

Routing entanglement in the quantum internet

Mihir Pant,^{1,2,*} Hari Krovi,² Don Towsley,³ Leandros Tassiulas,⁴
Liang Jiang,^{5,6} Prithwish Basu,⁷ Dirk Englund,¹ and Saikat Guha^{2,8}

¹*Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA*

²*Quantum Information Processing group, Raytheon BBN Technologies, 10 Moulton Street, Cambridge, MA*

³*College of Information and Computer Sciences, University of Massachusetts, Amherst, MA*

⁴*School of Engineering and Applied Science, Yale University, 17 Hill House Avenue, New Haven, CT*

⁵*Departments of Applied Physics and Physics, Yale University, New Haven, CT*

⁶*Yale Quantum Institute, Yale University, New Haven, CT*

⁷*Advanced Networking Systems, Raytheon BBN Technologies, 10 Moulton Street, Cambridge, MA*

⁸*College of Optical Sciences, University of Arizona, 1630 East University Boulevard, Tucson, AZ*

Remote quantum entanglement can enable numerous applications including distributed quantum computation, secure communication, and precision sensing. In this paper, we consider how a quantum network—nodes equipped with limited quantum processing capabilities connected via lossy optical links—can distribute high-rate entanglement simultaneously between multiple pairs of users (multiple flows). We develop protocols for such quantum “repeater” nodes, which enable a pair of users to achieve large gains in entanglement rates over using a linear chain of quantum repeaters, by exploiting the diversity of multiple paths in the network. Additionally, we develop repeater protocols that enable multiple user pairs to generate entanglement simultaneously at rates that can far exceed what is possible with repeaters time sharing among assisting individual entanglement flows. Our results suggest that the early-stage development of quantum memories with short coherence times and implementations of probabilistic Bell-state measurements can have a much more profound impact on quantum networks than may be apparent from analyzing linear repeater chains. This framework should spur the development of a general quantum network theory, bringing together quantum memory physics, quantum information theory, and computer network theory.

A *quantum network* can generate, distribute and process quantum information in addition to classical data [1]. The most important function of a quantum network is to generate long distance quantum entanglement, which serves a number of tasks including the generation of multiparty shared secrets whose security relies only on the laws of physics [2, 3], distributed quantum computing [4], improved sensing [5, 6], blind quantum computing (quantum computing on encrypted data) [7], and secure private-bid auctions [8].

Recent experiments have demonstrated *entanglement links*, viz., long-range entanglement established between quantum memories separated by a few kilometers using a point-to-point optical link [9]. As illustrated in Fig. 1, measurements performed at nodes in a quantum network can be used to glue together small entanglement links into longer-distance clusters. The nodes contain quantum memories that store qubits up to their coherence time, sources that generate photons entangled with the quantum memory to be sent to neighboring nodes, and local quantum processors that can perform multi-qubit joint measurements. Entanglement attempts between neighboring nodes are synchronized on a global clock. The *quantum routing protocol* dictates the measurements to be performed locally at each node in order to obtain the desired entanglement topology. Possible goals of a routing protocol could be to enable high rate entanglement among multiple user-pairs simultaneously,

or to generate multi-partite entanglement (entanglement between three or more parties).

The development of network algorithms and protocols for routing and scheduling information flows was critical for the creation of today’s Internet. We expect a similar development in algorithm/protocol design to be critical to design a versatile and high performance quantum network. Some results and intuitions from the theory of classical networks carry over into quantum networking. However, many new challenges arise due to the idiosyncrasies of quantum information. Unlike classical communications, where the rate can be increased by increasing transmit power, photon loss fundamentally limits the entanglement rate over any single link, which must decay exponentially with the length of optical fiber, regardless of the choice of quantum source, the transmit power or the detection strategy [10, 11]. Whereas copying of bits at a network node is common in multipath routing in classical networks [12, 13], copying a qubit is impossible because of the quantum no-cloning theorem [14, 15]. Unlike classical information flow, an entanglement flow does not have directionality. Rather, entanglement is generated across links all over the network and pieced together to form long-range entanglement. Quantum memories are much shorter lived and expensive compared to their classical counterparts making classical routing strategies such as disruption tolerant routing [16, 17]—where a packet is held by a node for until the desired next-hop link is up—much harder to mimic. Finally, distilling and shaping entanglement among a desired set of nodes from many copies of large

* mpant@mit.edu




FIG. 1. Examples of fusing small entangled clusters into larger ones using projective quantum measurements (green ovals) at nodes of a quantum network. Red circles represent qubits and black lines represent entanglement. (a) Two-qubit measurement (Bell state projection) used to connect two entangled links into a longer entangled link; (b) a three-qubit measurement (a GHZ state projection) fuses three clusters (two 2-qubit entangled links and one 3-qubit linear cluster) into one 4-qubit entangled cluster; (c) two adjacent nodes in a network performing a three-qubit GHZ measurement and a two-qubit Bell state measurement simultaneously. The measured qubits are lost, whereas the final entangled state of the unmeasured qubits upon successful completion of both measurements is the same regardless of the order or the simultaneity of the measurements. A quantum measurement at a node may succeed only with a probability, which is a function of the class of optical devices employed to realize the measurement (e.g., linear optics, single photon sources, and single photon detectors) and losses in devices. This figure does not show “failure outcomes”, i.e., the resulting entangled state if one or both measurements fail.

(potentially random) entangled clusters is a purely quantum problem that has no classical analogue.

In this paper, we present protocols for repeater nodes to support multiple simultaneous entanglement flows when every node is limited to the same quantum processing used in repeater chains: (probabilistic) two-qubit Bell state measurement (BSM), also called entanglement swapping. BSMs have been experimentally demonstrated in many physical systems [18–23]. Entanglement attempts between repeater nodes are probabilistic because they are connected via lossy optical links. In every clock cycle, pairs of neighboring repeater nodes attempt entanglement generation. The result of whether an entanglement generation was successful is transmitted back to the corresponding pair of nodes. The repeater nodes then make their local BSM decisions based on this ‘local’ link-state knowledge, i.e., the successes and failures of entanglements established across their nearest neighbor links. Measurements at different nodes can all be done in parallel because BSMs commute with one another.

One of our interesting findings is that *multi-path routing*, i.e. using multiple paths for routing entanglement in a quantum network, can enable long distance entanglement generation with a superior rate-vs.-distance *scaling* than a single linear repeater chain along the shortest path connecting Alice and Bob [24]. The rate along a single repeater chain falls off exponentially with distance [25, 26]. Multi-path routing reduces the exponent, resulting in an exponential increase in rate with increas-

ing distance. The quantum network we propose uses the same basic elements and operations (probabilistic BSMs) as a linear repeater chain but uses more repeaters. Note that increasing the number of repeaters in a repeater chain would not always increase the rate: given the total end-to-end distance, and given the losses at each node, there is an optimal number of repeaters between the end points of a flow that maximizes the rate, i.e., inserting more nodes along that linear path can actually diminish the rate [25, 26].

Another interesting result is that if the repeater nodes have ‘global’ link-state knowledge (knowledge of the state of all links in the network) and the entanglement generation probability is above a (percolation) threshold, multi-path routing enables long distance entanglement-generation at a rate that depends only linearly on the transmissivity η of a single link in the network, whereas the rate achieved by a linear repeater chain connecting Alice and Bob along the shortest path would decay as $\eta^{n_{\text{SP}}}$ where n_{SP} is the length of the shortest path. Pirandola recently showed that even a linear repeater chain can attain a rate that is proportional to η , but that requires repeater nodes equipped with error-corrected quantum processors [27]. We achieve the same feature (rate proportionality to η) by multipath routing with percolation, with a much simpler repeater. We also analyze repeater protocols to support multiple entanglement-generation flows. This analysis reveals that simple space-time-division multiplexing strategies that use local link-state knowledge at nodes can outperform the best rate-region (the set of rates simultaneously achievable by different flows) attainable by repeater nodes that simply time share among assisting individual flows.

Our work also opens up a number of new questions that remain unanswered. We abstract off the entanglement routing problem to the following parameters: G (network topology), p (probability of successful creation of an entangled pair across one link in a given time step), q (probability of a successful Bell measurement when attempted), S (number of parallel links across a network edge) and T (the number of time slots that a memory can hold a qubit before it decoheres). Even in this simplified model, finding the rate-region optimizing routing protocol remains open. The aforesaid abstraction applies when the only source of imperfection at each component (including the quantum memories) is pure loss. Since our protocol only requires a quantum memory to hold a qubit for one entanglement attempt between neighboring stations ($T = 1$), photon loss would indeed be the major source of imperfection in many implementations of the protocol. Accounting for more general errors would require purification of entanglement [28–30], which will require us to introduce the Fidelity of entanglement during intermediate steps of the routing protocol as an additional parameter, as was done by Jiang *et al.* [31]. Furthermore, we restricted our analysis only to 2-qubit measurements at repeater nodes. Multi-qubit unitary operations and multi-qubit measurements at repeater nodes




FIG. 2. Schematic of a general quantum repeater network. The large (green) circles represent ‘trusted’ nodes, which are connected via a classical network. The blue circles denote repeater stations, and the red circles inside them represent quantum memories. The dashed lines connecting the red circles are independent lossy optical (fiber) channels. In principle, all nodes in the network could be equipped with quantum repeaters (i.e., no trusted nodes), in which case depending upon the need, a node can be a consumer of shared entanglement, or act a router to conduit entanglement flows between other nodes.

(e.g., a 3-qubit GHZ projection across three locally held qubits) would require more complex repeaters than those in repeater chains, but may improve the achievable rates. Finally, it will be interesting to consider repeater protocols for distillation of multi-partite entanglement shared between more than two parties, and a repeater network that supports multiple simultaneous flows of multi-partite entanglement generation.

I. BACKGROUND

Let us consider a quantum network with topology described by a graph $G(V, E)$. Each of the $N = |V|$ nodes is equipped with a quantum repeater, and each of the $M = |E|$ edges is a lossy optical channel of range L_i (km) and power transmissivity $\eta_i \propto e^{-\alpha L_i}$, $i \in E$. Consider K source-destination (Alice-Bob) pairs (A_j, B_j) , $1 \leq j \leq K$, situated at (not necessarily distinct) nodes in V , each of which would like to generate maximally-entangled qubits (i.e., ebits) between themselves (and thus by definition not entangled with any other party, due to the monogamy property of entanglement), at the maximum rates possible R_j (ebits per channel use). The high-level objective is: *Given a class of quantum and classical operations at each of the repeater nodes of the underlying network, what operations should be performed at the repeater stations to maximize the rate region (R_1, R_2, \dots, R_K) simultaneously achievable by the entanglement flows?* More importantly, one would like to address networking questions such as: (a) what is the maximum rate-region attainable, (b) what is the trade-off between sum throughput and latency of the K entanglement flows, and (c) where should repeater nodes be placed, with constraints on devices (e.g., memories, sources, and detectors), to maximize the attainable rate region; all being subject to various practical considera-

tions. Ultimately one would like to develop explicit and efficient practical quantum routing protocols that employ quantum operations implemented via lossy and noisy devices, while only requiring local link-state knowledge and limited knowledge of the global network topology, analogous to the classical internet.

The entanglement-generation rate across a link of transmissivity η , in the absence of any repeater mediation, is limited to $-\log_2(1-\eta)$ ebits per mode, amounting to $\approx 1.44\eta$ ebits per mode when $\eta \ll 1$ [11][32]. The number of modes per second is a device-dependent constant, upper bounded by the maximum of the optical bandwidth of the source and the electrical bandwidth of the detector. Since $\eta \sim e^{-\alpha L}$ where L is the length of optical fiber, the ebits-per-mode rate also falls off exponentially with range L . Most analyses of repeater networks have been limited to linear chains, with the objective of outperforming the repeater-less bound [25, 26, 30, 31, 33–36]. Pirandola analyzed entanglement-generation capacities of repeater networks assuming ideal repeater nodes, i.e., those equipped with fully-error-corrected quantum processors and argued that for a single flow ($K = 1$), the maximum entanglement-generation rate R_1 reduces to the classical max-flow min-cut problem with edge e being associated with capacity $C(e) = -\log_2(1 - \eta(e))$ ebits per channel use [27], where $\eta(e)$ is the transmissivity of edge e . Pirandola subsequently argued that classical cut-set bounds with the above link capacity give outer bounds to the K -flow capacity region, but again, for ideal repeater nodes. Schoute and co-authors [37] developed routing protocols on specific network topologies and found scaling laws as functions of N , the number of qubits in the memories at nodes, and the time and space consumed by the routing algorithms, under the assumption that each link generates a perfect, lossless EPR pair in every time slot, and that the nodes’ actions are limited to (perfect) Bell-state measurements (BSMs). Acín




FIG. 3. Schematic of a square-grid topology. The blue circles represent repeater stations and the red circles represent quantum memories. Every cycle (time slot) of the protocol consists of two phases. (a) In the first (external) phase, entanglement is attempted between neighboring repeaters along all edges, each of which succeed with probability p (dashed lines). (b) In the second (internal) phase, entanglement swaps are attempted within each repeater node based on the successes and failures of the neighboring links in the first phase—with the objective of creating an unbroken end-to-end connection between Alice and Bob. Each of these internal connections succeed with probability q . Memories can hold qubits for $T \geq 1$ time slots.

and co-authors [38] have considered the problem of entanglement percolation where neighboring nodes share a perfect, lossless pure state. Further, van Meter and co-authors developed explicit networking protocols also restricted to pair-wise EPR pair generation and BSMs, but accounting for imperfect fidelities of the EPR pairs (and thus requiring purification over multiple imperfect pairs), and finite coherence times of the qubit memories [39].

II. ENTANGLEMENT ROUTING PROTOCOLS

A. Problem statement and notation

Consider a graph $G(V, E)$ that denotes the topology of the repeater network. See Fig. 2 for an illustration. Each node $v \in V$ is a repeater (blue circles), and each edge $e \in E$ is a physical link connecting two repeater nodes. $S(e) \in \mathbb{Z}^+$ is an integer edge weight, which corresponds to the number of parallel (single spatial, spectral or polarization mode) channels across the edge e (shown using blue lines). The number of memories at node v is $\sum_{e \in \mathcal{N}(v)} S(e)$ (see Fig. 2), where the sum is over $\mathcal{N}(v)$, the set of nearest neighbor edges of v , with $d(v) = |\mathcal{N}(v)|$ the degree of node v .

Time is slotted. We assume that each memory can hold a qubit perfectly for $T \geq 1$ time slots (T should be much smaller than the memory’s coherence time). Each time slot t , $t = 1, 2, \dots$, is divided into two phases: the “external” phase and the “internal” phase, which occur in that order. During the external phase, each of the $S(e)$

pairs of memories across an edge e attempts to establish a shared entangled (EPR) pair. An entanglement attempt across any one of the $S(e)$ parallel links across edge e succeeds with probability $p_0(e) \sim \eta(e)$ [10, 11], where $\eta(e) \sim e^{-\alpha L(e)}$ is the transmissivity of a lossy optical channel of length $L(e)$. Using two-way classical communication over edge $e(u, v)$, neighboring repeater nodes u, v learn which of the $S(e)$ parallel links (if any) succeeded in the external phase.

Let us assume that neighboring repeaters pick up to one successfully-created ebit (i.e., ignore multiple successes if any) as in Ref. [25, 26], in which case the probability that one ebit is established successfully across the edge e during the external phase is given by: $p(e) = 1 - (1 - p_0)^{S(e)}$. Let us also assume $S(e) = S, \forall e \in E$, which in turn gives us $p(e) = p, \forall e \in E$. While our results in this paper can be adapted to any network topology, we will henceforth use the 2D regular square grid topology (Fig. 3) to illustrate the performance of our routing algorithms.

One instance of the resulting external links created between repeater nodes after the external phase is shown in Fig. 3(b) using solid lines. In the internal phase, entanglement swap (BSM) operations are attempted locally at each repeater node between pairs of qubit memories. We associate these BSM attempts as *internal links*, i.e. links between memories internal to a repeater node, shown using dotted lines inside repeaters in Fig. 3(b). If $T > 1$, a repeater node can attempt a BSM between qubits held in two memories that were entangled with their respective neighboring node’s qubits in two different time slots. For minimizing the demands on memory coherence time [25, 26], we will assume $T = 1$. So, BSMs will always be attempted between two qubits in distinct memories that were entangled with their respective counterparts at their respective neighboring nodes in the same time slot. Each of these internal-link attempts succeed with probability q . Therefore, after the conclusion of one time slot, along a path comprising k edges (and thus $k - 1$ repeater nodes), one ebit is successfully shared between the end points of the path with probability $p^k q^{k-1}$. The maximum number of ebits that can be shared between Alice (say, node a) and Bob (say, node b) after one time slot is $\min\{d(a), d(b)\}$, assuming S is the same over all edges. For the square-grid topology shown, the maximum number of ebits that can be generated between Alice and Bob in each time slot is 4.

The remainder of the paper is dedicated to finding the optimal strategy for each repeater node in order to decide which locally-held qubits to attempt BSM(s) on during the internal phase of a time slot, based ideally only on knowledge of the outcomes (success or failure) of the nearest neighbor links, i.e., local link-state knowledge, during the respective preceding external phases. We will assume that each repeater node is aware of the overall network topology as well as the locations of the K Alice-Bob pairs. The goal of the optimal repeater strategy will be to attain the maximum entanglement-generation rate (if there is a single Alice and Bob, i.e., $K = 1$) or the

maximum rate-region for multiple flows (i.e., $K > 1$).

B. Multipath routing of a single entanglement flow

1. Entanglement routing with global link-state information

We begin with the assumption that global link-state knowledge is available at each repeater node, i.e., the state of every external link in the network after the external phase is known to every repeater in the network and can be used to determine the choice of which internal links to attempt within the nodes. Each memory can only be part of one entanglement swap, i.e., each red node can only be part of one internal edge. Consider the following greedy algorithm to choose the internal links: consider the subgraph induced by the successful external links and the repeater nodes (at the end of the external phase), and find in it the shortest path connecting Alice and Bob. If no connected path between Alice and Bob exists, no shared ebits are generated in that time slot. If a shortest path of length k_1 is found, all internal links along the nodes of that path are attempted, and the probability a shared ebit is generated by this path is the probability that all $k_1 - 1$ internal link attempts were successful, i.e., q^{k_1-1} . We then remove all the (external and internal) links of the above path from the subgraph, and find a shortest path connecting Alice and Bob in the pruned subgraph. Note that instead of removing the links of the first path from the subgraph, we could simply search for a shortest path in the original subgraph but one that is edge disjoint from the previous path. If such a path exists, we again attempt all internal links at the nodes of this path, so the probability the path contributes to the generation of an ebit between Alice and Bob is q^{k_2-1} where k_2 is the length of the second path; and so on.

The entanglement generation rate achieved using this greedy algorithm R_g is the sum of expected rates (in ebits per time slot) from these paths. Given the degree-4 nodes in a square grid topology, there can be a maximum of four edge disjoint paths between Alice and Bob. Fig. 3(b) illustrates our greedy algorithm. Given the set of external links created, the shortest path has length $k_1 = 4$, the next path has length $k_2 = 6$, and no further paths can be found. The two edge-disjoint paths are highlighted in green. Hence, the internal links depicted with the dashed lines in Fig. 3(b) are attempted and the expected number of shared ebits generated in this time cycle is: $q^{k_1-1} + q^{k_2-1}$. The net entanglement generation rate is the expectation of sums like the above (with up to four terms) over many random instantiations of the $(p, 1 - p)$ external-link creations during the external phase of many time slots. Evaluating this expected rate $R_g(p, q)$ achieved by the above routing strategy analytically as a function of the Alice-Bob distance (X_1, X_2) is difficult, even for a square-grid topology.

The intuition behind this simple greedy algorithm is

that the entanglement generation rate along a path of length k decays exponentially as q^{k-1} , suggesting that attempting internal links to facilitate connections along the shortest path first would optimize the expected rate. However, it is possible to draw random instances of successes of external links, where either one of the two possible options—(1) picking the shortest path (which disrupts all other paths) and (2) picking two edge disjoint (but longer) paths—could yield either a larger or a smaller expected rate than the other, depending upon the value of q . If q is larger than a threshold, option (2) would have a larger expected rate and vice versa. Finding the global optimal rule remains an open problem. It is easy however to prove that the greedy algorithm achieves a rate within a factor of four of the optimum algorithm employing global link-state knowledge, $R_{\text{opt}}(p, q)$. Let us denote the length of the shortest path between Alice and Bob with Manhattan distance (X, Y) in the induced subgraph after the external phase, as $n_{\text{SP}}(p)$. This quantity is of interest in percolation theory, and is not completely understood analytically. It undergoes a sharp transition (i.e., starts out large and suddenly jumps to a much smaller value) as p crosses p_c from below to above. Clearly, $R_g(p, q) \geq \mathbb{E}[q^{n_{\text{SP}}(p)-1}]$ since using the shortest path is the first step of the greedy algorithm. Furthermore, since the optimal rule can create entanglement over a maximum of four edge-disjoint paths in each time step, each of which must have a length no less than the length of the shortest path, $R_{\text{opt}}(p, q) \leq \mathbb{E}[4q^{n_{\text{SP}}(p)-1}] \triangleq R_{\text{opt}}^{\text{(UB)}}(p, q)$. Therefore, $R_{\text{opt}} \geq R_g \geq R_{\text{opt}}/4$, i.e., the greedy rule will achieve the same rate-vs.-distance scaling as the optimal algorithm that employs global link-state knowledge, and can be worse only by a small constant factor.

In Fig. 4(a) we plot $R_g(p, q)$ as a function of the Alice-Bob Manhattan distance (X, Y) on the square grid (measured in number of hops) with $q = 1$. When $p > p_c$, the bond percolation threshold of the underlying network ($p_c = 0.5$ for the square lattice), a *giant connected component* is formed by the external links alone at the end of the first (external) phase of a time slot. Recall that the rate along a length k path is $p^k q^{k-1}$, where $p \sim \eta$ is the transmissivity of each link. In the network case, when $p > p_c$ and $q = 1$, the p^k portion of the rate expression becomes immaterial for scaling with Alice-Bob distance, since percolation guarantees a connected path to exist between Alice and Bob along successful external links in each time slot. So, if $q = 1$, $R_g(p, q)$ remains essentially distance invariant. When $p < p_c$, the rate falls off exponentially with distance (even when $q = 1$). It is instructive to note here that the optimal rate (entanglement-generation capacity) achievable on a single length k path does not depend on k , and only on the transmissivity of the lossiest link in the path, i.e., $C \sim \eta$ [27], but achieving this requires infinite-coherence-time quantum memories and ideal quantum operations at nodes. The multi-path gain in the $p > p_c$ regime lets us achieve a distance-independent rate, but with memories whose co-

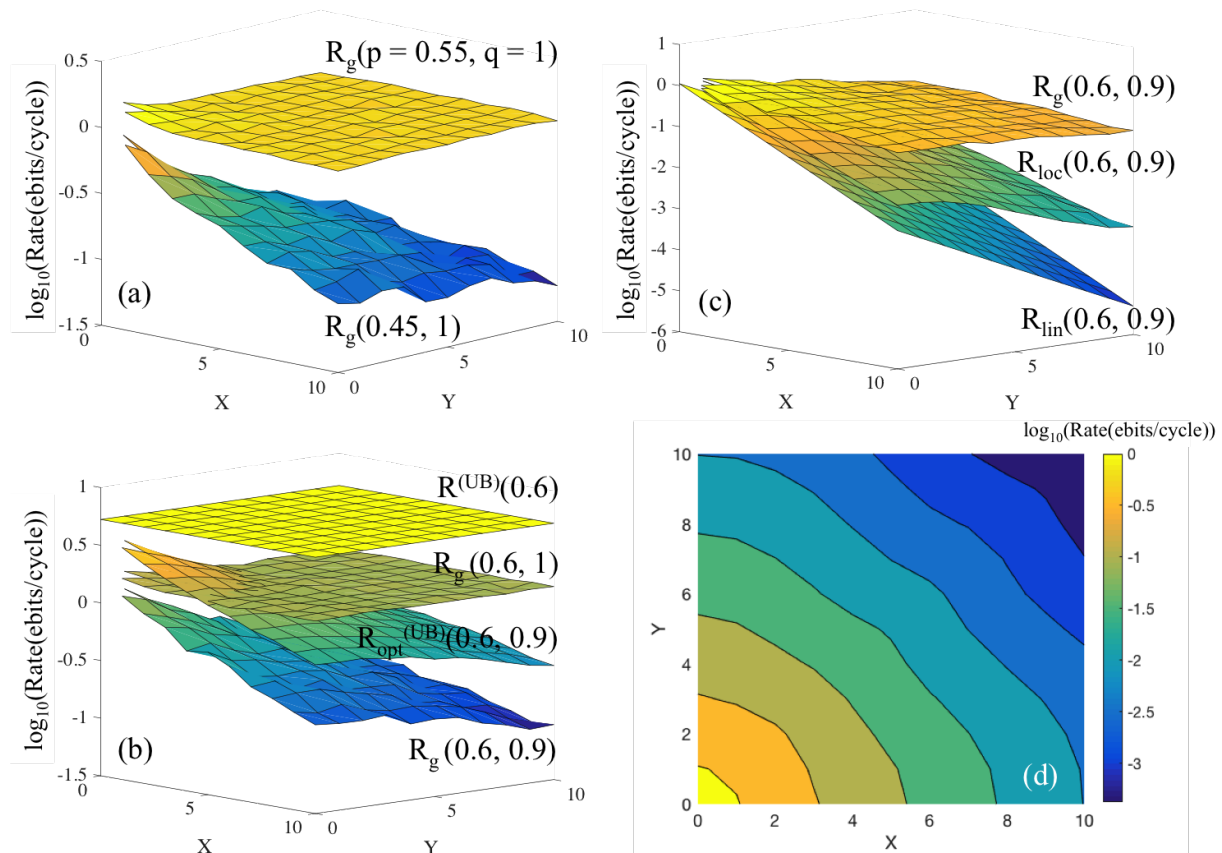


FIG. 4. Entanglement generation rate as a function of the Alice-Bob separation along X and Y (on a square grid) as a function of (p, q) ; (a) $R_g(p, q)$ is the rate attained by a global-knowledge-based protocol we propose where each node, in each time step, knows whether any link in the entire network succeeded or failed to establish entanglement. For the case of $q = 1$, R_g is distance independent when p is greater than the bond percolation threshold (0.5 for the square lattice) and decays exponentially if it is below the threshold. (b) $R^{(UB)}(0.6)$ is the distance-independent Pirandola rate upper bound for $p = 0.6$, achieving which requires perfect quantum processing at repeater nodes. $R_g(0.6, 1)$ is also distance independent, and within a factor 3.6 of $R^{(UB)}(0.6)$. With $q < 1$, e.g., $R_g(0.6, 0.9)$, the rate decays exponentially with distance. $R_{opt}^{(UB)}(0.6, 0.9)$ is an upper bound on the rate attainable with global-knowledge by any protocol. (c) R_{loc} is attained by a protocol we propose where each node, in each time step, only needs to know the link state of neighboring edges. The rate-distance scaling exponent of R_{loc} is clearly worse than R_g , but is significantly superior to that of a linear repeater chain along the shortest path, R_{lin} , demonstrating multi-path routing advantage even with local link-state knowledge. (d) Contour plot of the entanglement generation rate with the local rule when $p = 0.6$ and $q = 0.9$. Although the Alice to Bob distance along the network links is $X + Y$, there is a noticeable enhancement in the rate along the $X = Y$ direction because of more Alice-Bob paths of similar length.

herence times are no more than one time slot. The rates are calculated using monte-carlo simulations which results in some numerical noise that is insignificant compared to the difference between the plots, but is visible in $R_g(0.45, 1)$.

A general upper bound on the entanglement generation rate is given by the min-cut of the graph [27], and for a square lattice, is given by $R^{(UB)}(p) = -\log_2[(1-p)^4]$. $R^{(UB)}(0.6)$ is plotted in Fig. 4(b). The known methods for achieving $R^{(UB)}$ require infinite coherence time memories and error-corrected quantum processors at each node. For our implementation (assuming global link state knowledge), $R_g(0.6, 1)$ is also plotted in Fig. 4(b). Although our protocol only requires memories to hold

entanglement for one time step, the multi path advantage gives us the same constant rate-distance scaling and within a factor of ~ 3.6 of $R^{(UB)}(0.6)$. The assumption of perfect BSs is unrealistic and $q < 1$, in which case $R_g(p, q)$ falls off exponentially with distance; even when $p > p_c$, as seen in the plot for $R_g(0.6, 0.9)$. Finally, we plot the above discussed upper bound on R_g , $R_{opt}^{(UB)}(0.6, 0.9)$, which as expected has the rate-distance scaling of R_g , but larger by a factor less than 4.

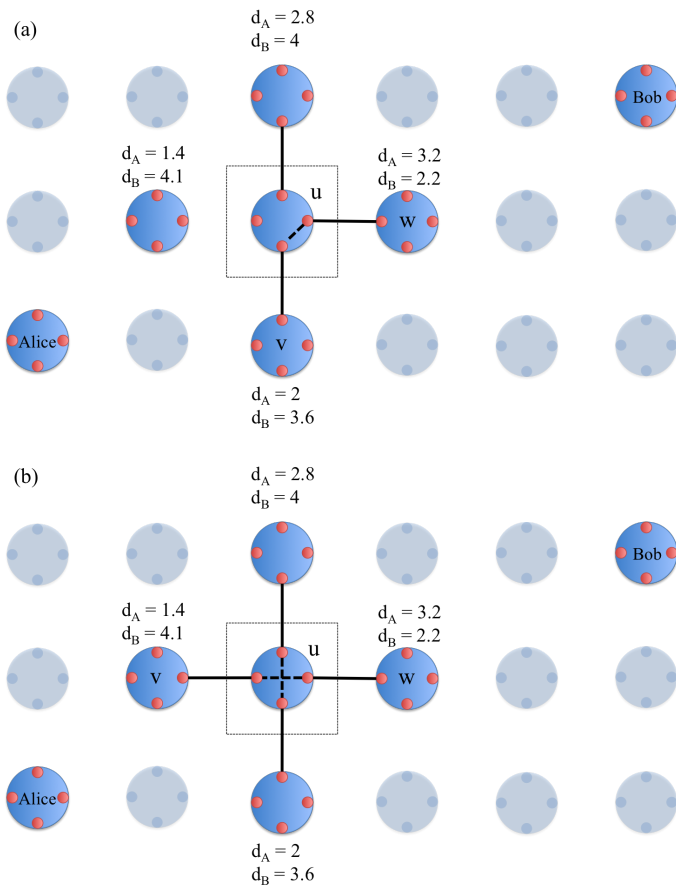


FIG. 5. The entanglement swap rule used at the repeater C in the dotted box in the case of local link-state knowledge. A and B are the repeaters closest to Alice and Bob, respectively, with a direct edge to C . (a) If two or three links are up, the memories linked to A and B undergo an entanglement swap. (b) If four links are up, the remaining two memories also undergo an entanglement swap.

2. Entanglement routing with local link-state information

$R_g(p, q)$, the rate attained by the protocol described in the previous subsection that employs global link-state knowledge, is re-plotted in Fig. 4(c). We also plot $R_{\text{lin}}(p, q) = p^{n_{\text{SP}}(1)} q^{n_{\text{SP}}(1)-1}$, the rate attained by a single linear repeater chain, where $n_{\text{SP}}(1)$ is the shortest-path length between Alice and Bob along the edges of the underlying square grid. The assumption of global link-state knowledge in large networks is unrealistic, as it requires memories whose coherence time increases with the network size due to the time required for the traversal of link-state information across the entire network. In this section, we describe a more realistic protocol in which knowledge of success and failure of an external link at each time slot is communicated only to the two repeater nodes connected by the link, as is the case in the analysis of many ‘second-generation’ linear repeater chains [25, 31, 33]. Repeater nodes need to decide on

which pair(s) of memories BSMS should be attempted (i.e., which internal links to attempt), based only on information about the states of external links adjacent to them. We assume that network topology and positions of Alice and Bob are known to each repeater station, and communicated classically beforehand.

Let us consider a local repeater rule illustrated in Fig. 5. The repeater u inside the dotted square box has to make a decision regarding which internal edges to attempt based on the information of which of the four neighboring external edges have been successfully created in the external phase. We associate d_A and d_B as the distance to Alice and Bob, respectively, at every repeater node. We use the \mathbb{L}^2 norm distance to Alice (resp. Bob) as d_A (resp. d_B). Of all the nearest neighbor nodes of u whose links to u were successful in that time slot, we label the one that has the minimum d_A as v . Similarly, the neighbor with a successful external link with u and the minimum d_B is labelled w . An internal link is attempted between the memories connected to v and w respectively, as shown in Fig. 5(a). If v and w are the same node, v (or w) is replaced by node u ’s nearest-neighbor node with the next smallest value of d_A (or d_B). The choice of whether to replace v or w is made in a manner that minimizes the sum of d_A and d_B from the eventually chosen two neighbors to connect. If all four external links are successful, an additional internal link is attempted between the remaining two memories as shown in Fig. 5(b). If only one of the neighboring external links is successful, no internal links are attempted, since this repeater node cannot be part of a path from Alice to Bob in that time slot. If two neighbors have the same values of d_A and d_B , an unbiased coin is tossed to determine the choice of v and w , to preserve symmetry in the protocol.

The entanglement generation rate $R_{\text{loc}}(p, q)$ achieved by the above described local rule is plotted in Fig. 4(c) and compared to $R_g(p, q)$ and $R_{\text{lin}}(p, q)$. We use $p = 0.6$ and $q = 0.9$, the same values used for the global-information rate plots in Fig. 4(b). As one expects, the rate-distance scaling of R_{loc} is worse than that of R_g . However, the rate-distance scaling exponent achieved by the local rule is superior to that of a linear chain, even though the physical elements employed to build the repeaters are identical. This is proven analytically in appendix B 2. Note that each of the three rates $R_g, R_{\text{loc}}, R_{\text{lin}}$ fall exponentially with distance, but the exponents are different. The scaling advantage of R_{loc} over R_{lin} arises because the local rule allows the entanglement-generation flow between Alice and Bob to find different (and potentially simultaneously multiple) paths in different time slots, and does not have to rely on all links along a linear chain to be successful. This is analogous to multi-path routing in a classical computer network. The contour plot in Fig. 4(d) further illustrates this point: there is a noticeable enhancement of R_{loc} along the $X = Y$ line because the diagonal direction contains the largest spatial density of possible paths between Alice and Bob. The scaling advantage over R_{lin} persists in any direction,




FIG. 6. (a) Multi-flow routing for two Alice-Bob pairs that lie along the sides of a 6×6 square, embedded in a 100×100 grid; (b) rate region (R_1, R_2) with different rules at repeater nodes, each employing local link-state knowledge, for $p = q = 0.9$. (c) Multi flow routing when the Alice-Bob paths cross (d) multiflow rate region for two local-knowledge rules.

i.e., along $Y = 0$ as well.

Sweeping over different values of p and q , we find that the multi-path advantage relative to a linear repeater chain increases as p decreases from unity, but there is little relative improvement as q is varied (see Appendix B 1).

Clearly, other distance metrics (e.g., \mathbb{L}^p norm for $p \geq 1$) can be used in lieu of the \mathbb{L}^2 norm in the algorithm described above. In Appendix A, we present a recursive numerical evaluation technique to find the rate-optimal distance metric, which can be applied to any network topology. For planar network topologies, the \mathbb{L}^2 norm appears near-optimal for our local routing algorithm.

An analytical enumeration of the expected number of edge-disjoint paths as a function of p between Alice and Bob separated by a given distance (X, Y) in a bond-percolation instance (i.e., with $p > p_c$) of a network is an open question, the solution of which will enable a firmer quantitative understanding of the multi-path advantage in entanglement generation in a repeater network.

C. Simultaneous entanglement flows

In this section, we consider simultaneous entanglement-generation flows between two Alice-Bob pairs, using local link state knowledge at all repeater nodes. Consider two pairs Alice 1 - Bob 1 (red nodes) and Alice 2 - Bob 2 (green nodes) as shown in the two scenarios in Fig. 6. In Fig. 6(a), the shortest paths




FIG. 7. A heat map plotting p_{usage} , the probability that a given repeater node is involved in a successful creation of a shared ebit generated between Alice and Bob, separated by 6 hops in an underlying square grid topology, when our local rule is employed. We assume $p = 0.9$ and $q = 0.9$.

connecting the two Alice-Bob pairs do not cross, but they do in Fig. 6(b). In both cases, they are placed at the four corners of a 6×6 square grid, embedded within a large square grid network. Denote by R_1 and R_2 the entanglement generation rates achieved by the two Alice-Bob pairs respectively. We first consider the case of non-intersecting flows shown in Fig. 6(a). A simple strategy is for every single repeater node (including the nodes labeled as the two Alices and Bobs) to use the local rule described in the previous section tailored to support the Alice 1-Bob 1 flow for a fraction, λ , of the time slots and to support the Alice 2-Bob 2 flow for the remaining $1 - \lambda$ fraction. For $p = q = 0.9$, the rate region attained by varying $\lambda \in [0, 1]$ is depicted with the blue line in Fig. 6(b), which we refer to as single-flow time-share. However, if every repeater with the exception of the Alices and Bobs carry out the above time-sharing strategy, even when all repeater nodes support flow 1, there is still some ‘left-over’ non-zero R_2 that is attained. This multi-flow time-share rate region is shown using the red line in Fig. 6(a).

In Fig. 7, for the case that Alice and Bob are separated by 6 hops on the square grid, we plot a color map of p_{usage} , the probability a given repeater node is involved in a successful creation of a shared ebit generated between Alice and Bob when our local rule is employed. We observe that only the repeaters lying in a small spatial region surrounding the straight line joining Alice and Bob are used significantly. This observation motivates a multi-flow spatial-division rule, in which we divide the network between the two flows, as shown in Fig. 6(a). Any repeater in the red shaded region follows the local rule tied to the Alice 1 - Bob 1 flow while repeaters in the green region operate with the local rule tied to the Alice 2

- Bob 2 flow. The placement of the boundary determines the rates R_1 and R_2 . The rate region attained is plotted with the yellow line in Fig. 6(b). This significantly outperforms time sharing. The two flows can co-exist and operate with a very small reduction from their individual best rates, because the repeaters they respectively benefit from the most form almost disjoint sets.

In the other extreme, we consider two Alice-Bob pairs, still separated by six hops, but with their shortest paths crossing as shown in Fig. 6(c). The rate region attained by multi-flow time sharing, shown by the line segment BC, still provides an improvement over single-flow time-sharing, shown by the line segment AD, as shown in Fig. 6(d). It is interesting to note that the maximum R_1 under multi-flow time sharing (point B) is slightly lower than maximum R_1 with the single-flow time-share rule (point A). This happens because unlike in single flow time-share, the nodes at Alice 2 and Bob 2 do not contribute to R_1 under multi-flow time-share. A point along AB represents time sharing between the strategies at points A and B. To further increase the rate, we adopt a multi-flow spatial division strategy in which nodes in the red region are configured to assist flow 1 and nodes in the green region are configured to assist flow 2. Varying the angle θ demarcating those regions results in the rate region shown by the yellow line in Fig. 6(d). This time, the improvement due to the spatial-division rule is not as pronounced, since the spatial regions corresponding to ‘useful’ repeater nodes for the two flows are not disjoint.

III. CONCLUSIONS

We proposed and analyzed quantum repeater protocols for entanglement generation in a quantum network in an architecture that uses the same elements as in linear repeater chains. We accounted for channel losses between repeater nodes and the probabilistic nature of entanglement swaps at each repeater stemming from device inefficiencies as well as the probabilistic nature of Bell-state measurements (e.g., due to inherent limitations of using linear optics and lossy detectors). The rate attained for a single entanglement-generation flow can far outperform that attainable over a linear repeater chain, even when the nodes only have local link-state knowledge, due to the multi-path routing advantage. We also proposed a modified version of our routing protocol for supporting simultaneous entanglement generation flows between multiple Alice-Bob pairs. We found multi-flow entanglement routing strategies that outperform the rate region attained when each repeater simply time shares among each flow. Our results suggest that building and connecting quantum repeaters in non-trivial network topologies could provide a substantial benefit over linear repeater

chains alone. Seen another way, given constraints on the number and quality of quantum memories, link losses between nodes, and limited and imperfect processing capabilities at repeater nodes, a 2D network topology can outperform the repeater-less rate-vs.-distance upper limits [10, 11] more easily than a linear repeater chain connecting the communicating parties.

Our work has also opened a number of new questions. Even in our simplified model—an abstraction that applies when the only source of imperfection at each component (including the quantum memories) is pure loss—the rate-optimal protocol remains open. Since our protocol only requires a quantum memory to hold a qubit for one entanglement attempt between neighboring stations, photon loss would indeed be the major source of imperfection in many implementations of the protocol. Accounting for more general errors would require purification of entanglement [28–30], i.e., converting several poorer-quality EPR pairs into a few good ones using local quantum operations and classical communication, accounting which will require us to introduce the Fidelity of shared entanglement at intermediate steps of the protocol. Furthermore, we restricted our analysis to the same operation used in the nodes of a linear repeater chain: 2-qubit measurements. Being able to perform multi-qubit unitary operations and multi-qubit measurements at repeater nodes (e.g., a 3-qubit GHZ projection across three locally held qubits) may improve the achievable rate regions. The idea of using a distance metric to choose the measurements at the repeater station could be used in protocols that use measurements of more than two qubits as well. Finally, it will be interesting to consider repeater protocols for the distillation of multi-partite entanglement shared between more than two parties, and a repeater network that can support multiple simultaneous flows of generation of multi-partite entanglement.

ACKNOWLEDGMENTS

SG would like to thank Stefano Pirandola, Zachary Dutton, and Dongning Guo for valuable discussions. MP, DE and LJ acknowledge support from the Air Force Office of Scientific Research MURI (FA9550-14-1-0052) and the Army Research Laboratory (ARL) Center for Distributed Quantum Information (CDQI). SG, HK, PB, MP and DE would like to acknowledge the Office of Naval Research program *Communications and Networking with Quantum Operationally-Secure Technology for Maritime Deployment* (CONQUEST), awarded under prime contract number N00014-16-C-2069. SG, DT, PB and LT acknowledge the ARL DAIS-ITA program; this work would not have occurred without the collaborations and ideas seeded by this program.

[1] H. J. Kimble, *Nature* **453**, 1023 (2008).

[2] A. K. Ekert, *Physical Review Letters* **67**, 661 (1991).

- [3] C. H. Bennett and G. Brassard, *Theoretical Computer Science* **560**, 7 (2014).
- [4] J. I. Cirac, A. K. Ekert, S. F. Huelga, and C. Macchiavello, *Physical Review A* **59**, 4249 (1999).
- [5] D. Gottesman, T. Jennewein, and S. Croke, *Physical Review Letters* **109**, 070503 (2012).
- [6] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, *Nature Physics* **10**, 582 (2014), arXiv:arXiv:1310.6045v1.
- [7] A. Broadbent, J. Fitzsimons, and E. Kashefi, in *2009 50th Annual IEEE Symposium on Foundations of Computer Science (IEEE, 2009)* pp. 517–526.
- [8] S. Guha, T. Hogg, D. Fattal, T. Spiller, and R. G. Beausoleil, *International Journal of Quantum Information* **06**, 815 (2008), arXiv:arXiv:0707.2051v1.
- [9] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Nature* **526**, 682 (2015), arXiv:1508.05949.
- [10] M. Takeoka, S. Guha, and M. M. Wilde, *Nature communications* **5**, 5235 (2014).
- [11] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, *Nature Communications* **8**, 15043 (2017), arXiv:1510.08863.
- [12] A. Vahdat and D. Becker, Technical report number CS-200006, Duke University, 1 (2000), arXiv:1001.3405.
- [13] S. Biswas and R. Morris, *Sigcomm* (2005), 10.1145/1090191.1080108.
- [14] W. K. Wootters and W. H. Zurek, *Nature* **299**, 802 (1982).
- [15] D. Dieks, *Physics Letters A* **92**, 271 (1982).
- [16] S. Jain, K. Fall, and R. Patra, in *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications - SIGCOMM '04*, Vol. 34 (ACM Press, New York, New York, USA, 2004) p. 145.
- [17] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, *Proceedings - IEEE INFOCOM* **00** (2006), 10.1109/INFOCOM.2006.228.
- [18] Z.-S. Yuan, Y.-A. Chen, B. Zhao, S. Chen, J. Schmiedmayer, and J.-W. Pan, *Nature* **454**, 1098 (2008).
- [19] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Nature* **497**, 86 (2013), arXiv:1212.6136.
- [20] S. Olmschenk, D. N. Matsukevich, P. Maunz, D. Hayes, L.-M. Duan, and C. Monroe, *Science* **323**, 486 (2009), arXiv:0907.5240.
- [21] J.-W. Pan, D. Bouwmeester, H. Weinfurter, and A. Zeilinger, *Physical Review Letters* **80**, 3891 (1998).
- [22] C. W. Chou, H. de Riedmatten, D. Felinto, S. V. Polyakov, S. J. van Enk, and H. J. Kimble, *Nature* **438**, 828 (2005).
- [23] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, *Nature* **449**, 68 (2007).
- [24] Pirandola recently showed [27], for an information-theoretic description of repeaters that are ideal fully-error-corrected universal quantum processors, that the optimal rate attainable for multi-path entanglement routing using such ideal repeaters is superior to the rate of a linear chain of ideal repeaters.
- [25] S. Guha, H. Krovi, C. A. Fuchs, Z. Dutton, J. A. Slater, C. Simon, and W. Tittel, *Physical Review A* **92**, 022357 (2015).
- [26] M. Pant, H. Krovi, D. Englund, and S. Guha, *Physical Review A* **95**, 012304 (2017), arXiv:1603.01353.
- [27] S. Pirandola, arXiv preprint (2016), arXiv:1601.00966.
- [28] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, *Physical Review Letters* **77**, 2818 (1996).
- [29] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, *Physical Review Letters* **76**, 722 (1996).
- [30] H.-J. Briegel, W. Dür, J. Cirac, and P. Zoller, *Physical Review Letters* **81**, 5932 (1998).
- [31] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, *Physical Review A* **79**, 032325 (2009).
- [32] The achievability of $-\log_2(1 - \eta)$ ebits per mode of secret communication rate over the lossy channel (with two way authenticated public classical communication) was first proven in 2009 by Pirandola *et al.* [40]. In 2014, Takeoka *et al.* proved an upper bound to the secret-key agreement capacity, $\log_2[(1 + \eta)/(1 - \eta)]$ ebits per mode [10], which equals $\approx 2.88\eta$ ebits per mode when $\eta \ll 1$, thereby establishing that the rate attained by *any* protocol must decay linearly with the channel’s transmissivity and hence exponentially with distance L in optical fiber (since $\eta \sim e^{-\alpha L}$). In 2015, Pirandola *et al.* proved an improved (weak converse) upper bound of $-\log_2(1 - \eta)$ ebits per mode, which established that as the secret key agreement capacity of the pure loss bosonic channel [11]. In 2017, Wilde *et al.* proved $-\log_2(1 - \eta)$ ebits per mode as a strong converse upper bound to the secret-key agreement capacity [41].
- [33] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, *Scientific Reports* **6**, 20463 (2016), arXiv:1509.08435.
- [34] N. Sinclair, E. Saglamyurek, H. Mallahzadeh, J. A. Slater, M. George, R. Ricken, M. P. Hedges, D. Oblak, C. Simon, W. Sohler, and W. Tittel, *Physical review letters* **113**, 053603 (2014).
- [35] K. Azuma, K. Tamaki, and H.-K. Lo, *Nature communications* **6**, 6787 (2015).
- [36] F. Ewert, M. Bergmann, and P. van Loock, *Physical Review Letters* **117**, 210501 (2016).
- [37] E. Schoute, L. Mancinska, T. Islam, I. Kerenidis, and S. Wehner, arXiv preprint, 1 (2016), arXiv:1610.05238.
- [38] A. Acín, J. I. Cirac, and M. Lewenstein, *Nature Physics* **3**, 256 (2007), arXiv:0612167 [quant-ph].
- [39] R. Van Meter, *Quantum Networking*. (Wiley, 2014).
- [40] S. Pirandola, R. García-Patrón, S. L. Braunstein, and S. Lloyd, *Physical review letters* **102**, 050503 (2009).
- [41] M. M. Wilde, M. Tomamichel, and M. Berta, *IEEE Transactions on Information Theory* **63**, 1792 (2017), arXiv:1602.08898.

Appendix A: Distance metric for the local routing rule using L^1 norm and recursion

Our entanglement routing protocol with local link-state information uses the ‘distance’ of neighboring repeater stations from Alice and Bob to decide which mem-




FIG. 8. Entanglement generation rates with different distance metrics. R_{L^1} and R_{L^2} are evaluated using the L^1 and L^2 norms respectively. The distance metric for R_{i1} (iteration 1) is calculated using R_{L^1} , and R_{i2} (iteration 2) is calculated using R_{i1} . R_{i2} and R_{L^1} are nearly indistinguishable as they coincide.

ories at a repeater should undergo entanglement swap attempts. The results presented in the paper use the L^2 norm as the distance metric. While the L^2 norm can be easily calculated for the square grid, it may not be easily generalizable for other (e.g., non-planar) topologies. Further, even though we do not prove the rate optimality of our local link-state routing protocol, given a network topology, it is not clear whether or not the L^2 norm is the optimal distance metric to be used in our protocol.

In order to adapt our algorithm for arbitrary network topologies, and also to find a near-optimal distance metric for our algorithm, we employ the following numerical recursive method. Our evaluation begins with calculating $R_{L^1}(\mathbf{n}_1, \mathbf{n}_2)$, the entanglement generation rate achieved when our local rule is used to route entanglement between nodes \mathbf{n}_1 and \mathbf{n}_2 , using the L^1 norm as the distance metric. In Fig. 8, we plot $R_{L^1}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y) , where X and Y are the distance (in hops) between \mathbf{n}_1 and \mathbf{n}_2 along the horizontal and vertical dimensions of the square grid, respectively. The rate-distance scaling exponent for R_{L^1} is worse than that of R_{L^2} , the rate attained by our protocol, using the L^2 norm as the distance metric. Next, for every repeater node \mathbf{n} , we define distances d_A and d_B to Alice \mathbf{A} and Bob \mathbf{B} respectively, with respect to the following new distance metric (let us name this metric $i1$): $d_A := 1/R_{L^1}(\mathbf{n}, \mathbf{A})$ and $d_B := 1/R_{L^1}(\mathbf{n}, \mathbf{B})$. We then calculate $R_{i1}(\mathbf{n}_1, \mathbf{n}_2)$, the entanglement generation rate achieved when our local rule is used with the $i1$ distance metric to route entanglement between every pair of nodes \mathbf{n}_1 and \mathbf{n}_2 . In Fig. 8, we plot $R_{i1}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y) . We see that the rate-distance scaling achieved by R_{i1} is even lower than that of R_{L^1} . However, when we go through the second iteration of the algorithm—i.e., define distance metric $i2$, under which




FIG. 9. $f(p, q)/pq$ quantifies the improvement in the scaling of $R_{\text{loc}}(p, q)$ with respect to $R_{\text{lin}}(p, q)$ with respect to the Alice-Bob Manhattan distance, n . $f(p, q)/pq$ increases as p is reduced in $[1, p_c]$ but changing q has a negligible effect.

$d_A = 1/R_{i1}(\mathbf{n}, \mathbf{A})$ and $d_B = 1/R_{i1}(\mathbf{n}, \mathbf{B})$, and use our local rule to evaluate $R_{i2}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y) —we find that the resulting rate R_{i2} is almost the same (visually indistinguishable in the plot) as R_{L^2} , the rate we obtained directly when using the L^2 norm as the distance metric. This suggests that: (a) for the square grid (and presumably for any planar network topology) the L^2 norm metric might be near-optimal for use within our local rule, and that (b) for any given network topology, one could potentially pre-compute the optimal distance metric by a recursive strategy on the given topology using the L^1 norm as the starting point. However, there are instances where our local rule does not give the rate-optimal local routing rule. As an example, when $p = 1$ and $q = 1$, it is possible to find four disjoint paths without any link-state knowledge (the links are all deterministic) and the optimal rate is four ebits/cycle for any location of Alice and Bob. However, the fact that we are trying to route every flow through the best possible path without any coordination between different flows leads to collisions, which results in a rate that is below the optimal rate of four ebits/cycle. Finding the rate-optimal local routing rule across different parameter values is left for future research.

Appendix B: Multipath rate advantage

1. Numerical Evaluation

The goal this subsection is to quantify the improvement in the rate-vs.-distance exponent achieved by our local rule over that of a linear chain along the shortest path, for all possible pairs of values of p and q . Fig. 4(c) shows this improvement, i.e., that of $R_{\text{loc}}(p, q)$




FIG. 10. Network used to prove the lower bound on entanglement generation rate with our local routing rule which shows that scaling of the rate with Alice-Bob manhattan distance for our rule is better than the scaling of the rate along a linear repeater chain along the shortest path between Alice-Bob.

compared to $R_{\text{lin}}(p, q)$, for $p = 0.6$ and $q = 0.9$. Clearly, $R_{\text{lin}}(p, q) = (pq)^n(p)/q \sim (1/q)[pq]^n$, where n is the Manhattan distance between Alice and Bob. We have numerically verified that $R_{\text{loc}}(p, q) \sim g(p, q)[f(p, q)]^n$ for n large. We hence quantify the rate improvement by numerically evaluating the ratio $f(p, q)/(pq)$ exhaustively for all $(p, q) \in [0, 1] \times [0, 1]$, using Monte Carlo simulations. The results are shown in Fig. 9, for configurations of Alice and Bob located along 45° with respect to the grid axes. We see that $f(p, q)/pq$ increases as p decreases in $[p_c, 1]$, but changing q has a negligible effect on this ratio.

2. Analytical lower bound on the rate achieved by the local routing rule

In this subsection, we derive an analytical lower bound on the entanglement generation rate attained by our local routing rule (using the L^2 norm as the distance metric), with the objective of demonstrating multi-path routing advantage, i.e., the rate-vs.-distance scaling attained by our local rule is strictly better than that attained by a linear repeater chain along the shortest path between Alice and Bob.

Consider routing entanglement between Alice and Bob located at (X, Y) and $(X+n, Y)$ respectively, i.e., n hops apart along the X dimension of the square lattice. We will evaluate a lower bound on R_{loc} by only evaluating the rate contributions from paths in which all the (external) links belong to the set of black dashed links shown in Fig. 10. The choice of internal links made at repeater nodes proceed as usual per our local rule. As a result, there are instances in which our local rule routes entanglement through paths comprising not just the black links, resulting in flows that do not contribute to our rate lower bound.

We will refer to Fig. 10 for the ensuing discussion. Recall that external links succeed (are ‘up’) with probability p and fail (are ‘down’) with probability $1 - p$, whereas internal links succeed with probability q . Consider $P(A \leftrightarrow v)$, the probability that there is a path between Alice A and repeater v that uses only black links.

$P(A \leftrightarrow v)$ includes the probability of making the required internal links to create a path between A and v , but not the probability of any internal links at the end points A or v . It is easy to see that in any given time step, there can be no more than one edge-disjoint path between A and v along the black dashed links, since link 8 must be part of the path. Let $l \left(\tilde{l} \right)$ be the event that the external link l is up (down). Further, note that at any given time step, of all the possible (0, 1 or 2) internal links attempted by our local rule at a repeater node, only one internal link, if successful, contributes to $A \leftrightarrow v$. Let $l - m$ be the event that the internal link attempted at a repeater node to connect external links l and m is successful. If links 1 and 8 are both up, node u attempts to connect those two links based on our local rule, regardless of the other links. If links 2, 3, 4 and 8 are up, but 1, 5, 6, 7 and 9 are down, u attempts to connect 4 and 8, z attempts to connect 3 and 4 and y attempts to connect 2 and 3. Considering these two possibilities, we have

$$\begin{aligned} P(A \leftrightarrow v) &> \Pr(1, 8, 1-8) \\ &+ \Pr(2, 3, 4, 8, \tilde{1}, \tilde{5}, \tilde{6}, \tilde{7}, \tilde{9}, 2-3, 3-4, 4-8) \\ &= [p + p^3(1-p)^5q^2]pq \\ &= p'pq, \end{aligned} \quad (\text{B1})$$

where $p' = p + p^3(1-p)^5q^2 > p$.

$P(v \leftrightarrow x)$ is the probability that there is a path between v and x that uses only black links (the probability of internal link successes at the end points v and x are not included). $P(v \leftrightarrow x)$ and $P(A \leftrightarrow v)$ are not independent events because they both involve link 9. $P(v \leftrightarrow x|A \leftrightarrow v)$ is the probability that there exists a path along black dashed lines between v and x given that a path along black dashed lines exists between A and v . We now show that $P(v \leftrightarrow x|A \leftrightarrow v) > P(v \leftrightarrow x)$.

$$\begin{aligned} P(v \leftrightarrow x|A \leftrightarrow v) &= P(v \leftrightarrow x|A \leftrightarrow v, 9)\Pr(9|A \leftrightarrow v) + \\ &P(v \leftrightarrow x|A \leftrightarrow v, \tilde{9})\Pr(\tilde{9}|A \leftrightarrow v) \\ &= P(v \leftrightarrow x|9)\Pr(9|A \leftrightarrow v) + \\ &P(v \leftrightarrow x|\tilde{9})\Pr(\tilde{9}|A \leftrightarrow v) \\ &= P(v \leftrightarrow x|9)\left(1 - \Pr(\tilde{9}|A \leftrightarrow v)\right) + \\ &P(v \leftrightarrow x|\tilde{9})\Pr(\tilde{9}|A \leftrightarrow v) \\ &= \Pr(\tilde{9}|A \leftrightarrow v) \times \\ &\left(P(v \leftrightarrow x|\tilde{9}) - P(v \leftrightarrow x|9)\right) \\ &+ P(v \leftrightarrow x|9) \end{aligned} \quad (\text{B2})$$

where $P(v \leftrightarrow x|A \leftrightarrow v, 9) = P(v \leftrightarrow x|9)$ and $P(v \leftrightarrow x|A \leftrightarrow v, \tilde{9}) = P(v \leftrightarrow x|\tilde{9})$ because link 9 being up or down is the only probabilistic event that influences both $P(A \leftrightarrow v)$ and $P(v \leftrightarrow x)$. Further,

$$\begin{aligned}
P(A \leftrightarrow x) &= P(A \leftrightarrow v)qP(v \leftrightarrow x|A \leftrightarrow v) \\
&> P(A \leftrightarrow v)P(v \leftrightarrow x)q \\
&= (P(A \leftrightarrow v))^2 q \\
&= (p'pq)^2 q,
\end{aligned} \tag{B4}$$

$$\begin{aligned}
P(v \leftrightarrow x) &= P(v \leftrightarrow x|9)\Pr(9) + P(v \leftrightarrow x|\tilde{9})\Pr(\tilde{9}) \\
&= P(v \leftrightarrow x|9)\left(1 - \Pr(\tilde{9})\right) + \\
&\quad P(v \leftrightarrow x|\tilde{9})\Pr(\tilde{9}) \\
&= \Pr(\tilde{9})\left(P(v \leftrightarrow x|\tilde{9}) - P(v \leftrightarrow x|9)\right) \\
&\quad + P(v \leftrightarrow x|9).
\end{aligned} \tag{B3}$$

Comparing B2 and B3, $\Pr(\tilde{9}|A \leftrightarrow v) = \Pr(\tilde{9})\Pr(A \leftrightarrow v|\tilde{9})/\Pr(A \leftrightarrow v) > \Pr(\tilde{9})$ because $\Pr(A \leftrightarrow v|\tilde{9}) > \Pr(A \leftrightarrow v)$ following equation B1. Similarly, $\left(P(v \leftrightarrow x|\tilde{9}) - P(v \leftrightarrow x|9)\right) > 0$. Hence, $P(v \leftrightarrow x|A \leftrightarrow v) > P(v \leftrightarrow x)$.

From Fig. 10, we can see that in order to get a path along black dashed lines from A to x , there must be a path along black dashed lines from A to v and from v to x , and the internal link at v must succeed. Therefore,

where we use symmetry between $A \leftrightarrow v$ and $v \leftrightarrow x$ in the third line. Repeating this for all repeaters between Alice and Bob, it is easy to see that

$$\begin{aligned}
R_{\text{loc}} > P(A \leftrightarrow B) &> p'^{\lceil n/2 \rceil} p^{\lfloor n/2 \rfloor} q^{n-1} \\
&\geq \left(\sqrt{p'p}\right)^n q^{n-1} \\
&= \left[\left(\sqrt{p'p}\right)q\right]^n q^{-1} \\
&= (pq)^{\beta n} q^{-1},
\end{aligned} \tag{B5}$$

where $\lceil n/2 \rceil$ is the smallest integer greater than or equal to $n/2$ and $\lfloor n/2 \rfloor$ is the largest integer smaller than or equal to $n/2$. The second inequality uses the fact that $p' > p$ and $n > 0$. $\beta = \log\left[\left(\sqrt{p'p}\right)q\right]/\log[pq] < 1$ because $p < p' < 1$ and $q < 1$.

Therefore, since $R_{\text{loc}} > (pq)^{\beta n} q^{-1}$ with $\beta < 1$ and $R_{\text{lin}} = (pq)^n q^{-1}$, the exponent in the scaling with n is smaller in R_{loc} compared to R_{lin} , i.e. the rate-vs.-distance scaling is better with multi-path routing. Using a similar reasoning, it is easy to see that the same is true even when Alice and Bob are at located at different Y coordinates. It should be noted that the lower bound we derive here is not meant to be tight (see Section B1 for a full numerical evaluation of the exponents for R_{loc} and R_{lin}). The only purpose of this subsection was to prove that the rate-vs.-distance scaling for entanglement routing strictly benefits from multi-path routing.