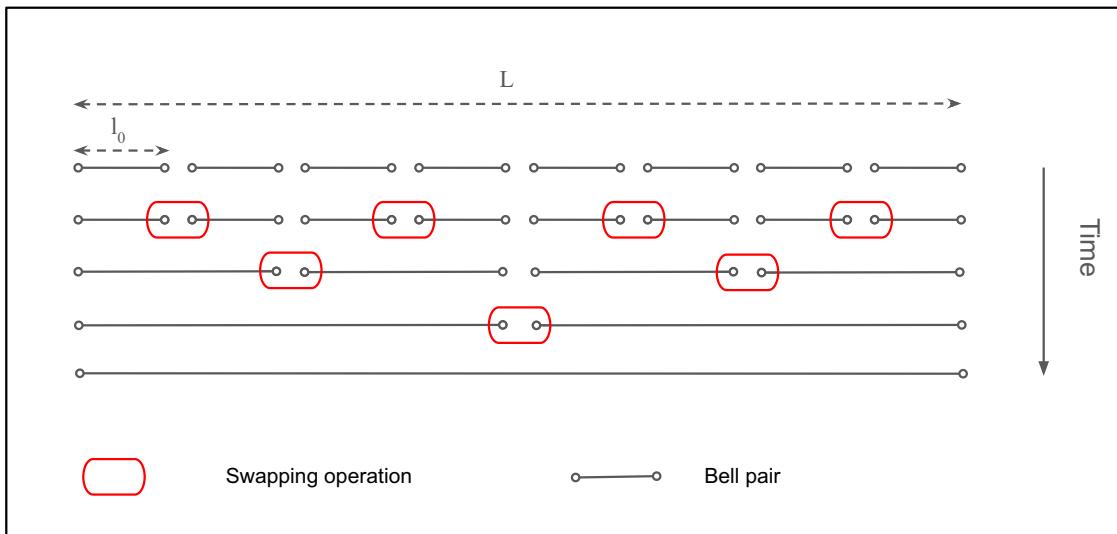


Fig. 3 | Schematic of a nested swapping protocol. We start with N elementary segments, each of length l_0 . The red boxes represent a swapping operation. Based on the decision variable \mathcal{D}_i , distillation may be performed before a swapping operation.

Each swapping operation doubles the length of the link. This process is repeated until an end-to-end link is established.

and swapping, allowing for a detailed analysis of the overall system performance. By integrating these elements, our recursive approach offers a significant improvement over previous models, enabling more accurate predictions and better optimization of quantum communication protocols.

Elementary link generation. In this setup, we use a meet-in-the-middle protocol with spatial (or spectral) multiplexing. Each repeater has an ensemble of emitters located on either side. This ensemble has a generation frequency ν , where it generates M photons entangled with M emitters in ν^{-1} time. This generation cycle of producing M multiplexed pairs in ν^{-1} time is referred to as a *burst*. One photon from each entangled photon pair thus generated is then sent to a Bell state analyzer located exactly midway between any two repeater stations, hereinafter referred to as the midpoint analyzer. We further assume emissions are synchronized across the chain, and the photons arrive at the midpoint station at the same time, and are able to effectively remove the “which-path-information.” The midpoint analyzer then performs parallel BSMs on all M incoming photons from each side with a success probability of $1/2$, and communicates the result to both repeater stations. There is also a need for an optical switch to separate the various multiplexing modes to the respective detectors at the analyzers. This switching operation could be lossy depending on the multiplexing or the detectors used. Recent advances in detection technology³⁹ have shown promise for building large detector arrays with spatial resolution that can potentially allow us to forego the need for optical switches. While an explicit analysis has not been done, our proposed protocol is also compatible with a midpoint source (MPS) scheme like Zero-Added Loss Multiplexing³¹, and will provide similar results. We also assume that time-bin dual-rail encoding is used for each multiplexed channel, primarily since it allows for protection against depolarization of the photon in the channel. However, if a polarization-based encoding is used, the elementary link generation equations will have to be updated to accommodate relevant noise models.

Link propagation. Two-way repeaters use a swap operation for link propagation. A swap operation involves performing a controlled-NOT (CNOT) gate on the halves of two Bell pairs situated at a middle repeater, and measuring the involved qubits at the middle repeater to create a longer link. Depending upon the technology used, this swap operation may be probabilistic or deterministic; however, for simplicity, we only consider deterministic swapping operations in our setup.

In our protocol, we use a nested swap strategy based on the Innsbruck Protocol^{40,41}. The Innsbruck protocol involves a series of entanglement swaps where qubits initially entangled with nearby nodes are used to establish entanglement with more distant nodes through intermediary swaps. This yields a nested structure where the network is divided into $N = 2^n$ segments (see Fig. 3). This nested swap procedure allows for the establishment of long-range entanglement connections between nodes that are not directly adjacent. By recursively applying entanglement swap operations, the protocol facilitates the generation of entangled links across the entire network. Prior works have studied swapping schedules other than nested swapping, like Swap ASAP^{42,43}, sequential generation and swapping^{43,44}, hybrid strategies⁴⁵, among others. However, nested schemes perform better than several other swapping schemes, especially in settings where the swapping operations are probabilistic, and repeaters are equipped with distillation, by providing an entanglement distribution rate that decays polynomially rather than exponentially in distance³⁵. Moreover, a nested swapping schedule allows for node synchronization for generation, swapping, and distillation operations. It is because of these reasons, coupled with an ease of analysis, that we have chosen a nested swapping schedule for our protocol.

Our protocol diverges from the standard Innsbruck protocol in two ways—(1) distillation may or may not be performed at each level depending on the expected quality of the Bell states (see subsection “Recursive formulation of the probability distribution”) (2) all links are created in parallel with multiple links shared between two stations, in a single burst with no interaction between bursts. In our protocol, we consider swaps as deterministic operations that allow us to save on the associated classical communication time costs. However, it is important to note that unless one has perfect elementary links, swaps, even in the case of perfect gate and measurement operations, cause an exponential decay in the fidelity with each swap⁴⁶. In the absence of a means to improve fidelity, this exponential decay makes quantum networks based solely on swaps impractical for long-range communications. Subsections “Gate operations” and “Measurement operations” in the Supplementary Methods explain the models used for gate and measurement operations, and Supplementary Eq. (9) has been used for modeling the swapping operations.

Distillation. The degradation in entanglement fidelity due to imperfect operations (such as swaps) or imperfect memories used for storing the entangled qubits can be mitigated by using distillation. In a distillation operation, a larger number (n) of Bell pairs is transformed to achieve a

smaller number (k) of higher fidelity Bell pairs. While significant improvements have been made in the field of distillation^{47–50}, in this paper, for ease of analysis, we use one of the basic schemes called the DEJMPS protocol ($n = 2$, $k = 1$)⁵¹ (see Supplementary Eq. (7) in the “Distillation” subsection in Supplementary Methods). The DEJMPS protocol is typically employed in an iterative fashion until a threshold fidelity is achieved or a threshold number of Bell pairs have been exhausted to create a higher fidelity pair. However, since most of our evaluation focuses on high-input-fidelity Bell pairs, we have considered only a maximum of a single round of distillation before swapping. Our protocol is easily modified to account for multiple rounds of distillation, i.e., performing as many distillation rounds until the fidelity of all available Bell pairs reaches above the threshold fidelity, or we run out of multiplexed Bell pairs available.

A critical consideration is the temporal overhead introduced by executing DEJMPS. In deterministic distillation protocols, classical communication is needed to relay Pauli correction information between parties. However, this exchange does not introduce latency, allowing operations to proceed without delay. In contrast, probabilistic distillation protocols impose stricter timing requirements, as the success or failure of the distillation must be communicated to determine subsequent actions. This requirement can create a significant bottleneck in two-way architectures. To address these temporal costs, one approach is to operate in a “blind” mode⁴¹. However, in our protocol, the distillation operations are conducted with “informed” decisions^{20,25}.

Deterministic distillation protocols might offer better performance than probabilistic protocols like DEJMPS on metrics such as secret-key rates and memory usage over time. However, deterministic distillation protocols typically require a large number of input entangled states for creating a higher fidelity state. Another possibility is to utilize a combination of deterministic and probabilistic distillation schemes with high-fidelity output states.

Another important consideration in two-way repeaters is the decision whether to distill or not before performing a swap. Given a higher initial fidelity, it is possible to perform multiple swaps and increase the length of the link before the fidelity drops to a level where distillation might be required. For a nested two-way protocol, this decision is made at each nesting level, determining whether one or multiple rounds of distillation are required to deliver Bell pairs end-to-end. Typically, this decision is not reactive and does not depend on real-time measurement outcomes or classical information from other network nodes. Instead, it is pre-computed prior to run-time based on factors such as the expected fidelity, number of multiplexed channels available, operational noise, and the number of repeaters.

As discussed in subsection “Multiplexing for more than one success,” and further explained in subsection “Recursive formulation of the probability distribution,” it is possible to keep track of the probability distribution of the number of Bell pairs. This probability distribution can be further optimized over the decision to distill based on a key service metric (such as the secret-key rate per shot or fidelity threshold). In the proposed protocol, this decision parameter is pre-determined for all nested levels and is taken to be a static network-wide agreement. This decision can be made using any rule that might be suitable for the application and the metric to be optimized. As examples, we have included two case scenarios in our analysis for deciding when to perform distillation—(1) comparing expected secret-key rates with and without distillation for a chain with 2^{n-i} segments, where i is the nesting level, and $n = \log_2 N$ for an N elementary segment linear network (see Eq. (12) in subsection “Recursive formulation of the probability distribution”). This is primarily driven by the fact that secret-key rate combines both throughput and fidelity into a single metric, making it a useful metric to optimize. This rule has been referred to in this manuscript as the SKR rule. (2) comparing the fidelity of links to a pre-determined fidelity threshold, where the decision to distill is contingent upon the link quality being less than the threshold (referred to in this manuscript as the F_{th} rule).

This approach could potentially be useful in a scenario where the quality of links above a certain threshold is desired as a service metric⁴². Both of these policies can be further optimized with an objective to maximize the end-to-end secret-key rate or number of Bell pairs while considering various possible distillation schedules and protocols. Furthermore, since we consider distillation, we need to keep track of the Bell pairs sacrificed when a round of distillation is performed. For a general n' -to- k' distillation scheme, each successive distillation round reduces the number of available Bell pairs by a factor of at least $1/\lceil n'/k' \rceil$ —see Supplementary Eqs. (14) and (15) in the Supplementary Materials for details.

Quality of memory. The temporal costs associated with classical communication with distillation require Bell states to be held in long-lived memories that do not undergo significant decoherence. Degradation in memory quality is usually characterized by T_1 and T_2 times. The T_1 time denotes the thermal relaxation time—the time it takes for the excited state $|1\rangle$ to relax back to the ground state $|0\rangle$. The T_2 time is the dephasing time that captures the loss of coherence due to dephasing in a quantum memory. In this manuscript, we only consider dephasing noise (see the “Memory Decoherence” subsection in the Supplementary Methods in the Supplementary Materials).

Termination. The protocol concludes once one or more Bell pairs are successfully established between the end stations. However, there may be cases where repeater stations lack enough Bell pairs to perform distillation. We explore various termination strategies for such scenarios when the number of Bell pairs in a segment drops below a certain threshold (R_i) for any nesting level i . If the static distillation schedule—pre-determined based on a distillation rule such as a fidelity threshold F_{th} or the SKR rule outlined in this manuscript—requires at least one round of distillation at the current or higher nesting levels, the protocol must adapt accordingly. Using this framework, we propose three potential termination strategies:

- *Strategy 1:* Repeater stations in the affected segment send classical messages instructing Alice, Bob, and all intermediate repeaters to halt all operations related to the burst. As these messages propagate, the repeater stations release the memory resources associated with the burst.
- *Strategy 2:* A variation of the first strategy involves the repeaters performing entanglement swaps and notifying their counterpart stations to perform additional swaps without distillation. Here, counterparties refer to repeaters that share a Bell pair with the initiating repeater. This approach allows repeaters holding k links to proceed similarly, ultimately establishing k end-to-end links. However, because this strategy bypasses the static distillation schedule, the established links will likely be of lower quality.
- *Strategy 3:* Another option is to allow unaffected segments to proceed without interruption. Specifically, repeater stations can continue performing distillation and swapping operations in segments where the number of Bell pairs $k \geq R_i$. Repeater stations learn about the failure on a segment as and when they do, minimizing classical communication time compared to Strategy 1. Also, this strategy requires fewer resources than Strategy 2, reducing the need for gates and other operations, while freeing up memory for future bursts.

The threshold R_i can be set based on the minimum number of Bell pairs required for the chosen n' -to- k' distillation scheme (such that $R_i \geq n'$). Alternatively, this threshold may be optimized using the static distillation decision schedule, ensuring that termination does not occur at any nesting level. We have selected *Strategy 3* because it minimizes classical communication time, potentially avoids unnecessary gate and measurement operations, and frees up memory resources for future bursts. However, this strategy sets a lower bound on performance. While computationally challenging, a more efficient termination strategy that optimizes resource usage could be developed in future work.

Recursive formulation of the probability distribution

We consider a linear network with $N = 2^n$ links. Let $M = m2^{n+1}$ denote the number of multiplexed channels available at each elementary link in a single shot, with $m \geq 1$. Note, this assumption can be relaxed if distillation is not required on all levels, to M taken to be less than $m2^{n+1}$. We consider a nested pumping distillation protocol that performs at most one distillation round at each level.

Let Y_i denote the number of Bell pairs on a segment at level i , with $p_{i,k} = P(Y_i = k)$ denoting the probability of having exactly k Bell pairs at level i . Let π_0 denote the probability that a link-level Bell pair generation attempt succeeds, and let $\bar{\pi}_0 = 1 - \pi_0$. Then,

$$p_{0,k} = \binom{M}{k} \pi_0^k \bar{\pi}_0^{M-k}, \quad k = 0, 1, \dots, M. \quad (1)$$

Multiplexing with exactly one success—2G-NC. For 2G-NC, the focus is on creating at least one bell pair in all segments. Further, using a similar setup assumed in¹⁴, it is assumed that there are no distillation operations, and since swapping operations are deterministic, we perform a network-wide swap simultaneously on the single link created, thus creating an end-to-end bell pair with a success probability,

$$\begin{aligned} \text{Success Prob.}_{(2\text{GNC})} &= \left(\sum_{k=1}^M p_{0,k} \right)^N, \quad k = 0, 1, \dots, M \\ &= \left(\sum_{k=1}^M \binom{M}{k} \pi_0^k \bar{\pi}_0^{M-k} \right)^N, \text{ using Eq. (1)} \\ &= \left(1 - p_{0,0} \right)^N = \left(1 - \bar{\pi}_0^M \right)^N. \end{aligned} \quad (2)$$

Multiplexing with more than one success, along with distillation. In order to calculate $p_{i,k}$, we first determine the effect of a distillation operation at level i whenever it is performed. This is captured by $q_{i,k}$,

$$q_{i,k} = P(X_i = k) = \sum_{j=2k}^{M/2^i} p_{i,j} \binom{\lfloor j/2 \rfloor}{k} d_i^k \bar{d}_i^{\lfloor j/2 \rfloor - k}, \quad (3)$$

$$i \in \{0, \dots, n\}; k = 0, 1, \dots, \lfloor M/2^{i+1} \rfloor,$$

where X_i denotes the number of Bell pairs produced by one distillation step at level i when performed, and d_i is the probability of a successful distillation step. This equation can be extended to reflect the case when no distillation is performed at that level. That is,

$$q_{i,k} = P(X_i = k) = \begin{cases} \sum_{j=2k}^{M_i} p_{i,j} \binom{\lfloor j/2 \rfloor}{k} d_i^k \bar{d}_i^{\lfloor j/2 \rfloor - k}, & \mathcal{D}_i = 1, \quad i = 0, \dots, n; k = 0, 1, \dots, \lfloor M_i/2 \rfloor, \\ p_{i,k}, & \mathcal{D}_i = 0 \end{cases} \quad (4)$$

where $M_i = \lfloor M/2^{\sum_{j=0}^{i-1} \mathcal{D}_j} \rfloor$ for $i > 0$, and \mathcal{D}_i is the indicator function for distillation,

$$\mathcal{D}_i = \begin{cases} 0, & \text{if no distillation at level } i \\ 1, & \text{if distillation at level } i \end{cases}, \quad i = 0, \dots, n.$$

To note, for the current analysis, we have a static value of $\mathcal{D}_i = 0$, when $i = n$. Now,

$$p_{i,k} = \begin{cases} (q_{i-1,k})^2 + 2q_{i-1,k} \sum_{j=k+1}^{M_i} q_{i-1,j}, & \mathcal{D}_i = 1 \\ (p_{i-1,k})^2 + 2p_{i-1,k} \sum_{j=k+1}^{M_i} p_{i-1,j}, & \mathcal{D}_i = 0 \end{cases}, \quad i = 1, \dots, n; k = 0, 1, \dots, M_i. \quad (5)$$

We now consider a protocol that terminates whenever $Y_i < R_i$ where R_i denotes a termination threshold for each level i segment. Let,

$$\begin{aligned} r_i &= \text{Probability of reset at level } i \text{ conditioned on reaching level } i, \\ p'_{i,k} &= P(Y_i = k \mid \text{no reset at levels } 0, 1, \dots, i), \\ q'_{i,k} &= \text{Probability of having } k \text{ distilled pairs conditioned on reaching level } i. \end{aligned}$$

Now,

$$r_0 = \sum_{j=0}^{R_0-1} \binom{M}{j} \pi_0^j \bar{\pi}_0^{M-j}, \quad (6)$$

and

$$p'_{0,k} = \begin{cases} \left(\binom{M}{k} \pi_0^k \bar{\pi}_0^{M-k} \right) / (1 - r_0), & k \geq R_0, \\ 0, & k < R_0 \end{cases}, \quad \forall k \in \{0, 1, \dots, M\}. \quad (7)$$

Now, as defined earlier,

$$\begin{aligned} q'_{i,k} &= P(Y_i = k \mid \text{no reset at levels } 0, 1, \dots, i), \\ &= \begin{cases} \sum_{j=2k}^{M_i} p'_{i,j} \binom{\lfloor j/2 \rfloor}{k} d_i^k \bar{d}_i^{\lfloor j/2 \rfloor - k}, & \mathcal{D}_i = 1, \\ p'_{i,k}, & \mathcal{D}_i = 0 \end{cases}, \quad \forall k \in \{0, 1, \dots, \lfloor M_i/2 \rfloor\}. \end{aligned} \quad (8)$$

$$r_i = \begin{cases} \sum_{l=0}^{R_{i-1}-1} \left((q'_{i-1,l})^2 + 2q'_{i-1,l} \sum_{j=l+1}^{M_i} q'_{i-1,j} \right), & \mathcal{D}_i = 1, \\ 0, & \mathcal{D}_i = 0 \end{cases}, \quad \text{s.t. } R_i \leq R_{i-1} \quad \forall i \in \{1, \dots, n\}. \quad (9)$$

Now,

$$p'_{i,k} = \begin{cases} \left((q'_{i-1,k})^2 + 2q'_{i-1,k} \sum_{j=k+1}^{M_i} q'_{i-1,j} \right) / (1 - r_i), & \mathcal{D}_i = 1 \quad \& \quad k \geq R_i, \\ (p'_{i-1,k})^2 + 2p'_{i-1,k} \sum_{j=k+1}^{M_i} p'_{i-1,j}, & \mathcal{D}_i = 0 \quad \& \quad k < R_i, \\ 0, & \forall k \in \{0, \dots, M_i\}. \end{cases} \quad (10)$$

Now, let f_i be the probability of reset at level i ,

$$f_i = \begin{cases} 1 - (1 - r_0)^N, & i = 0 \\ (1 - (1 - r_i)^{\frac{N}{2^i}}) \prod_{j=0}^{i-1} (1 - r_j)^{\frac{N}{2^j}}, & \text{otherwise.} \end{cases} \quad (11)$$

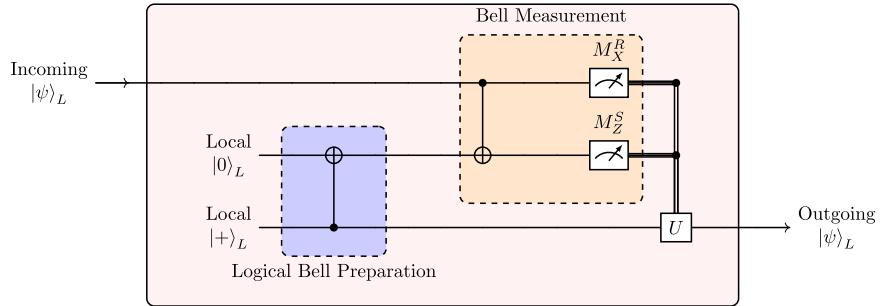
The decision to distill \mathcal{D}_i can be computed using any criterion best suited for the application. As examples, we have considered two conditions—(1)

$$\begin{cases} \mathcal{D}_i = 1 \\ F_{th} \geq F_i \text{ where } F_i \text{ is the fidelity at the } i\text{th level and } F_{th} \text{ is a pre-decided} \\ \mathcal{D}_i = 0 \end{cases}$$

threshold fidelity, (2) $SKR_i^{\uparrow} \geq SKR_i$, where SKR_i is the secret-key rate at level i without distillation, and SKR_i^{\uparrow} is the secret-key rate at level i after distillation. Using Supplementary Eq. (12) in the Supplementary Materials, this SKR -based decision rule can be further elaborated as,

$$\begin{aligned} \mathcal{D}_i &= 1 \\ SKR_i^{\uparrow} &\geq SKR_i \\ \mathcal{D}_i &= 0 \\ &\quad \mathcal{D}_i = 1 \\ &\quad \text{i.e., } r_{\text{secure}}(\rho_i^{\uparrow}) \cdot \mathbb{E}(Y_i^{\uparrow}) \geq r_{\text{secure}}(\rho_i) \cdot \mathbb{E}(Y_i), \\ &\quad \mathcal{D}_i = 0 \end{aligned} \quad (12)$$

Fig. 4 | Teleportation-based error correction (TEC) circuit for the one-way quantum repeater architecture. R and S denote the incoming and locally prepared code blocks, respectively, each encoded using the same (n, m) quantum parity code (QPC). Syndrome measurements are denoted by $M_{\{X, Z\}}$. The incoming logical qubit $|\psi\rangle_L$, which may have suffered photon loss during transmission through the optical fiber, is coupled to a freshly prepared logical Bell pair via transversal CNOT operations and logical Bell measurements. The local logical states are prepared using a fault-tolerant encoding circuit. Based on the measurement outcomes, a Pauli correction—denoted by the unitary U —is applied to the untouched half of the Bell pair, yielding a clean outgoing logical state that preserves the original quantum information. This state is then downloaded onto photonic qubits and transmitted to the next repeater node.



where ρ_i is the two-qubit state shared between two parties before distillation at level i , Y_i^\uparrow denote the number of Bell pairs on a segment at level i after one round of distillation, ρ_i^\uparrow is the state after distillation at level i , $r_{\text{secure}}(\rho)$ denotes the secret-key fraction for a two-qubit state ρ , and

$$\begin{aligned} \mathbb{E}(Y_i) &= \prod_{j=0}^i (1 - r_j)^{\frac{N}{2^j}} \cdot \sum_{k=R_i}^{M_i} k \cdot p'_{i,k} \\ \mathbb{E}(Y_i^\uparrow) &= \prod_{j=0}^i (1 - r_j)^{\frac{N}{2^j}} \cdot \sum_{k=R_i}^{\lfloor M_i/2 \rfloor} k \cdot q'_{i,k}. \end{aligned} \quad (13)$$

To note, for both cases listed above, we do not distill when equality holds.

One-way repeater architecture

First proposed by Munro et al.^{13,52}, one-way repeater architectures use near-deterministic methods to handle loss and operational errors. One fundamental difference between this scheme and a two-way scheme is its requirement for only forward or one-way classical communication. This need for one-way classical communication can further be eliminated if recovery operations on the errors accumulated in the preceding segment are performed at each repeater, resulting in only forward flow of quantum information. To tackle errors, in a one-way scheme, the quantum state to be transmitted is encoded in a logical qubit (qudit) using several physical qubits (qudits). For our analysis, we only consider a Bell pair of which one qubit is encoded and transmitted, while the other stays with the initiating party. Depending upon the nature of errors, there are various quantum error correction codes that can be used to encode the Bell pair and protect against these specific anticipated errors. For our comparison, we focus on Quantum Parity Codes (QPCs)^{13,52}. QPCs are generalized Shor Codes, and are capable of supporting teleportation-based error correction (TEC). A general form of an (n, m) QPC encodes the logical qubits as $|0\rangle_L = (|+\rangle_L + |-\rangle_L)/\sqrt{2}$ and $|1\rangle_L = (|+\rangle_L - |-\rangle_L)/\sqrt{2}$, with

$$|+\rangle_L = \frac{1}{2^{n/2}} (|0\rangle^{\otimes m} + |1\rangle^{\otimes m})^{\otimes n}; |-\rangle_L = \frac{1}{2^{n/2}} (|0\rangle^{\otimes m} - |1\rangle^{\otimes m})^{\otimes n}.$$

QPCs can be used to recover any encoded state under erasure noise as long as the following two conditions are met—(1) at least one qubit must arrive for each sub-block; (2) at least one sub-block must arrive with no loss. QPCs are loss tolerant⁵⁴, which is particularly useful to counter erasure losses in the optical fiber. They can also be prepared fault-tolerantly¹⁵. This error tolerance makes QPCs well suited for the one-way architecture, which must correct for loss and gate errors in the absence of heralding or two-way feedback. The logical teleportation procedure implemented at each repeater

allows the encoded Bell state to be re-encoded and forwarded, while also correcting for errors introduced in the previous segment. This process proceeds recursively along the chain, enabling end-to-end entanglement to be generated without the need for any backward classical signaling.

To achieve this feed-forward functionality, a TEC scheme is used. In this scheme, each repeater node corrects errors from the preceding segment and prepares a clean, re-encoded logical state to be forwarded downstream. This process allows for pipelined, segment-wise error correction and preserves the fidelity of end-to-end entanglement across a chain of repeaters. The TEC protocol works by transferring the quantum state of an incoming encoded block to a freshly prepared logical qubit, using a combination of entanglement, measurement, and Pauli frame updates. The flow of operations at each repeater follows these four steps:

1. Photon loss detection: A quantum non-demolition measurement is first applied to the incoming logical qubit $|\psi\rangle_L$ to detect photon losses. This allows identification of missing physical qubits without disturbing the surviving ones.
2. Preparation of logical Bell pair: Locally, a fresh encoded Bell state is prepared using a $|+\rangle_L$ and a $|0\rangle_L$ block. One half of this pair will eventually carry the teleported state forward.
3. Entangling operations and logical Bell measurement: The surviving photons of the incoming state are coupled to the $|0\rangle_L$ block via transversal controlled-NOT (CNOT) gates. A logical Bell measurement is then performed: the incoming block is measured in the logical X -basis and the entangled block in the logical Z -basis, implemented via individual measurements on the physical qubits.
4. Correction via Pauli frame: Based on the measurement outcomes, an appropriate Pauli correction is applied—either physically or virtually—to the untouched $|+\rangle_L$ block. This restores the logical state and yields a high-fidelity outgoing qubit, ready for the next transmission segment.

A schematic circuit diagram of the TEC operation is shown in Fig. 4. This segment-wise error correction process repeats recursively across all repeater nodes, ensuring that errors do not accumulate along the chain. For a detailed explanation on syndromes in QPCs, see Namiki et al.⁵⁵. In our analysis, we assume the codes are prepared fault-tolerantly, and require the same setup as outlined in Muralidharan et al.¹⁵ and Namiki et al.⁵⁵. For simplicity, similar to the two-way protocol outlined earlier, we assume spatial multiplexing such that all incoming photons from a block arrive at the repeater at the same time or with negligible time delay. The scheme is compatible with spectral and time-bin multiplexing; however, in those cases, switching losses may need to be explicitly accounted for. Since error correction takes place locally at each repeater, the scheme places more demanding requirements on repeater placements, operation fidelities and speeds, particularly as the code size increases. Nonetheless, for future

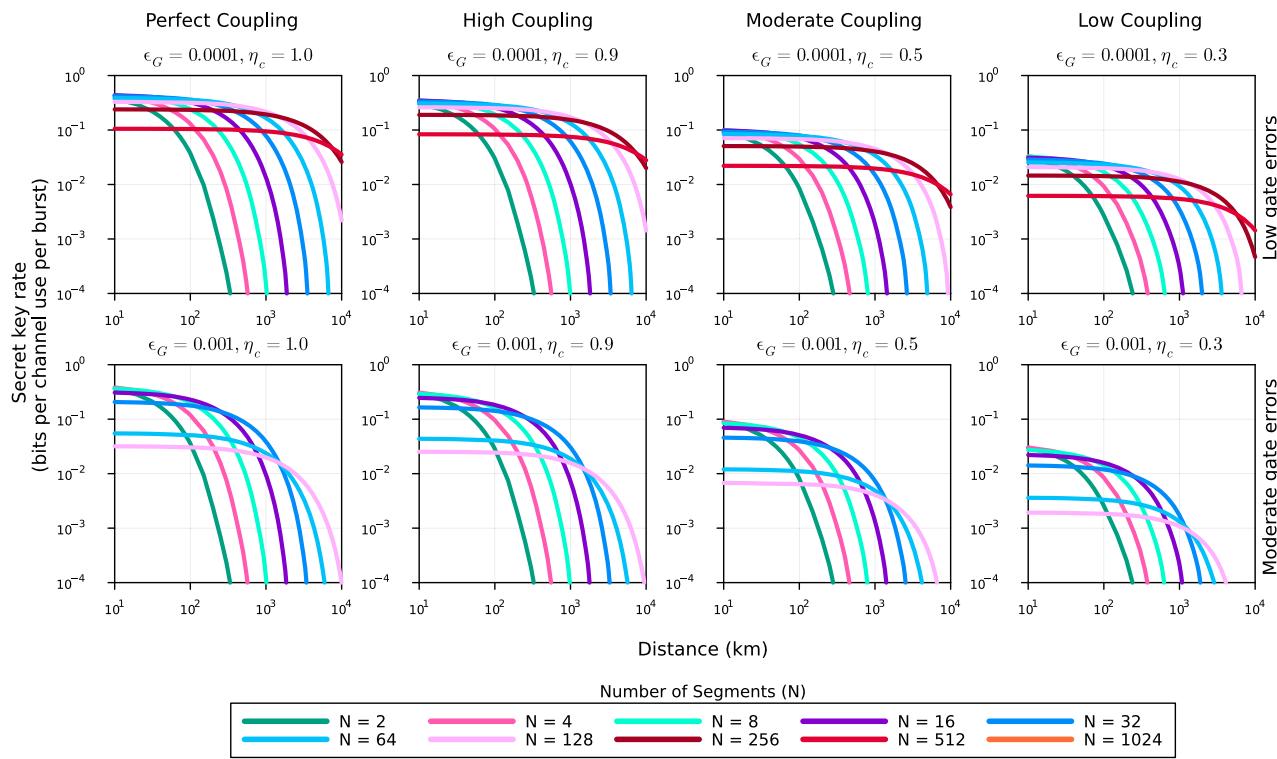


Fig. 5 | Performance of multiplexed two-way protocol (MTP) with distance using secret-key rate as the metric. The number of segments is shown in different colors and denoted by N . The plots in the top row consider a low gate error scenario with a gate error rate (ϵ_G) of 10^{-4} or 0.01%, and the bottom row plots show the performance with moderate gate errors ($\epsilon_G = 10^{-3}$ or 0.1%). The different columns show the performance in different coupling regimes, starting with a perfect coupling ($\eta_c = 1$),

with the coupling coefficient reducing when moving from left to right ($\eta_c \in \{1, 0.9, 0.5, 0.3\}$). In this setup, we used the protocol based on one-way BB84⁶⁹ Secret-key rate to inform the distillation decision-making process, allowing a maximum of one round of distillation at any level of nesting. Also, in this setup, no distillation is performed at the end level of nesting.

networks where high-quality operations can be realized, QPC-based one-way schemes offer excellent potential for low-latency transmission. Furthermore, similar to the strategy outlined in ref. 15, we use codes that deliver the highest key rate using the least number of qubits for our comparison.

Results

Parameter regime

We develop a model of our protocol in “Methods” and use it to compare its performance with that of the one-way scheme. As outlined in subsections “Distillation” and “Recursive formulation of the probability distribution” in “Methods,” we optimize the decision to perform distillation using a service metric (e.g., secret-key rate or fidelity threshold). Although the model yields entanglement delivery rate and average fidelity, we will use secret-key rate as our metric throughout this section. Our choice of secret-key rate as the primary performance metric is guided by the necessity of establishing a network capable of consistently delivering high-quality Bell pairs at a rapid pace. Secret-key rate combines fidelity (link quality) and entanglement delivery rate (quantity) into a single metric, making it a straightforward choice for measuring performance. See the Secret-key Rate subsection in the Supplementary Methods section of Supplementary Materials for details.

In Fig. 5, we show the secret-key rate of the MTP across different ranges of coupling coefficients, gate and measurement noise, for distances up to 10^4 km. We note that the coupling efficiency and gate errors affect secret-key rate in qualitatively different ways. Secret-key rate decreases as coupling efficiency decreases in a uniform manner over all segment lengths, whereas increasing gate errors asymmetrically affect more segmented networks.

In this section, we outline the parameter regime and the model assumptions for the results presented in subsections “Performance evaluation using secret-key rates” and “Comparison of resource costs,” where

we compare the performance of the MTP with one-way schemes. For our comparison, we have selected different QPCs optimized for specific distances for different parameter settings of coupling efficiency, gate, and measurement noise. Furthermore, in our analysis we mainly focus on the parameter regime, demonstrated in prior works to be advantageous for one-way repeater schemes; specifically, as identified by Muralidharan et al.¹⁴, this regime corresponds to coupling efficiency $\eta_c \geq 0.9$, gate error $\epsilon_G \leq 10^{-3}$, and gate time $t_G \leq 10^{-9}$ s.

For this comparison, keeping in line with the analysis in ref. 14, we assume high initial fidelity (computed as $F = 1 - 1.125\epsilon_G$), gate error $\epsilon_G \in \{10^{-4}, 10^{-3}\}$, measurement error $\xi = 0.25\epsilon_G$, coupling coefficient $\eta_c \in \{1, 0.9\}$. Moreover, in this analysis, we have chosen a realizable decoherence time with $T_2 = 1$ s. This parameter choice has been experimentally demonstrated in hardware platforms like trapped ions⁵⁶ and Rydberg atoms⁵⁷. In our analysis, we consider optimal architectures for both one- and two-way schemes. For the one-way scheme, for each distance, an optimal (n, m) QPC is chosen that minimizes the total number of qubits required to deliver a unit secret-key, with the search parameters constrained to $n \leq 70$, $m \leq 20$, and the inter-repeater spacing constrained between 1 and 4 km. For the multiplexed two-way scheme (MTP), a maximum of 1024 segments and 1024 multiplexed channels have been considered, primarily to limit computational costs. For 2G-NC, we use a numerical search for selecting the optimal number of multiplexed channels (ranging between 1 and 1024) that minimizes the total number of qubits required over the linear network to deliver a unit secret-key. Furthermore, to make this comparison, for both QPC and MTP, only the envelopes of the best performing configuration (in terms of inter-repeater spacing, and additionally for QPC the specific (n, m) code), evaluated using secret-key rate (SKR) per channel use per burst as the performance metric, have been considered.

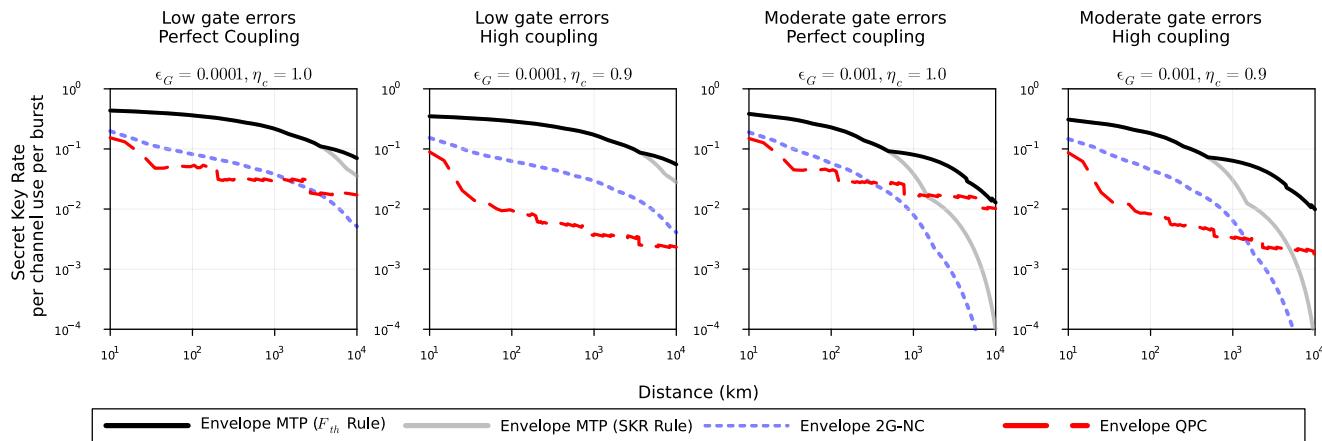


Fig. 6 | Performance comparison between one- and two-way schemes using the secret-key rate as the metric. The red dashed line shows the performance by the optimal Quantum Parity Codes (QPC), the blue dotted line represents the optimal performance for the non-encoded second generation scheme “2G-NC,” the solid lines are the envelope for the secret-key rates for multiplexed two-way scheme (MTP) with two different distillation rules with black solid line representing the fidelity threshold rule, and the gray solid line representing the SKR rule. To note, for the fidelity threshold rule, we use $F_{th} = 0.95$. For both MTP schemes and the 2G-NC,

the envelope has been taken over with a different number of elementary segments varying between 2 and 1024. For each distance, a specific (n, m) QPC is chosen, optimizing for the total number of qubits required with the search parameters constrained to $n \leq 70$, $m \leq 20$. For the MTP schemes and the 2G-NC, a maximum of 1024 multiplexed channels has been considered. Compared to the QPC and the 2G-NC, the MTP schemes deliver better secret-key rates per channel use per burst in all parameter regimes.

Summary of assumptions

- Our protocol assumes unconstrained availability of quantum emitters/memories at all repeaters. In our analysis, we treat memory availability as an upper-bound resource, characterized by the total number of memories per repeater ($2Mv$), where M denotes the number of multiplexing channels and v is the source frequency. This would reduce memory requirements considerably. Note that we require a memory availability corresponding to the worst-case scenario of all links succeeding and all other probabilistic events succeeding. In a nested distillation scheme, the actual number of memories in practice used per cycle depends on the repeater’s position in the chain, with repeaters located at half-way typically storing Bell pairs for longer durations, and storage times decreasing with each recursive midpoint. Furthermore, this assumes a worst-case scenario where all elementary links and distillation operators succeed. In practice, one would want to have enough memories to guarantee that all states can be kept with sufficiently high probability. Although we recognize that optimizations in swap scheduling could potentially reduce the memory overhead, our present analysis serves as an upper-bound study, with detailed memory scheduling optimizations deferred to future work.
- We assume high cooperativity for the cavity-enhanced memories. Cavities have shown promise in realization of quantum networks^{58,59}, enabling implementations of efficient multi-qubit gates^{60–62}, fast storage and readouts⁶³.
- We assume switching to be perfect, for both inter-memory connectivity at the repeater, and in case of spectral or temporal multiplexing, the switching required to connect incoming photons to the appropriate memory. Furthermore, we do not account for delays associated with performing CNOT gate operations between any two qubits. These delay costs can be non-trivial with current hardware technology, especially if the relative distance between the selected qubits on the register is large. The primary reason for not accounting for these delays is that we assume fast gate operations for both one-way and two-way schemes. Both one-way and two-way schemes will be proportionally hit by temporal costs associated with these two-qubit gate operations.
- We assume detectors are perfect and that the probability of success of a BSA is exactly 1/2 at the midpoint stations.

- We assume that detectors and quantum memories can be reset in a time smaller than the inverse of the source frequency v . We further assume that qubit readout times are sufficiently fast such that they are negligible with respect to the inverse of the source frequency v . This allows for pipeline operations where the only bottleneck is the source’s ability to generate bursts. This assumption is motivated by the requirement for the network to operate in a steady state. In our model, a continuous and robust stream of end-to-end Bell pairs is generated by ensuring that memories are quickly reset once a burst of operations has concluded. By requiring the reset time to be smaller than the inverse of the source frequency, we ensure ready availability of memories for the subsequent cycle, thereby reducing unnecessary memory overhead.
- We assume that the optical losses in the fiber in transit are the same for all frequencies in the case where spectral multiplexing is used. We further assume that the optical fiber does not contribute to any other form of noise except erasure. We assume the speed of light in fiber to be 200,000 km/s.
- We assume deterministic swap gates. Several proposals that use high cooperativity cavities have shown potential for achieving such gates^{60,64}.
- Memory decoherence time T_2 is assumed to be 1 s^{17–19}.
- We assume fast gate and measurement operations (gate and measurement time $t_G \ll 10^{-9}$ s).
- Measurement errors ξ are assumed to be a quarter of gate errors ϵ_G i.e., $\xi = 0.25\epsilon_G$ ¹⁴.
- Elementary link fidelity is estimated to be $1 - 1.25\epsilon_G$ ¹⁴ using depolarized states for elementary link Bell pairs.
- Fiber attenuation length has been taken as 20 km.

Costs

Prior works have mostly focused on memory-constrained regimes, and have considered memories as the most significant cost factor^{14,20,25}. However, promising developments in multiple hardware platforms since^{17,18} have weakened these assumptions. It is critical that better cost metrics be considered to evaluate the performance of different repeater architectures. While a detailed cost analysis (accounting for environmental noise, hardware requirements, labor, physical infrastructure, and software development and upkeep) would be the most appropriate approach, we believe the following high-level metrics can still serve as a guidepost for comparing quantum network deployments:

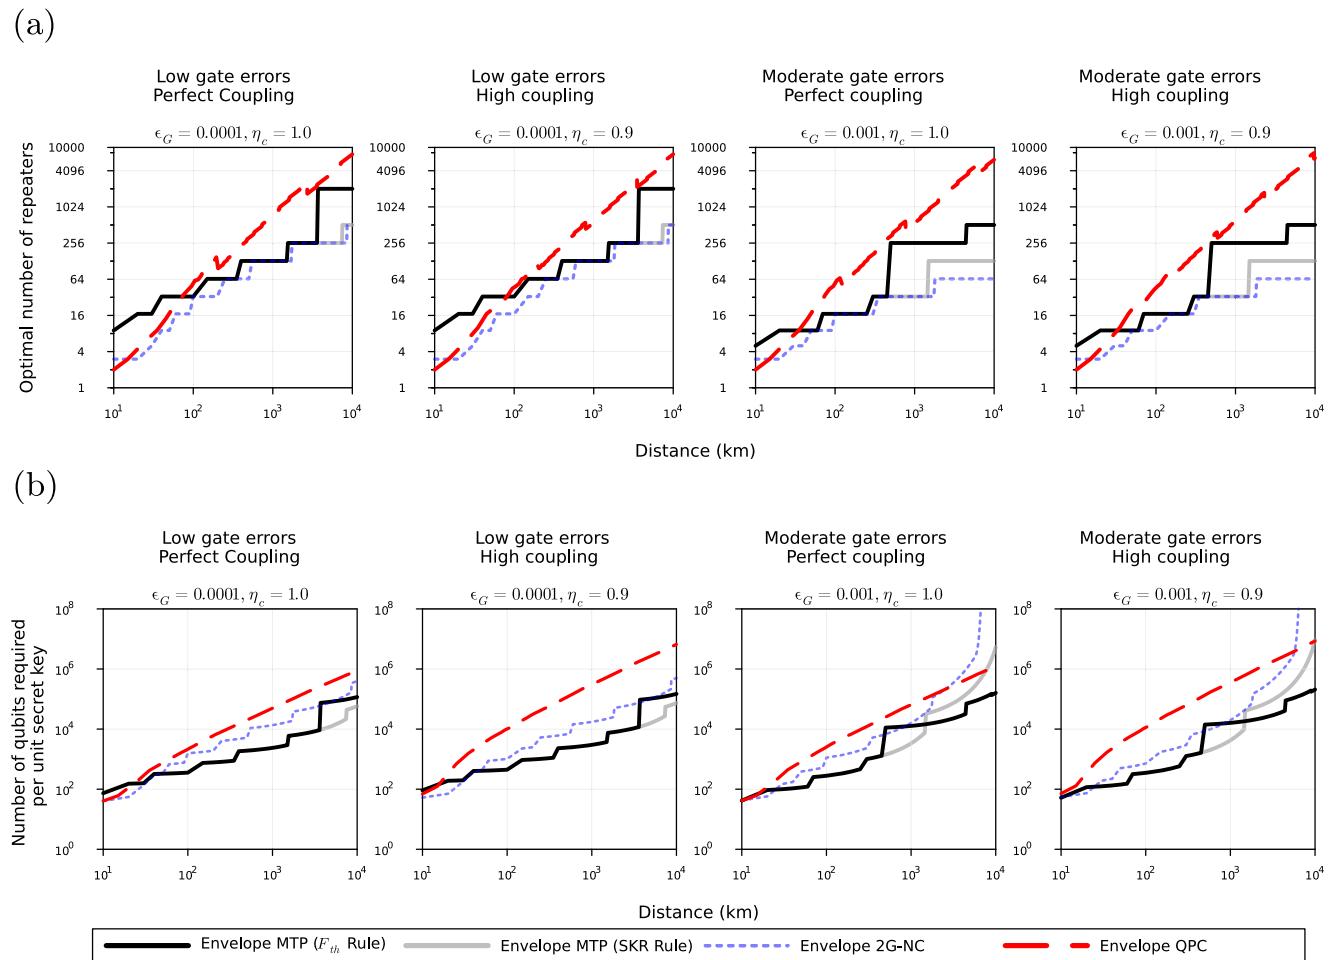


Fig. 7 | Comparing hardware costs between one- and two-way quantum repeaters.

Number of **a** repeaters, and **b** qubits required per burst for each unit secret-key delivered for optimal performance for one- and two-way repeater architectures. The red dashed line shows the number of repeaters required for the optimal Quantum Parity Code (QPC), the blue dotted line shows the number of repeaters required for optimal performance for 2G-NC, the black and the gray solid lines are the envelopes for the number of repeaters required for the optimal performing multiplexed two-way schemes (MTP) using a $F_{th} = 0.95$ and a SKR-based distillation decision rule

respectively. To note, we do not consider the ancilla qubits required for state preparation or teleportation-based error correction for QPC, and the estimation presented here is a lower bound. For all long-distance parameter regimes considered, the MTP requires significantly less number of repeaters and qubits than the QPC. Compared to the 2G-NC, the MTP (SKR rule) scheme requires a similar number of repeaters but less number of qubits for delivering a unit secret-key. To note, MTP using the F_{th} rule requires slightly more repeaters than the MTP based on the SKR rule, and the 2G-NC protocol, but a lower number of repeaters than the QPC.

- Cost of repeater installations, including acquisition of land and physical infrastructure, maintenance, temperature requirements, among others. This is captured in our metric of the number of repeaters required for delivery.
- Memory costs, including initialization costs and residence times. Here, residence time refers to the duration for which a memory remains engaged. Since one-way and two-way repeater architectures require vastly different types of quantum memories in the parameter regime of fast gate operations and readout, we capture these costs using the metric of the number of qubits required to deliver a single Bell pair.
- Number of 2-qubit gates, circuit size, and usage. In our analysis, we have only considered 2-qubit gates as the appropriate measure, since both QPC and two-way nested schemes will require 2-qubit gate operations.
- Number of measurement operations. In ref. 65, the authors identify the number of measurement operations as a potential candidate for evaluating link costs in the context of routing in quantum networks. The authors use simulations to establish the relationship between measurement count and overall network performance, highlighting this metric's potential for assessing resource consumption when determining the optimal path for data transmission.

- Cost of operating a repeater, including energy and ongoing maintenance. Although not explicitly considered in our current analysis, these operational costs are expected to play a significant role in practical deployments, and a detailed evaluation is left for future work.

Performance evaluation using secret-key rates

We use two different flavors of the MTP based on two different distillation decision rules—(1) a rule based on secret-key rate (2) a rule based on fidelity threshold (see subsection “Multiplexing with more than one success, along with distillation” in “Methods” and Eq. (12) for details), where we distill when the fidelity drops below the threshold value. For the fidelity threshold, we have used the threshold of 0.95, based on a visual search on a small set of threshold values (see Supplementary Note 1 in Supplementary Materials for performance plots using fidelity thresholds other than 0.95). As shown in Fig. 6, the MTP repeater schemes outperform the protocol based on the optimal QPCs for all considered parameter regimes. These differences in performance range between one to two orders of magnitude, depending on the gate errors and coupling efficiencies considered. In the case of moderate gate errors and long distances, understanding this gain in the context of associated costs, as analyzed in subsection “Comparison of resource costs” in Results, is important. We also find that the MTP outperforms the

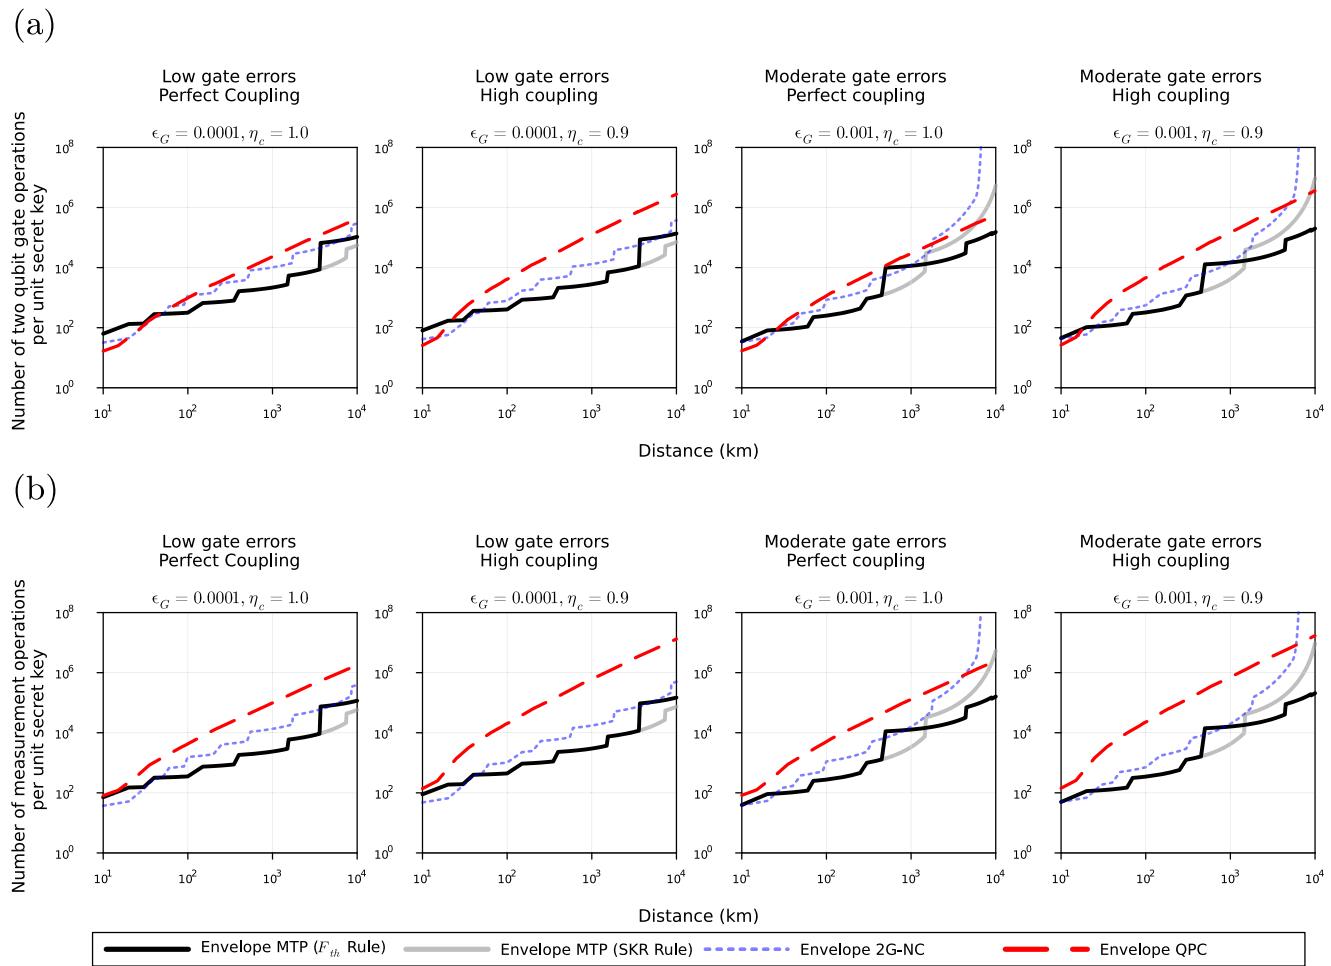


Fig. 8 | Comparing operation costs between one- and two-way quantum repeaters. Number of a two-qubit gates, and b measurement operations required per burst for each unit secret-key delivered for one- and two-way repeater architectures. The red dashed line shows the number of repeaters required for the optimal Quantum Parity Code (QPC), the blue dotted line shows the number of repeaters required for optimal performance for 2G-NC, the black and the gray solid lines are the envelopes for the highest deliverable secret-key rate for multiplexed two-way schemes (MTP)

using a $F_{th} = 0.95$ and a SKR-based distillation decision rule respectively. To note, for QPCs, we have not considered gate operations required for state preparation or gate operations on ancilla qubits, and the estimation presented here is a lower bound. For (almost) all parameter regimes considered, the optimal QPC-based protocol requires a higher number of two-qubit gates and measurement operations than the MTP.

multiplexing protocol aimed at single elementary link generation (referred to as 2G-NC in ref. 14) across the entire parameter regime considered in the manuscript. To note, both the SKR and the $F_{th} = 0.95$ rule are probably non-optimal, and an optimal distillation schedule can be achieved using a numerical search. Furthermore, since we have limited the number of distillation rounds to a maximum of one per nesting level, potential improvements in the performance could be made if this constraint is relaxed.

Comparison of resource costs

In this subsection, we compare different costs, i.e., the number of repeaters, the number of two-qubit gates, and the number of measurement operations. Figures 7 and 8 demonstrate that the resources required for the one-way schemes are significantly higher than the equivalently multiplexed two-way schemes. Fig. 7a compares the number of repeaters required for the optimal secret-key rates shown in Fig. 6. As shown in the plots, the QPC-based system requires a significantly larger number of repeaters compared to the multiplexed two-way (MTP) system. This difference in the required number of repeaters becomes more pronounced as imperfections in coupling and gate efficiency increase in the system. Note that 2G-NC requires a slightly less or equal number of repeaters compared to the MTP schemes in most parameter regimes. In Fig. 7b, we compare the qubit resources

required to deliver a unit secret-key for different gate and coupling efficiencies. In this analysis, we consider a lower bound on the number of qubit resources required for QPC, since we do not consider the ancilla qubits required for state preparation and teleportation-based error correction. We observe from the graphs that the QPC-based system requires more qubit resources for all parameter regimes considered. However, if ancilla qubits are included, it is likely that the two-way scheme will perform even better. It should also be noted that the number of qubits required per unit secret-key delivered has been used as the metric of comparison in Muralidharan et al.¹⁴. Using the number of qubits required as the sole cost metric, it might be straightforward to see the attractiveness of the MTP compared to the QPC and the 2G-NC protocol. Figure 8a, b presents the estimated number of measurements and 2-qubit gate operations per unit secret-key delivered required to maximize secret-key rate as a function of distance. As with the number of repeaters and qubits, we find that the QPC-based one-way scheme requires significantly more gate and measurement operations across most parameter regimes considered, with the possible exception of low gate errors for distances $\lesssim 50$ km. For all other considered parameter regimes, these differences in gates and measurement costs range between one and two orders of magnitude, with the MTP posing lower resource requirements compared to both QPC and the 2G-NC.

Conclusion

The rapid development of quantum technologies has spurred efforts to establish robust quantum networks. The choice of repeater architecture significantly impacts the scalability and reliability of these networks. A comprehensive comparison between different repeater architectures is essential to understand their strengths and weaknesses under varying conditions, including error rates, resource availability, and communication latency. Such an analysis can guide the design of practical quantum networks by highlighting where specific architectures excel or falter, and providing insights into the trade-offs between performance and technological complexity.

Pioneering work by Muralidharan et al.¹⁴ compared one-way and two-way schemes, identifying parameter regimes where each scheme could be advantageous. However, Muralidharan's setup assumes a memory-constrained regime and does not utilize the full power of multiplexing. Studies considering multiplexing have focused on maximizing the success probability of a single elementary link or have not incorporated nested purification^{14,21–23}. Where such considerations have been made, distillation operations have often been assumed deterministic^{20,25}. This study aims to clarify the performance expectations of multiplexed two-way and one-way repeater architectures, providing a framework to make informed decisions when selecting the optimal architecture based on application requirements.

In this manuscript, we consider a two-way protocol that leverages the power of multiplexing with an application-aware decision parameter for distillation. Additionally, we present a thorough evaluation of performance differences between one-way and MTPs using relevant metrics such as the secret-key rate, number of repeaters, qubits, and gate and measurement operations. Focusing on the regime identified in prior work as favorable to one-way schemes, we demonstrate that the multiplexed two-way repeater scheme, in an unconstrained memory regime, outperforms one-way schemes even under conditions previously believed to favor the latter. Furthermore, these performance gains can be realized with lower resource requirements, making two-way schemes a more attractive alternative.

While our findings suggest that multiplexed two-way schemes are potentially a near-universal choice across various parameter regimes, the performance achieved in our analysis may be sub-optimal. Our study focused on basic protocols and requires further exploration to identify additional areas for improvement. For instance, our current analysis does not utilize any adaptive mechanisms at the link level for decision-making to optimize overall performance. Additionally, we only considered a basic DEJMPS protocol with a maximum of a single round performed at any nested level. These simplifying assumptions preserve analytical clarity but likely leave room for significant performance improvements. More advanced distillation schemes^{50,66,67} could improve performance while simultaneously lowering resource requirements. Moreover, our distillation scheduling may not be optimal and could be improved to enhance performance and reduce costs. We also assume perfect synchronization of emissions and neglect delays due to local gate connectivity, which may require further attention in experimental implementations.

Our choice of physical parameters—such as long memory coherence times, efficient photon coupling, and high-fidelity gates—reflects recent experimental achievements in platforms like diamond NV centers, silicon T-centers, and other solid-state systems^{37,38,68}. Although engineering all these capabilities simultaneously in a scalable system remains challenging, they are not mutually incompatible. Therefore, our parameter set represents an optimistic but feasible near-term scenario that has been partially realized in current experimental platforms.

Building on the framework proposed in this study, future work can focus on improving network performance by introducing optimizations in distillation scheduling or selecting more efficient distillation protocols. Extending this framework to an asynchronous setup or a connectionless protocol might provide interesting insights and potential improvements. Another important direction is to consider architectures involving free-space or satellite-based links, which may differ in coupling efficiencies, loss characteristics, and synchronization challenges. A detailed exploration of hybrid terrestrial-space

quantum networks would help assess the practicality of repeater architectures in global-scale scenarios and is a compelling direction for future study. We recommend exploring these extensions in future research.

Data availability

All the data presented in this paper can be generated with a customized code (see “Code availability” statement).

Code availability

The code used to generate these data is available from the corresponding author upon reasonable request.

Received: 9 December 2024; Accepted: 5 July 2025;

Published online: 16 July 2025

References

1. Degen, C., Reinhard, F. & Cappellaro, P. Quantum sensing. *Rev. Mod. Phys.* **89**, 035002 (2017).
2. Eldredge, Z., Foss-Feig, M., Gross, J. A., Rolston, S. L. & Gorshkov, A. V. Optimal and secure measurement protocols for quantum sensor networks. *Phys. Rev. A* **97**, 042337 (2018).
3. Ge, W., Jacobs, K., Eldredge, Z., Gorshkov, A. V. & Foss-Feig, M. Distributed quantum metrology with linear networks and separable inputs. *Phys. Rev. Lett.* **121**, 043604 (2018).
4. Guo, X. et al. Distributed quantum sensing in a continuous-variable entangled network. *Nat. Phys.* **16**, 281–284 (2020).
5. Caleffi, M. et al. Distributed quantum computing: a survey. *Computer Networks* **254**, 110672 (2024).
6. de Andrade, M. G., Panigrahy, N. K., Dai, W., Guha, S. & Towsley, D. Universal quantum walk control plane for quantum networks. Preprint at <http://arxiv.org/abs/2307.06492> (2023).
7. Pirandola, S. et al. Advances in quantum cryptography. *Adv. Opt. Photonics* **12**, 1012 (2020).
8. Pan, D. et al. The evolution of quantum secure direct communication: On the road to the qinternet. *IEEE Commun. Surv. Tutor.* **26**, 1898–1949 (2024).
9. Hillery, M., Bužek, V. & Berthiaume, A. Quantum secret sharing. *Phys. Rev. A* **59**, 1829–1834 (1999).
10. Markham, D. & Sanders, B. C. Graph states for quantum secret sharing. *Phys. Rev. A* **78**, 042309 (2008).
11. Pirandola, S., Laurenza, R., Ottaviani, C. & Banchi, L. Fundamental limits of repeaterless quantum communications. *Nat. Commun.* **8**, 15043 (2017).
12. Takeoka, M., Guha, S. & Wilde, M. M. Fundamental rate-loss tradeoff for optical quantum key distribution. *Nat. Commun.* **5**, 5235 (2014).
13. Munro, W. J., Azuma, K., Tamaki, K. & Nemoto, K. Inside quantum repeaters. *IEEE J. Sel. Top. Quantum Electron.* **21**, 78–90 (2015).
14. Muralidharan, S. et al. Optimal architectures for long distance quantum communication. *Sci. Rep.* **6**, 20463 (2016).
15. Muralidharan, S., Kim, J., Lütkenhaus, N., Lukin, M. D. & Jiang, L. Ultrafast and fault-tolerant quantum communication across long distances. *Phys. Rev. Lett.* **112**, 250501 (2014).
16. Niu, D., Zhang, Y., Shabani, A. & Shapourian, H. All-photonic one-way quantum repeaters with measurement-based error correction. *npj Quantum Inf.* **9**, 1–9 (2023).
17. The Next Decade in Quantum Computing—and How to Play. <https://www.bcg.com/publications/2018/next-decade-quantum-computing-how-play> (2020).
18. Wang, P. et al. Single ion qubit with estimated coherence time exceeding one hour. *Nat. Commun.* **12**, 233 (2021).
19. Dudin, Y. O., Li, L. & Kuzmich, A. Light storage on the time scale of a minute. *Phys. Rev. A* **87**, 031801 (2013).
20. Razavi, M., Piani, M. & Lütkenhaus, N. Quantum repeaters with imperfect memories: cost and scalability. *Phys. Rev. A* **80**, 032301 (2009).

21. Guha, S. et al. Rate-loss analysis of an efficient quantum repeater architecture. *Phys. Rev. A* **92**, 022357 (2015).
22. Dhara, P., Patil, A., Krovi, H. & Guha, S. Subexponential rate versus distance with time-multiplexed quantum repeaters. *Phys. Rev. A* **104**, 052612 (2021).
23. Dhara, P., Linke, N. M., Waks, E., Guha, S. & Seshadreesan, K. P. Multiplexed quantum repeaters based on dual-species trapped-ion systems. *Phys. Rev. A* **105**, 022623 (2022).
24. Collins, O. A., Jenkins, S. D., Kuzmich, A. & Kennedy, T. A. B. Multiplexed memory-insensitive quantum repeaters. *Phys. Rev. Lett.* **98**, 060502 (2007).
25. Razavi, M., Thompson, K., Farmanbar, H., Piani, M. & Lütkenhaus, N. Physical and architectural considerations in quantum repeaters. In *Proc. Quantum Communications Realized II* Vol. 7236, 18–30. <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7236/723603/Physical-and-architectural-considerations-in-quantum-repeaters/10.1117/12.811880.full> (SPIE, 2009).
26. Huie, W., Menon, S. G., Bernien, H. & Covey, J. P. Multiplexed telecommunication-band quantum networking with atom arrays in optical cavities. *Phys. Rev. Res.* **3**, 043154 (2021).
27. Rozpędek, F., Seshadreesan, K. P., Polakos, P., Jiang, L. & Guha, S. All-photonic Gottesman-Kitaev-Preskill-qubit repeater using analog-information-assisted multiplexed entanglement ranking. *Phys. Rev. Res.* **5**, 043056 (2023).
28. Li, B., Goodenough, K., Rozpędek, F. & Jiang, L. Generalized quantum repeater graph states. *Phys. Rev. Lett.* **134**, 190801 (2025).
29. Munro, W. J., Harrison, K. A., Stephens, A. M., Devitt, S. J. & Nemoto, K. From quantum multiplexing to high-performance quantum networking. *Nat. Photonics* **4**, 792–796 (2010).
30. Piparo, N. L., Hanks, M., Gravel, C., Nemoto, K. & Munro, W. J. Resource reduction for distributed quantum information processing using quantum multiplexed photons. *Phys. Rev. Lett.* **124**, 210503 (2020).
31. Chen, K. C. et al. Zero-added-loss entangled-photon multiplexing for ground- and space-based quantum networks. *Phys. Rev. Appl.* **19**, 054029 (2023).
32. Piparo, N. L., Munro, W. J. & Nemoto, K. Quantum multiplexing. *Phys. Rev. A* **99**, 022337 (2019).
33. Nishio, S., Lo Piparo, N., Hanks, M., Munro, W. J. & Nemoto, K. Resource reduction in multiplexed high-dimensional quantum Reed-Solomon codes. *Phys. Rev. A* **107**, 032620 (2023).
34. Nishio, S. et al. Multiplexed quantum communication with surface and hypergraph product codes. *Quantum* **9**, 1613 (2025).
35. Duan, L.M., Lukin, M. D., Cirac, J. I. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414**, 413–418 (2001).
36. Sangouard, N., Simon, C., de Riedmatten, H. & Gisin, N. Quantum repeaters based on atomic ensembles and linear optics. *Rev. Mod. Phys.* **83**, 33–80 (2011).
37. Davidson, O., Yoge, O., Poem, E. & Firstenberg, O. Single-photon synchronization with a room-temperature atomic quantum memory. *Phys. Rev. Lett.* **131**, 033601 (2023).
38. Zhang, S. et al. Realization of a programmable multipurpose photonic quantum memory with over-thousand qubit manipulations. *Phys. Rev. X* **14**, 021018 (2024).
39. Miki, S., Yamashita, T., Wang, Z. & Terai, H. A 64-pixel NbTiN superconducting nanowire single-photon detector array for spatially resolved photon detection. *Opt. Express* **22**, 7811 (2014).
40. Dür, W., Briegel, H.-J., Cirac, J. I. & Zoller, P. Quantum repeaters based on entanglement purification. *Phys. Rev. A* **59**, 169–181 (1999).
41. Hartmann, L., Kraus, B., Briegel, H.-J. & Dür, W. Role of memory errors in quantum repeaters. *Phys. Rev. A* **75**, 032310 (2007).
42. Kozlowski, W., Dahlberg, A. & Wehner, S. Designing a quantum network protocol. In *Proc. 16th International Conference on emerging Networking EXperiments and Technologies* 1–16. <https://dl.acm.org/doi/10.1145/3386367.3431293> (ACM, 2020).
43. Pourousef, S., Shapourian, H. & Towsley, D. Analysis of asynchronous protocols for entanglement distribution in quantum networks. Preprint at <http://arxiv.org/abs/2405.02406> (2024).
44. Xiao, Z. et al. A connectionless entanglement distribution protocol design in quantum networks. *IEEE Netw.* **38**, 131–139 (2024).
45. Bacciottini, L., Lenzini, L., Mingozzi, E. & Anastasi, G. REDiP: ranked entanglement distribution protocol for the quantum internet. *IEEE Open J. Commun. Soc.* **5**, 397–411 (2024).
46. Briegel, H.-J., Dür, W., Cirac, J. I. & Zoller, P. Quantum repeaters: the role of imperfect local operations in quantum communication. *Phys. Rev. Lett.* **81**, 5932–5935 (1998). Publisher: American Physical Society.
47. Rozpędek, F. et al. Optimizing practical entanglement distillation. *Phys. Rev. A* **97**, 062333 (2018).
48. Krastanov, S., Albert, V. V. & Jiang, L. Optimized entanglement purification. *Quantum* **3**, 123 (2019).
49. Jansen, S., Goodenough, K., de Bone, S., Gijswijt, D. & Elkouss, D. Enumerating all bilocal Clifford distillation protocols through symmetry reduction. *Quantum* **6**, 715 (2022).
50. Goodenough, K. et al. Near-term n to k distillation protocols using graph codes. *IEEE J. Sel. Areas Commun.* 1–1. <https://ieeexplore.ieee.org/document/10479665/> (2024).
51. Deutsch, D. et al. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Phys. Rev. Lett.* **77**, 2818–2821 (1996).
52. Munro, W. J., Stephens, A. M., Devitt, S. J., Harrison, K. A. & Nemoto, K. Quantum communication without the necessity of quantum memories. *Nat. Photonics* **6**, 777–781 (2012).
53. Jiang, L. et al. Quantum repeater with encoding. *Phys. Rev. A* **79**, 032325 (2009). Publisher: American Physical Society.
54. Ralph, T. C., Hayes, A. J. F. & Gilchrist, A. Loss tolerant optical qubits. *Phys. Rev. Lett.* **95**, 100501 (2005).
55. Namiki, R., Jiang, L., Kim, J. & Lütkenhaus, N. Role of syndrome information on a one-way quantum repeater using teleportation-based error correction. *Phys. Rev. A* **94**, 052304 (2016).
56. Pino, J. M. et al. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature* **592**, 209–213 (2021).
57. Bluvstein, D. et al. Logical quantum processor based on reconfigurable atom arrays. *Nature* **626**, 58–65 (2024).
58. Reiserer, A. & Rempe, G. Cavity-based quantum networks with single atoms and optical photons. *Rev. Mod. Phys.* **87**, 1379–1418 (2015).
59. Brekenfeld, M., Niemietz, D., Christesen, J. D. & Rempe, G. A quantum network node with crossed optical fibre cavities. *Nat. Phys.* **16**, 647–651 (2020).
60. Borne, A., Northup, T. E., Blatt, R. & Dayan, B. Efficient ion-photon qubit SWAP gate in realistic ion cavity-QED systems without strong coupling. *Opt. Express* **28**, 11822–11839 (2020).
61. Asaoka, R., Utsugi, T., Tokunaga, Y., Kanamoto, R. & Aoki, T. Optimization of a cavity-QED system for fast two-qubit gates. In *Proc. 2021 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC) 1*. <https://ieeexplore.ieee.org/document/9542055/> (IEEE, 2021).
62. Solak, L. O. R., Rossatto, D. Z. & Villas-Boas, C. J. Universal quantum computation using atoms in cross-cavity systems. *Phys. Rev. A* **109**, 062620 (2024).
63. Kollath-Bönig, J. S. et al. Fast storage of photons in cavity-assisted quantum memories. *Phys. Rev. Appl.* **22**, 044038 (2024).
64. Borregaard, J., Kómár, P., Kessler, E. M., Lukin, M. D. & Sørensen, A. S. Long-distance entanglement distribution using individual atoms in optical cavities. *Phys. Rev. A* **92**, 012307 (2015).

65. Van Meter, R., Satoh, T., Ladd, T. D., Munro, W. J. & Nemoto, K. Path selection for quantum repeater networks. *Netw. Sci.* **3**, 82–95 (2013).
66. Gu, A., Leone, L., Goodenough, K. & Khatri, S. Constant overhead entanglement distillation via scrambling. Preprint at arXiv <https://doi.org/10.48550/arXiv.2502.09483> (2025).
67. La Corte, L., Goodenough, K., Maity, A. G., Santra, S. & Elkouss, D. Bayesian optimization for repeater protocols. In *2025 International Conference on Quantum Communications, Networking, and Computing (QCNC)* 135–142 (2025).
68. Bradley, C. et al. Robust quantum-network memory based on spin qubits in isotopically engineered diamond. *npj Quantum Inf.* **8**, 122 (2022).
69. Bennett, C. H. & Brassard, G. Quantum cryptography: public key distribution and coin tossing. *Theor. Comput. Sci.* **560**, 7–11 (2014).

Acknowledgements

We thank Saikat Guha and Alireza Shabani for their valuable comments on the project idea and intermediate results. We are especially grateful to Liang Jiang and Filip Rozpędek for their insightful feedback and suggestions on the manuscript. Additionally, we thank the Manning College of Information and Computer Sciences at the University of Massachusetts Amherst for providing access to their High Performance Computation Cluster. The authors acknowledge funding support from the NSF-ERC Center for Quantum Networks grant EEC-1941583 and DOE Grant AK000000018297.

Author contributions

D.T. formulated the general problem. P.M. designed the protocols and conducted evaluations. P.M. and D.T. contributed to the analysis. K.G. contributed to discussions. P.M. wrote the manuscript with inputs from all co-authors.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02222-x>.

Correspondence and requests for materials should be addressed to Prateek Mantri.

Peer review information *Communications Physics* thanks Alessandro Laneve, Chen-Long Li, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025