

Continuously distributing entanglement in quantum networks with regular topologies

Lars Talsma¹,[✉] Álvaro G. Iñesta^{1,2,3} and Stephanie Wehner^{1,2,3,*}

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

²*EEMCS, Quantum Computer Science, Delft University of Technology, 2628 CD Delft, The Netherlands*

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



(Received 3 March 2024; revised 18 June 2024; accepted 1 July 2024; published 19 August 2024)

Small interconnected quantum processors can collaborate to tackle quantum computational problems that typically demand more capable devices. These linked processors, referred to as quantum nodes, can use shared entangled states to execute nonlocal operations. As a consequence, understanding how to distribute entangled states among nodes is essential for developing hardware and software. We analyze a protocol where entanglement is continuously distributed among nodes that are physically arranged in a regular pattern: a chain, a honeycomb lattice, a square grid, and a triangular lattice. These regular patterns allow for the modular expansion of networks for large-scale distributed quantum computing. Within the distribution protocol, we investigate how nodes can optimize the frequency of attempting entanglement swaps, trading off multiple entangled states shared with neighboring nodes for fewer states shared with non-neighboring nodes. We evaluate the protocol's performance using the virtual neighborhood size—a metric indicating the number of other nodes with which a given node shares entangled states. Employing numerical methods, we find that nodes must perform more swaps to maximize the virtual neighborhood size when coherence times are short. In a chain network, the virtual neighborhood size's dependence on swap attempt frequency differs for each node based on its distance from the end of the chain. Conversely, all nodes in the square grid exhibit a qualitatively similar dependence of the virtual neighborhood size on the swap frequency.

DOI: [10.1103/PhysRevA.110.022429](https://doi.org/10.1103/PhysRevA.110.022429)

I. INTRODUCTION

A quantum network is a system of interconnected quantum devices that extends beyond the capabilities of classical networks [1]. Such devices, known as quantum *nodes*, can be connected over long and short distances. Leveraging quantum-mechanical effects like entanglement, these nodes enable a variety of technologies. For instance, the quantum internet aims to facilitate quantum communication between any two points on Earth [2,3], allowing applications such as quantum key distribution [4,5] and secure access to remote quantum computers [6]. Over short distances, dense arrays of closely connected nodes can cooperate to solve challenging quantum computational problems by distributing the workload among them [7,8].

In a network of nodes connected over short distances, we assume that we can design the topology, unlike in long-distance networks whose topologies might have logistical constraints such as the location of cities. Specifically, we investigate network topologies where nodes form a regular pattern, which constitutes a modular and scalable architecture. In particular, we consider networks with a *regular topology* where nodes are regularly spaced and connected over the same physical distance, each having the same number of physical neighbors d . The nodes can form a chain for $d = 2$, a honeycomb lattice for $d = 3$, a square lattice for $d = 4$, and a

triangular lattice for $d = 6$ (Fig. 1). These lattices regularly tile the plane.

Such a modular network design offers a scalable approach to, e.g., distributed quantum computing. In such a setting, nodes can implement nonlocal operations using entangled states shared with other nodes, allowing for universal quantum computation [9–11]. Consequently, nodes can scale the number of qubits available for computation by sharing such bipartite states—*entangled links*—with many different nodes. Additionally, since entangled links are consumed in nonlocal operations, nodes require many links if they want to implement many nonlocal operations.

To distribute many links among many different nodes, we consider a protocol for *continuous distribution* (CD) of entanglement, where nodes continuously distribute links among them [12]. Compared to *on-demand* protocols where nodes explicitly request entanglement [13–15], CD protocols do not involve a routing problem to establish links among the nodes requesting them [12,13,16]. Solving routing problems can be computationally demanding for large quantum networks. Hence, in networks with many nodes, employing a CD protocol could be a suitable approach to distributing entangled links at a high rate.

We investigate the performance of a CD protocol that aims to distribute many entangled links among many different nodes in quantum networks with a regular topology, employing the quantum network model of Ref. [12]. In such networks, pairs of nodes are physical neighbors when they are connected by physical channels such as optical fibers [17].

*Contact author: s.d.c.wehner@tudelft.nl




FIG. 1. Quantum networks with a regular topology. Regularly spaced quantum nodes connected over identical physical channels each have d physical neighbors and form a chain, a honeycomb lattice, a square grid, and a triangular lattice, as exemplified by these fragments of larger networks.

Nodes can store quantum information in the form of qubits, and nodes can entangle these qubits to form entangled links. We assume that nodes can share multiple links and that each node has a “large-enough” number of quantum memories with finite coherence times (see Sec. III A for details). We assume that all nodes are identical and that they are connected over identical physical channels (in particular, all neighboring nodes are at the same physical distance). Moreover, each node has the same number of physical neighbors in the resulting networks. Other unit cells can tile the plane (see Fig. 2(a) in Ref. [18]), although these generally involve nodes with different numbers of physical neighbors or physical channels with different lengths. Therefore, we restrict ourselves to the regular patterns from Fig. 1. Last, we consider networks with and without boundaries—*finite* and *infinite* networks.

To distribute useful entanglement among the networked nodes, (i) physical neighbors can generate shared entangled links in a heralded fashion [19], (ii) remote nodes can transform two links shared with an intermediary node into a longer link in an entanglement swap [20], and (iii) nodes can discard links when their quality has decreased too much due to decoherence and entanglement swaps [12,21,22]. Following Ref. [12, Algorithm 1], we encapsulate these operations in a simple CD protocol that discretizes time and prescribes the nodes what to do in each time step. Within the protocol, the nodes can adjust how frequently they attempt swaps to modify the distribution of entangled links: If no swaps are performed, then entanglement will only be shared among physical neighbors; if swaps are often performed, then entanglement will mostly be shared among physically distant nodes.

With the objective of distributing many entangled links among many different nodes, we evaluate the performance of the CD protocol using the performance metrics introduced in Ref. [12]. At a specific time, the *virtual neighborhood size* indicates the number of nodes any node shares entangled links with, and the *virtual node degree* indicates how many entangled links any node shares with other nodes. The virtual neighborhood size and virtual node degree explicitly consider the time dependence of entangled links as links can be created and removed over time. For example, in a distributed quantum computing setting, a large virtual neighborhood size means that a node can perform nonlocal operations with many different nodes, and a large virtual node degree means that a node can implement many nonlocal operations. Other approaches to analyzing entanglement distribution in quantum networks include evaluating the time it takes to distribute end-to-end

entanglement among specific pairs of nodes [15,23–26]. However, such a metric is better suited to evaluate the performance of on-demand protocols where the goal is to optimize the time it takes to generate entanglement among a set of end nodes. The virtual neighborhood size and virtual node degree are more suitable for evaluating CD protocols, for example capturing the goal of distributing many links among many different nodes.

In this paper, we employ numerical methods to evaluate the performance of a CD protocol distributing entanglement in regular-topology quantum networks to maximize the virtual neighborhood size and virtual node degree. Our main findings offer design heuristics for CD protocols in quantum networks with regular topologies:

(i) When coherence times are short, swaps must be performed more frequently to maximize the virtual neighborhood size. Intuitively, nodes must make good use of the links before the links are cut off.

(ii) The impact of network boundaries on the protocol’s performance depends on the network topology. In a finite chain, the dependence of the virtual neighborhood size on the swap attempt frequency is different for each node depending on the node’s distance to the edge of the chain. In contrast, for networks with a square-lattice topology, the virtual neighborhood size of all nodes behaves similarly as a function of the swap frequency.

This paper is structured as follows. In Sec. II, we present the quantum network model by discussing networks with a regular topology and how nodes can distribute entanglement in such networks. Furthermore, we adopt a simple CD protocol to facilitate entanglement distribution and define the performance metrics. Subsequently, in Sec. III, we use these metrics to evaluate the performance of the CD protocol in quantum networks with a regular topology. Finally, in Sec. IV, we reflect on the results and discuss potential future work.

II. QUANTUM NETWORK MODEL

We introduce our quantum network model in this section. We first present the physical topologies of the networks we have investigated. Then, we discuss how quantum nodes can use entanglement generation, entanglement swaps, and the removal of low-fidelity links to distribute useful entanglement in regular networks. Accordingly, we adopt an entanglement distribution protocol that the nodes use for distributing entanglement among them. Finally, we discuss the performance metrics that we have used to evaluate the performance of this protocol.

We adopt the quantum network model of Ref. [12] (see Fig. 2 for an illustration). In this model, *nodes* can generate, process, and store quantum information in the form of *qubits*. Such nodes can send quantum information to each other over *physical channels*. Nodes connected via physical channels are physical neighbors. Two nodes can share any number of entangled qubits, where we refer to these shared bipartite states as *entangled links*. Nodes can employ qubit platforms such as nitrogen-vacancy centers in diamond [19,27] and trapped ions [28,29] and be connected over physical channels such as optical fibers [17] and free space [30,31].




FIG. 2. Example quantum network. Nodes can share any number of entangled links, either directly generating entanglement over physical channels or swapping entanglement to create longer-distance links.

A. Network topology

We consider arrays of nodes connected over short distances for large-scale quantum networks. In this scenario, we assume that we are free to design the network topology (in contrast to long-distance networks for quantum communication whose topologies might have logistical constraints such as the location of cities or are utilizing existing optical fibers). Hence, to modularly scale the network size, we investigate network topologies where nodes form a regular pattern.

In particular, we consider quantum networks with a *regular topology* where nodes are regularly spaced and connected over identical physical channels such that each node has the same number d of physical neighbors. We say that such a network has a *physical node degree* d . The quantum nodes form a chain for $d = 2$, a honeycomb lattice for $d = 3$, a square lattice for $d = 4$, and a triangular lattice for $d = 6$ (Fig. 1). The honeycomb, square, and triangular lattices tile the plane regularly.

We investigate *infinite* and *finite* quantum networks, that is, networks without and with boundaries. The virtual neighborhood size and virtual node degree of nodes in an infinite network approximate, on average, those of nodes far from the network boundaries in large-scale quantum networks (see Sec. III B for details on how “far”). Furthermore, infinite networks provide a convenient platform for performance analysis as each node behaves equivalently due to the network’s translational symmetries. We note that nodes on the boundary of finite networks have fewer than d physical neighbors.

B. Network dynamics

To distribute useful entanglement, (i) physical neighbors attempt to generate entanglement in a heralded fashion, (ii) nodes swap entanglement to convert short-distance entangled links into longer links, and (iii) nodes discard low-fidelity entangled links that are of insufficient quality for their intended purpose. We briefly discuss these three operations and summarize the network model parameters in Table I (see Appendix A for more details).

1. Generating entanglement

Two physical neighbors can attempt to generate a shared entangled link in a heralded fashion [19]. The entanglement generation attempt heralds success with a probability p_{gen} and fails to generate an entangled link with a probability $1 - p_{\text{gen}}$. To model quantum noise, we apply a depolarizing channel (a

TABLE I. Quantum network model parameters. The parameters T , t_{cut} , F_{new} , M , and F_{min} must satisfy (2).

Physical topology	
d	Physical node degree, number of physical neighbors
Hardware	
p_{gen}	Probability of successfully heralding entanglement
p_{swap}	Probability of successfully swapping entanglement
T	Coherence time (exponential decay rate of fidelity)
F_{new}	Entanglement generation fidelity
Software	
t_{cut}	Cutoff time
M	Maximum swap distance
F_{min}	Minimum required entangled link fidelity
Protocol	
q	Probability of attempting an entanglement swap

worst-case noise model) [32] to the Bell state $|\phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. Then, nodes generate entangled links of the Werner form [33]

$$\rho = \frac{4F - 1}{3} |\phi^+\rangle\langle\phi^+| + \frac{1 - F}{3} \mathbb{I}_4, \quad (1)$$

where $F \equiv \langle\phi^+|\rho|\phi^+\rangle$ is the fidelity [34] of the generated state ρ to the target state $|\phi^+\rangle$ and \mathbb{I}_4 the four-dimensional identity. We assume that all nodes generate entangled Werner states with the same fidelity $F = F_{\text{new}}$.

2. Swapping entanglement

Two nodes that are not connected by a physical channel can create shared entangled links by swapping entanglement via an intermediary node [20]. For example, suppose that nodes A and B do not share a physical channel. However, nodes A and B both share an entangled link with an intermediary node I with fidelities F_{AI} and F_{BI} . The intermediary node I can do a Bell-state measurement on the qubits storing the links. Then, using local operations and classical communication, the nodes perform an entanglement swap. Specifically, the nodes transform the initial two links into a new link between nodes A and B with fidelity $F_{\text{AB}} \leq F_{\text{AI}}, F_{\text{BI}}$ [15,35]. Nodes successfully execute the swap with a probability p_{swap} and fail to generate a longer link with a probability $1 - p_{\text{swap}}$ (consuming the initial links).

3. Discarding entanglement

To ensure that entangled links are of sufficient quality for, e.g., distributed quantum computing applications, nodes discard low-fidelity entangled links. We consider two fidelity-decreasing processes.

Qubits interact with their environment, and the fidelity of the links they store decreases over time—links *decohere*. We assume that the link fidelity decays exponentially with time, where we characterize the decay rate by an abstract coherence time T . To ensure that the fidelity of all entangled links exceeds some threshold fidelity F_{min} , nodes discard entangled links that are stored longer than a *cutoff time* t_{cut} [15,21,22]. In particular, nodes monitor the *age* of the entangled links—the

ALGORITHM 1. Example CD protocol.

Inputs: A quantum network with an arbitrary configuration of entangled links. The network has a regular topology characterized by the physical node degree d . The hardware is described by p_{gen} , p_{swap} , T and F_{new} , and we choose t_{cut} , M and F_{min} (see Table 1 for details). Nodes can tune the swap attempt probability q to improve the protocol's performance.

Output: A quantum network with an updated configuration of entangled links.

Algorithm:

1. **Cutoff time:** Nodes discard entangled links with ages equal to the cutoff time t_{cut} . Nodes first apply cutoffs to ensure they do not use old links later in the protocol.
2. **Entanglement generation:** Physical neighbors attempt to generate entangled links and successfully herald a link of fidelity F_{new} with a probability p_{gen} .
3. **Entanglement swapping:** All nodes simultaneously perform the following steps:
 - 3.1 Nodes randomly choose a link from their memory.
 - 3.2 Nodes choose a second link randomly from the set of links stored in differently oriented qubits.
 - 3.3 Nodes attempt to swap the two entangled links with a probability q and succeed with a probability p_{swap} to create a longer link. When nodes do not attempt to swap, the initial links are not used in further attempts. When the swap fails, nodes discard the initial links.

Nodes repeat steps 3.1–3.3 until no more swaps are possible. Nodes are unaware of the swaps of other nodes.
4. **Maximum swap distance:** The nodes communicate which swaps they have attempted. Then, nodes discard entangled links that have been generated from more than M short-distance links (generated between physical neighbors).

time elapsed since the creation of the link—and subsequently discard links with an age equal to the cutoff time.

The entangled link fidelity generally decreases with the number of swaps it has been involved in. Again, to ensure that the fidelity of all entangled links exceeds some minimum fidelity F_{min} , nodes discard entangled links that are the fusion of more than M short-distance links (generated between physical neighbors) [12]. We refer to M as the *maximum swap distance*.

Combining these requirements, nodes that generate entangled links with a fidelity F_{new} and demand links with a minimum fidelity F_{min} must satisfy the relation [15]

$$t_{\text{cut}} \leq -T \ln \left[\frac{3}{4F_{\text{new}} - 1} \left(\frac{4F_{\text{min}} - 1}{3} \right)^{\frac{1}{M}} \right]. \quad (2)$$

C. Entanglement distribution protocol

To distribute entanglement among the nodes, we employ a simplified version of the CD protocol from Ref. [12, Algorithm 1]. In our CD protocol (Algorithm 1), physical neighbors generate entangled links in a heralded fashion, non-neighboring nodes swap entanglement via intermediary nodes, and nodes discard links when their fidelity has decreased too much.

In the CD protocol, nodes attempt entanglement swaps with a probability q . This probability is the only protocol

parameter that nodes can “tune” to improve the performance of the entanglement distribution process. During each iteration of the CD protocol, nodes attempt swaps until no more swaps are possible.

We assume that nodes distribute preshared entanglement, so we omit the consumption of links in applications (in contrast to the CD protocol of Ref. [12, Algorithm 1]). The CD protocol discretizes time and defines the operations that all nodes perform simultaneously. The coherence time T and cutoff time t_{cut} are expressed in units of this discretized time. Furthermore, we assume that each qubit can only generate entanglement with a fixed neighboring node. Then, nodes label the qubit *orientation* as the direction of the physical neighbor. Nodes only swap entangled links from qubits with different orientations to avoid “unnecessary” swaps. In a chain, for example, a node only swaps pairs of links where one link is stored in a left-oriented qubit and the other in a right-oriented qubit. This prevents nodes from creating links between two qubits in the same node (for example, preventing a node from swapping two links with their left physical neighbor) or links that could have been generated directly via a physical channel. As all nodes implement the protocol simultaneously, there is no time to communicate and coordinate more elaborate swap strategies. Last, the protocol does not consider entanglement distillation (it could be included by modifying the network parameters; see Appendix A).

D. Performance metrics

We evaluate the performance of the CD protocol (Algorithm 1) in quantum networks with regular topologies. In applications such as distributed quantum computing, the quantum nodes would likely benefit from (i) entangled links with many different nodes to scale the number of qubits available for computation and (ii) many entangled links with other nodes to implement many nonlocal operations. With these objectives, we employ performance metrics for quantum networks as introduced by Ref. [12]:

Definition 1 ([12]). The *virtual neighborhood* of node i at time t , $V_i(t)$, is the set of nodes that share an entangled link with node i at time t . Two nodes are *virtual neighbors* if they share at least one entangled link. The *virtual neighborhood size* is defined as $v_i(t) \equiv |V_i(t)|$.

Definition 2 ([12]). The *virtual node degree* of node i at time t , $k_i(t)$, is the number of entangled links connected to node i at time t .

These performance metrics capture the objective of distributing many entangled links between many nodes, and explicitly incorporate the time-dependent dynamics of quantum networks. In a distributed quantum computing setting, a large $v_i(t)$ indicates that nodes share entangled links with many different nodes (increasing the number of qubits available for computation), and a large $k_i(t)$ means that nodes can implement many nonlocal operations.

The performance metrics are stochastic processes, i.e., (time) sequences of random variables. Considering this time dependence, we investigate the steady-state expected value of the performance metrics to learn more about the long-term behavior of the network. When a quantum network is running the CD protocol of Algorithm 1, Ref. [12] showed that there

is a unique steady-state value for the expected virtual neighborhood size,

$$v_i \equiv \lim_{t \rightarrow \infty} \mathbb{E}[v_i(t)], \quad (3)$$

and the expected virtual node degree,

$$k_i \equiv \lim_{t \rightarrow \infty} \mathbb{E}[k_i(t)]. \quad (4)$$

We employ discrete-time network simulations that implement the CD protocol (Algorithm 1) to estimate v_i and k_i ; see Appendix B for details.

III. PERFORMANCE OF CD PROTOCOLS IN REGULAR-TOPOLOGY NETWORKS

In this section, we evaluate the performance of CD protocols in quantum networks with regular topologies. Nodes can optimize the performance of the entanglement distribution protocol by varying the probability of attempting swaps, q . In Sec. III A, we investigate the influence of network parameters (the coherence time and the entanglement generation fidelity) on the protocol's performance in *infinite* regular networks (i.e., without boundaries). Then, in Sec. III B, we investigate the influence of network boundaries on the virtual neighborhood size of (finite) chains and square lattices.

A. Infinite networks

The behavior of nodes in infinite networks approximates that of nodes far from the network boundaries in large-scale quantum networks. Furthermore, infinite networks present a convenient setting for performance analysis as all nodes behave equivalently due to the network's translational symmetries. We investigate the influence of varying network parameters (the coherence time T and the entanglement generation fidelity F_{new}) on the performance metrics using infinite networks.

We assume that nodes generate entanglement deterministically ($p_{\text{gen}} = 1$). Although entanglement generation is generally probabilistic, deterministic entanglement distribution protocols can guarantee the delivery of entangled states at specified time intervals [36,37]. Deterministic entanglement generation provides a convenient analysis platform as each node generates the same number of entangled links. In Appendix A, we provide a study of the effect of $p_{\text{gen}} < 1$ on the protocol's performance. For $p_{\text{gen}} = 1, \frac{1}{2}, \frac{1}{4}$, we observe that the optimal swap attempt probability q scales approximately linearly with p_{gen} .

Furthermore, we assume that nodes execute swaps deterministically ($p_{\text{swap}} = 1$). Qubit platforms such as nitrogen-vacancy centers in diamond [38] and trapped ions [39] can realize Bell-state measurements that succeed deterministically to facilitate the entanglement swaps. Deterministic swaps are convenient for analysis as we do not have to consider failed swaps. When swaps do fail, the virtual neighborhood size and the virtual node degree decrease (see Appendix A).

Last, we assume that nodes have a “large-enough” number of memories. In particular, nodes can store all generated entangled links until they discard them when the links age to the cutoff time; that is, we assume that nodes have at least $d t_{\text{cut}}$ qubits. In our numerical experiments, this corresponds to

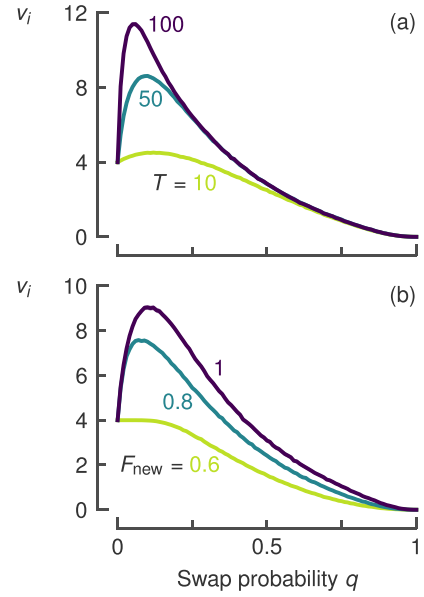


FIG. 3. The optimal probability of attempting to swap q depends on the coherence time T and entanglement generation fidelity F_{new} . Virtual neighborhood size v_i of a node in an infinite square lattice ($d = 4$) as a function of q for (a) $T = 10, 50, 100$ time steps (cutoff time $t_{\text{cut}} = 2, 11, 22$ time steps) and (b) $F_{\text{new}} = 0.6, 0.8, 1$ (maximum swap distance $M = 1, 2, 4$); both colored from light to dark. For longer T , links live longer before nodes discard them (longer t_{cut}), meaning that nodes store more links and can have more virtual neighbors, increasing the maximum v_i . When links live long, swapping conservatively results in more virtual neighbors (lower optimal q). A higher F_{new} increases the maximum v_i as links can be involved in more swaps before nodes discard them (increased M). Then, swapping more frequently results in more virtual neighbors (higher optimal q). We consider (a) $F_{\text{new}} = 0.9$ ($M = 3$) and (b) $T = 50$ time steps ($t_{\text{cut}} = 11$ time steps). We assume that nodes generate entanglement and execute swaps deterministically ($p_{\text{gen}}, p_{\text{swap}} = 1$) and that nodes require a minimum link fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained with network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

the order of 10–100 memories. Currently, experiments attain memories of, e.g., 10 qubits in diamond nitrogen-vacancy centers [40], which we expect to increase in the near future to the values we use in our simulations.

Our network simulations implement the cutoff time t_{cut} and the maximum swap distance M . Then, to vary the coherence time T , we find a t_{cut} that satisfies condition (2); similarly, to vary the entanglement generation fidelity F_{new} , we find an M that satisfies (2). Illustratively, if the coherence time T is short, then nodes discard links quickly (short t_{cut}) regardless of F_{new} and M . Similarly, if F_{new} is too low, then swapping two links results in a new link with fidelity $F' < F_{\text{min}}$, i.e., nodes should not attempt swaps ($M = 1$) regardless of T and t_{cut} .

Generally, as a function of increasing swap probability, the virtual neighborhood size starts at $v_i = d$ ($q = 0$). Then v_i increases to a maximum before converging to zero as $q \rightarrow 1$ (Fig. 3 for a square-lattice network, $d = 4$). For $q = 0$, $v_i = d$ as nodes only share entangled links with their physical neigh-

bors (and discard them when the links reach the cutoff time t_{cut}). Then, as q increases, nodes attempt swaps and can share links with non-neighboring nodes. However, as nodes attempt more swaps (larger q), they also consume more links in swaps and discard more links for being involved in too many swaps (M). At first, nodes gain more new virtual neighbors as q increases. Then, at some q , losing links (due to a combination of consuming them in swaps and removing low-fidelity ones) balances out this gain in virtual neighbors; v_i reaches a maximum. As the swap probability q increases further, consumption of links in swaps and low-fidelity link removal outweigh the creation of new virtual neighbors, resulting in a decreasing v_i . Finally, as $q \rightarrow 1$, $v_i \rightarrow 0$ as nodes discard all links for being involved in too many swaps.

We see that the swap probability q that maximizes the virtual neighborhood size v_i , the *optimal* q , depends on the coherence time T and the entanglement generation fidelity F_{new} (Fig. 3). As T increases, entangled links live longer before nodes cut them off (longer t_{cut}). Consequently, nodes store more entangled links and can share entangled links with a larger set of nodes, resulting in an increased v_i . Note that v_i is bounded by a function of the cutoff time t_{cut} or the maximum swap distance M (see Appendix B). For short T and t_{cut} , nodes quickly discard links for living too long. Then, swapping frequently (relatively high q) increases v_i during the short lifetime of the links. In contrast, for relatively long T and t_{cut} , the nodes benefit from swapping more conservatively as nodes consume fewer links in swaps and discard fewer links that have been swapped too many times. That is, the optimal q decreases for longer T .

As F_{new} increases, entangled links can be involved in more swaps before nodes discard them (larger M). Consequently, nodes can attempt swaps more often (higher q) and share entangled links with nodes that are further away, resulting in an increased virtual neighborhood size v_i . That is, the optimal q increases with increasing F_{new} . However, increasing F_{new} results in diminishing gains of the virtual neighborhood size v_i . When there is “enough” time for links to exist (large T , t_{cut}), the maximum swap distance M limits v_i . In that case, v_i approaches its upper bound, i.e., nodes share entanglement with all nodes it can potentially share entanglement with. In contrast, when links exist for a limited time, nodes likely do not share entanglement with all nodes they potentially could.

As the physical node degree d increases, the maximum virtual neighborhood size v_i also increases (see Appendix C). In particular, increasing $d = 2$ to $d = 3$ increases the maximum value of v_i by more than the ratio of node degrees ($3/2$). We note that, for increasing swap distance M , the bound on v_i grows quicker in networks with $d = 3$ than those with $d = 2$ (see Appendix B). Increasing the physical node degree to $d = 4, 6$ shows diminishing returns.

The virtual node degree k_i , i.e., the total number of entangled links connected to node i , decreases monotonically for increasing swap probability q for all physical node degrees d and network parameters (see Appendix C). Nodes achieve a maximum $k_i = d t_{\text{cut}}$ when not attempting swaps ($q = 0$) as they only lose links due to cutoffs (t_{cut}). For $q > 0$, nodes also consume links in swaps or discard links for being involved in too many swaps. For $q = 1$, nodes discard all links for being involved in too many swaps; hence, $k_i = 0$.

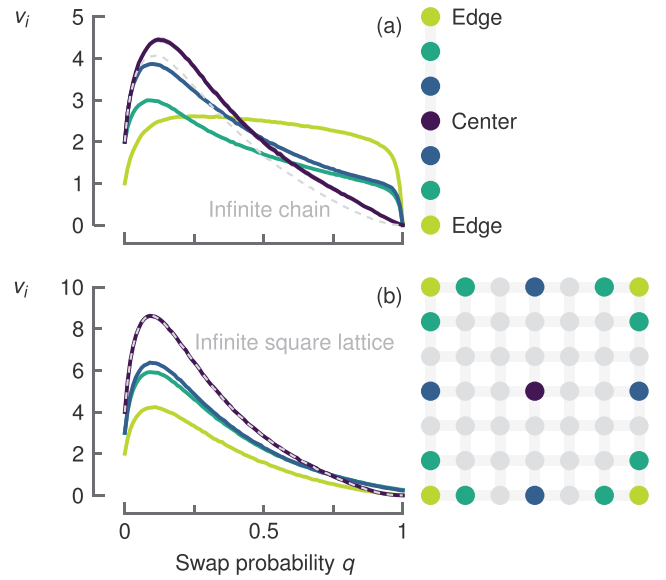


FIG. 4. In a chain, the dependence of the virtual neighborhood size v_i on the probability of attempting swaps q is different for each node depending on its distance to the edge; conversely, in a square lattice, the dependence is qualitatively the same across all nodes. For nodes near and far from the network boundary (colored from light to dark), v_i as a function of q in a finite (a) chain ($d = 2$) and (b) square lattice ($d = 4$). Edge nodes in a chain only have one physical neighbor and cannot implement swaps. Other nodes can implement swaps but may have a limited potential neighborhood and an asymmetric number of links oriented in either direction, resulting in fewer virtual neighbors. In a finite square lattice, all nodes can implement swaps, and nodes with the same number of physical neighbors have similar v_i . We consider a coherence time $T = 50$ time steps (cutoff time $t_{\text{cut}} = 11$ time steps) and an entanglement generation fidelity $F_{\text{new}} = 0.9$ (maximum swap distance $M = 3$). We assume that nodes generate entanglement and execute swaps deterministically ($p_{\text{gen}}, p_{\text{swap}} = 1$) and that nodes require a minimum link fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained with network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

B. Finite networks

We see that the effect of network boundaries on the virtual neighborhood size v_i depends on the topology (Fig. 4). Specifically, the behavior of v_i depends on the node’s distance to the edge of a chain, while, in a square lattice, v_i behaves qualitatively the same across all nodes.

In a finite chain, the virtual neighborhood size v_i of nodes with the same number of physical links to the center node (nodes symmetric around the center) behaves equivalently. Edge nodes only have one physical neighbor, meaning that $v_i = 1$ when nodes do not attempt swaps ($q = 0$). Furthermore, according to the CD Protocol (Algorithm 1), these edge nodes cannot implement swaps and hence do not consume entangled links in swaps (recall that nodes attempt swaps with entangled links stored in qubits with different orientations, one left-oriented and one right-oriented link in a chain). For $q > 0$, v_i of the edge node increases initially as other nodes in the chain attempt swaps. Then, v_i stabilizes for a wide range

of swap probabilities before $v_i \rightarrow 0$ as $q \rightarrow 1$ because nodes discard all links for being involved in too many swaps.

Nodes between the edge nodes—*interior* nodes—have virtual neighborhood sizes v_i that are qualitatively more similar to that of nodes in an infinite chain. However, the edges do have an influence. Specifically, compared to nodes in an infinite chain, v_i of interior nodes near the edge converges to zero more slowly as $q \rightarrow 1$. We try to explain this behavior with an example. First, note that interior nodes near the edge have an asymmetric number of nodes that they can share entanglement with to their left and right (recall that nodes in a chain swap one left-oriented and one right-oriented link). Now, suppose that all nodes decide to swap entanglement, except for one node that is closer than M physical links to the edge. In that case, this node discards one link that is involved in too many swaps. The other link (oriented toward the nearby edge) cannot have been involved in too many swaps and is not discarded. Hence, at the end of the protocol step, the node near the edge that did not attempt to swap entanglement retains one link, i.e., $v_i \neq 0$. Such a scenario is more likely to happen for large q . This results in a larger v_i of nodes near the edge in a finite chain compared to v_i of infinite chain nodes as $q \rightarrow 1$. However, when $q = 1$, all nodes attempt swaps; then nodes discard all links for being involved in too many swaps ($v_i = 0$). Last, nodes closer to the edge have fewer nodes they can share entanglement with, resulting in a smaller maximum v_i .

Nodes that are M or more physical links away from the chain edge (e.g., the center node in Fig. 4) do have a symmetric number of potential virtual neighbors in both directions. For these nodes, the virtual neighborhood size v_i behaves similarly as in the case of an infinite chain. Intuitively, such nodes are far enough from the boundary to not experience edge effects. Moreover, nodes that are precisely M physical links away from the edge (the center node in Fig. 4) have a higher virtual neighborhood size v_i compared to nodes in an infinite chain.

In a finite square lattice, the virtual neighborhood size v_i of nodes with the same physical node degree d_i is quantitatively similar (as boundary nodes in finite regular networks have fewer physical neighbors than interior nodes, we refer to the physical node degree of a specific node i). Corner nodes have two physical neighbors ($d_i = 2$), side nodes (on the boundary but not in the corner) have three ($d_i = 3$), and interior nodes (not on the boundary) have four ($d_i = 4$). In contrast to the finite chain, boundary nodes in a finite square lattice can implement swaps, resulting in a qualitatively similar v_i behavior of all nodes. However, boundaries do have an influence in that nodes have limited potential virtual neighborhoods when they are closer to the boundary. For instance, side nodes closer to a corner have a smaller v_i than side nodes further away from a corner. The maximum v_i of corner nodes is approximately one-half of the maximum v_i of the interior nodes, and the maximum v_i of the side nodes is about three-quarters of the maximum v_i of the interior nodes. Last, v_i of the side nodes does not converge to exactly zero as $q \rightarrow 1$ because side nodes generate an uneven number ($d_i = 3$) of entangled links per time slot, meaning that there is a nonzero probability that the nodes do not involve each link in too many swaps.

In both finite chains and square lattices, the virtual node degree k_i still monotonically decreases to zero as the swap probability $q \rightarrow 1$ (see Appendix C). However, when nodes do not attempt swaps ($q = 0$), boundary nodes have a smaller virtual node degree, $k_i = d_i t_{\text{cut}}$. In the finite chain, due to the CD protocol (Algorithm 1), edge nodes do not consume links in swaps because they cannot attempt swaps, meaning that k_i converges to zero slowly. Additionally, by the same explanation as for the virtual neighborhood size v_i , k_i of nodes near the edge decreases to zero slower than nodes in an infinite chain.

IV. DISCUSSION

We have adopted a simple protocol for the CD of entanglement among networked nodes. The nodes can optimize the probability of attempting entanglement swaps to improve the performance of the CD protocol. Using numerical methods, we have evaluated the protocol's performance in networks where nodes form a regular pattern. We have employed performance metrics that explicitly consider the time dependence of the entangled states. In particular, the virtual neighborhood of any node is the set of nodes it shares entanglement with, and the virtual node degree of any node is the number of entangled states it shares with other nodes. A large virtual neighborhood size indicates that a node shares entangled states with many different nodes, and, for instance, by using entanglement as a means of nonlocal coupling, many qubits are available for computation in distributed quantum computing. In this setting, a sizable virtual node degree indicates that a node can implement many nonlocal operations.

We present our findings as heuristics for the design of CD protocols in quantum networks with regular topologies. First, we observe that the swap attempt probability that maximizes the virtual neighborhood size depends on network parameters like the coherence time and entanglement generation fidelity. In particular, this optimal swap attempt probability decreases for a longer coherence time and increases for a higher entanglement generation fidelity. Second, the maximum virtual neighborhood size increases with the number of physical neighbors d per node—for example, going from a chain ($d = 2$) to a honeycomb lattice ($d = 3$) increases this maximum by more than the ratio of physical neighbors per node ($3/2$). Expanding the network to square-lattice ($d = 4$) and triangular-lattice ($d = 6$) topologies shows diminishing returns. Last, we see that the influence of network boundaries depends on the network topology. In a chain of nodes, being on or near a network boundary fundamentally alters the performance metrics. In contrast, in a square lattice, the metrics behave qualitatively the same across all nodes. Moreover, the performance metrics of nodes in a finite chain only start approaching those of nodes in an infinite chain when they are further from the boundary than the maximum swap distance M .

We have limited our analysis to regular topologies in one and two dimensions, but we could extend it to three-dimensional regular networks, optimally filling rooms with hypothetical future clusters of quantum computing nodes. We have investigated the optimal swap attempt probability while relating the cutoff time with the coherence time and

the maximum swap distance with the entanglement generation fidelity; given a setup with a specific coherence time and entanglement generation fidelity, future research could simultaneously optimize over the swap probability, cutoff time, and maximum swap distance to maximize the virtual neighborhood size. Furthermore, our CD protocol delivers preshared entangled links. That is, our analysis omitted the consumption of entangled links associated with, for example, implementing nonlocal operations in distributed quantum computing. Such consumption could alter the optimal swap probability and the maximum virtual neighborhood size. Future work could implement link consumption like the analysis of a network with a tree topology from Ref. [12].

Further analysis could involve more elaborate CD protocols. For example, the protocol could pair entangled links used in swaps more efficiently (instead of randomly) to increase the virtual neighborhood size, potentially using information about the topology of the network. We assumed that qubits have a “large-enough” number of quantum memories; if that assumption is not met, then it would be interesting to investigate how a protocol could optimally utilize the limited number of memories. Last, protocols could subject the performance metrics to certain constraints. For instance, nodes could demand some minimum virtual neighborhood size v_{\min} to ensure a minimum number of qubits available for computation or demand a minimum virtual node degree k_{\min} to ensure that nodes can implement a minimum number of nonlocal operations. Nodes could also require the ratio k_i/v_i to attain some minimum value, resulting in a multiobjective optimization problem (similar to meeting the *quality-of-service* requirements investigated by Ref. [12]).

The data and the code to generate, process, and plot the data can be found in Ref. [41].

ACKNOWLEDGMENTS

We thank Janice van Dam, Francisco Ferreira da Silva, and Bethany Davies for their feedback on this paper. L.T. acknowledges S.W.’s financial support to realize this project. Á.G.I. acknowledges financial support from the Netherlands Organisation for Scientific Research (NWO/OCW), as part of the Frontiers of Nanoscience program. S.W. acknowledges support from an NWO VICI grant.

Á.G.I. defined the project, and L.T. developed and analyzed the network simulations. L.T. prepared the paper. Á.G.I. and S.W. supervised the project and provided active feedback at every stage of the project.

APPENDIX A: QUANTUM NETWORK MODEL DETAILS

In this Appendix, we provide more details on the operations and associated parameters we consider in our quantum network model as discussed in Sec. II and motivate the choices of parameters used in Sec. III.

1. Generating entanglement

Two physical neighbors herald the successful generation of entanglement [19] with a probability p_{gen} , and raise a failure

flag with a probability $1 - p_{\text{gen}}$. To model quantum noise, we apply a depolarizing channel (a worst-case noise model) [32] to the Bell state $|\phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. As a result of the depolarizing channel, the initial state $|\phi^+\rangle$ is unaffected with some probability x . However, with a probability $1 - x$, the initial state depolarizes to the completely mixed state $\mathbb{I}_4/4$, where \mathbb{I}_4 is the four-dimensional identity [32]. Consequently, nodes generate entangled links of the Werner form [33]

$$\rho = \frac{4F - 1}{3} |\phi^+\rangle\langle\phi^+| + \frac{1 - F}{3} \mathbb{I}_4, \quad (1 \text{ revisited})$$

where $F = F(\rho, |\phi^+\rangle) \equiv \langle\phi^+|\rho|\phi^+\rangle = \frac{3}{4}x + \frac{1}{4}$ is the fidelity [34] of the generated state ρ to the target state $|\phi^+\rangle$.

Generally, heralded entanglement generation attempts succeed probabilistically, but protocols can guarantee the generation of entangled links at specified intervals [36,37]. For example, experiments successfully generate entanglement with a probability $p_{\text{gen}} \approx 5 \times 10^{-5}$ between nitrogen-vacancy (NV) centers in diamond [37,42]. Such a low probability of generating entanglement would require many entanglement generation attempts per heralded entangled link, resulting in demanding simulation requirements compared to deterministic generation. However, protocols can perform batches of these intrinsically probabilistic entanglement generation attempts to provide deterministic entanglement generation at prespecified times [36,37]. Such protocols can make a trade-off between entanglement generation rates and entanglement generation fidelity, for example, generating entangled links of fidelity $F \approx 0.8$ at a rate of 6 Hz or, prioritizing the number of entangled links, achieving a generation rate of 39 Hz with $F \approx 0.6$ [36]. Such a robust, deterministic entanglement generation protocol can be part of the link layer in a quantum network stack [37,43].

Suppose that we generate entanglement with a low probability p'_{gen} and that the generated entangled link has a coherence time T . Like the entanglement generation protocol discussed above, we use these entanglement generation attempts (with a low probability of success) to create a protocol that generates entanglement with a higher probability of success p_{gen} over a batch of many individual attempts. Recall that the CD protocol (Algorithm 1) discretizes time in units associated with the rate of the entanglement generation protocol (p_{gen}). That is, as we combine several entanglement generation attempts (p'_{gen}), the entanglement generation protocol (p_{gen}) takes longer and the unit of time increases in the CD protocol. Accordingly, the coherence time T of the entangled link associated with p_{gen} decreases in units of this discretized time. If we design protocols with $p_{\text{gen}} = 1, \frac{1}{2}, \frac{1}{4}$, then we observe that the optimal swap attempt probability q increases approximately linearly with p_{gen} [Fig. 5(a)]. As the discretized time step in the CD protocol depends on the time it takes the entanglement generation protocol (p_{gen}) to attempt entanglement delivery between physical neighbors, we scale the coherence time T (and cutoff time t_{cut}) inversely proportional to p_{gen} , e.g., $T(p_{\text{gen}} = 1) = 2T(p_{\text{gen}} = \frac{1}{2})$. Then, for the range of entanglement generation success probability that we investigated ($p_{\text{gen}} = 1, \frac{1}{2}, \frac{1}{4}$), we see that the optimal q for $p_{\text{gen}} = 1$ is approximately double that of the optimal q when

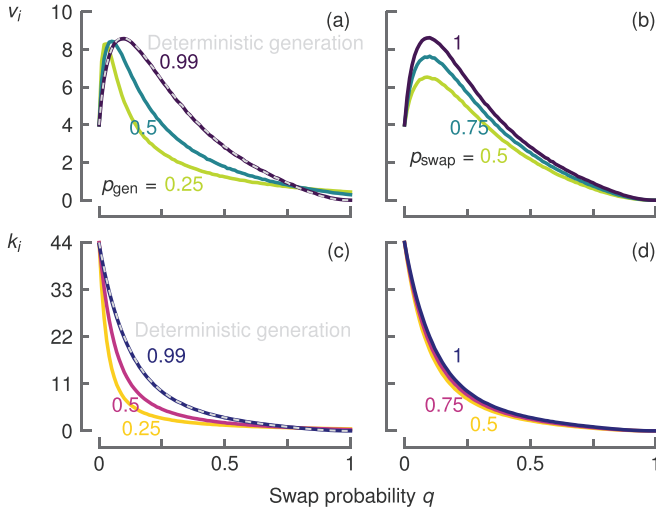


FIG. 5. Probabilistic generation of entanglement shifts the optimal probability of attempting swaps q while swaps with nonzero probability of failure decrease the maximum virtual neighborhood v_i . [(a) and (b)] Virtual neighborhood size v_i and [(c) and (d)] virtual node degree k_i as a function of q . We vary [(a) and (c)] the probability of successfully generating entanglement $p_{\text{gen}} = 0.25, 0.5, 0.99$ and [(b) and (d)] the probability of successfully executing swaps $p_{\text{swap}} = 0.5, 0.75, 1$; both colored from light to dark. For the range of p_{gen} investigated, we observe that the optimal q scales with p_{gen} (a). For example, the optimal q when $p_{\text{gen}} = \frac{1}{2}$ is approximately half of the optimal q when $p_{\text{gen}} = 1$. The maximum k_i remains approximately the same but converges to zero more quickly for lower p_{gen} (c). When swaps succeed probabilistically ($p_{\text{swap}} < 1$), nodes waste links in failed swaps. This means that nodes store fewer links, i.e., lower k_i (d). Furthermore, nodes have fewer opportunities to share links with remote nodes, resulting in a lower maximum virtual neighborhood size v_i as p_{swap} decreases (b). For probabilistic entanglement generation, we scale the coherence time T (and cutoff time t_{cut}) inversely proportion to p_{gen} ($T = 50, 50, 100, 200$ time steps, $t_{\text{cut}} = 11, 11, 22, 44$ time steps for $p_{\text{gen}} = 1, 0.99, 0.5, 0.25$). For probabilistic swaps, we consider $T = 50$ time steps ($t_{\text{cut}} = 11$ time steps). For both cases, we consider an entanglement generation fidelity $F_{\text{new}} = 0.9$ (maximum swap distance $M = 3$) and a threshold fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained using network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

$p_{\text{gen}} = \frac{1}{2}$ and that the optimal q for $p_{\text{gen}} = \frac{1}{2}$ is approximately double that of the optimal q when $p_{\text{gen}} = \frac{1}{4}$.

Last, we note that the proof that there exists a unique steady-state value for the expected number of virtual neighbors and the expected virtual degree of any node by Ref. [12] is under the assumption that entanglement generation is probabilistic ($p_{\text{gen}} < 1$). However, in this paper, we generally assume that $p_{\text{gen}} = 1$ (deterministic) for a simplified analysis of the results. Reference [12] expects that there exists a unique steady state for $p_{\text{gen}} = 1$; we elaborate on this expectation in Appendix B. Additionally, we note that the results of the performance metrics are almost indistinguishable for $p_{\text{gen}} = 0.99$ and $p_{\text{gen}} = 1$ [Figs. 5(a) and 5(c)].




FIG. 6. Entanglement swap. Physical neighbors can directly generate entangled links, while non-neighboring nodes can generate links by swapping entanglement via an intermediary node.

2. Swapping entanglement

Two nodes not connected via a physical channel may create an entangled link by swapping entanglement with an intermediary node [20]. For example, suppose that two nodes A and B do not share a physical channel but are both physical neighbors of an intermediary node I. Nodes A and B can directly herald entangled links (of the Werner form (1) with fidelities F_{Ai} and F_{Bi}) with node I over these physical channels. To implement the swap, the intermediary node performs a Bell-state measurement on its entangled qubits. Then, using local operations and classical communication, the three nodes swap entanglement, transforming the initial links into an entangled link between nodes A and B of the Werner form with fidelity [15,35]

$$F_{\text{AB}} = F_{\text{Ai}}F_{\text{Bi}} + \frac{(1 - F_{\text{Ai}})(1 - F_{\text{Bi}})}{3} \leq F_{\text{Ai}}, F_{\text{Bi}}. \quad (\text{A1})$$

Swapping entanglement via an intermediary node is illustrated in Fig. 6.

We assume that the nodes successfully swap entanglement with a probability p_{swap} and fail with a probability $1 - p_{\text{swap}}$. On failure, the initial links are destroyed and no link is produced. Experimentally, the Bell-state measurements that facilitate the entanglement swaps can be deterministic ($p_{\text{swap}} = 1$, for example, in diamond NV centers [38,42]) or probabilistic (generally $p_{\text{swap}} = \frac{1}{2}$ using linear optics [44], but $p_{\text{swap}} > \frac{1}{2}$ is possible using ancillary photons [45,46]). When entanglement swaps succeed probabilistically, nodes can lose entangled links in failed swaps. This means that nodes will store fewer links (decreased k_i) and hence will share links with a smaller set of virtual neighbors, i.e., smaller v_i [Figs. 5(b) and 5(d)]. Last, we assume that links formed in a swap assume the age of the oldest initial link (the time elapsed since the creation of the link).

3. Discarding entanglement

Qubits interact with their environment and the fidelity of the links they store decreases over time; we say that the links decohere. We model this decoherence as the successive application of a depolarizing channel, a worst-case noise model [32]. Then, during a time interval Δt , the fidelity $F(t)$ of a Werner state (1) at time t evolves according to [15]

$$F(t + \Delta t) = \frac{1}{4} + \left[F(t) - \frac{1}{4}\right]e^{-\Delta t/T}, \quad (\text{A2})$$

where T is an abstract coherence time that characterizes the exponential decay rate of the fidelity.

To ensure that the fidelity of all entangled links exceeds some threshold fidelity F_{\min} , nodes discard entangled links that are stored longer than a *cutoff time* t_{cut} [15,21,22]. Similarly, nodes discard links that have been formed in the fusion of more than M short-distance links (generated between physical neighbors)—i.e., discard links that exceed the maximum swap distance M [12]. We recall that physical neighbors generate entangled links of the Werner form (1) with a fidelity F_{new} and that new links created in an entanglement swap assume the age of the oldest initial link. Then, given an abstract coherence time T that characterizes the exponential decay rate of the fidelity, nodes must satisfy the relation [15]

$$t_{\text{cut}} \leq -T \ln \left[\frac{3}{4F_{\text{new}} - 1} \left(\frac{4F_{\min} - 1}{3} \right)^{\frac{1}{M}} \right]. \quad (2 \text{ revisited})$$

Suppose that the entangled link fidelity is not large enough for its intended purpose, e.g., distributed quantum computing applications. In that case, quantum nodes can use entanglement distillation protocols to turn low-fidelity entangled links into links of higher fidelity using local operations. The entangled links that we consider—of the Werner form (1)—are entangled for a fidelity $F > \frac{1}{2}$ ($x > \frac{1}{3}$) [47]. Then bipartite distillation protocols [48,49] can distill multiple initial entangled links of fidelity $F > \frac{1}{2}$ to a new link of fidelity $F' > F$. In this way, nodes ensuring a minimum fidelity $F_{\min} > \frac{1}{2}$ can generate higher-fidelity links if their application requires so. In general, we choose the lower bound $F_{\min} = \frac{1}{2}$ to analyze performance at the extreme of “useful” links (note that we require $F_{\min} = \frac{1}{2} + \epsilon$ for a tiny $\epsilon > 0$; however, this ϵ would have an insignificant influence on calculating Inequality (2), so we omit it for simplicity).

Although implementing such distillation is outside the scope of this work, distillation can be incorporated into our model. For example, at the entanglement generation level, we could integrate distillation in the deterministic entanglement generation protocol we discussed above. To account for the time needed to implement the distillation, the unit of discretized time increases and, consequently, T and t_{cut} decrease in terms of this discretized time. Additionally, we should adjust p_{gen} and F_{new} according to the results of the distillation protocol. For existing links, we could integrate distillation (as part of an application) in the CD protocol (Algorithm 1) (see also the CD protocol of Ref. [12] and their discussion of integrating distillation in Appendix A).

From the parameters related by Inequality (2), we adopt the cutoff time t_{cut} and the maximum swap distance M as simulation parameters. If we want to investigate various coherence times T , then we associate values of t_{cut} that satisfy Inequality (2) and use those values in our simulations. Similarly, if we want to vary the entanglement generation fidelity F_{new} , then we associate values of M that satisfy Inequality (2). In addition to the remarks in Sec. II B, we require

$$\frac{3}{4F_{\text{new}} - 1} \left(\frac{4F_{\min} - 1}{3} \right)^{\frac{1}{M}} < 1, \quad (A3)$$

since the cutoff must be positive. If we assume $F_{\min} = \frac{1}{2}$ and want some $M > 1$, then we see that there is some minimum F_{new} to satisfy Inequality (A3) regardless of the T . Lower

ALGORITHM 2. Steady-state estimation.

Inputs:

- (1) $\bar{X}_N(t)$: the sample mean of the stochastic process $\{X(t)\}$ observed over N realizations at times $t = t_0, t_1, \dots, t_{\ell-1}$;
- (2) a, b : the minimum and maximum values of the stochastic process $\{X(t)\}$;
- (3) w : size of the steady-state window.

Output:

- (1) Assessment of whether $\bar{X}_N(t)$ has attained a steady state in the time window $W = \{\ell - w, \ell - w + 1, \dots, \ell - 1\}$.

Algorithm:

- 1: Define the error $\varepsilon' \leftarrow 3(b - a)/\sqrt{N}$.
 - 2: Define the steady-state window $W \leftarrow \{\ell - w, \ell - w + 1, \dots, \ell - 1\}$.
 - 3: Calculate the size of the interval of confidence $\Delta_{ij} \leftarrow 2\varepsilon' - |\bar{X}_N(t_i) - \bar{X}_N(t_j)|, \forall i, j \in W \text{ and } i \neq j$.
 - 4: If the interval of confidence $\Delta_{ij} < \frac{3}{2}\varepsilon'$, then **abort** (steady state not found). Otherwise, declare **success** (steady state present in the time window).
-

values of M result in lower required values of F_{new} (assuming constant F_{\min}) and vice versa, motivating our choice to relate the values of F_{new} and M and those of T and t_{cut} .

Experimentally, multiqubit nodes that combine communication and memory qubits can reach coherence times in the order of seconds [50]. In the CD protocol, we discretize time in units associated with the generation rate of protocols that guarantee the generation of entangled links at specified intervals. As discussed above, such protocols currently deliver (on the order of) tens of links per second. Accordingly, we employ values of the coherence times T that may be feasible in the near future, i.e., coherence times of tens to hundreds of time steps. Using the entanglement generation fidelities we discussed above, we retrieve reasonable values for the maximum swap distance M via Inequality (2).

Last, we assume that quantum nodes have a “large-enough” number of memories. When not attempting entanglement swaps ($q = 0$), nodes can store all entangled links until they discard the links when they age to the cutoff time. This translates into nodes storing at most $d t_{\text{cut}}$ entangled links, where d is the physical degree of the node. With the moderate values of the cutoff time t_{cut} we employ in this work, the number of qubits required per node is relatively close to experimentally achieved values [40].

APPENDIX B: DATA SAMPLING

In this Appendix, we present our data sampling technique to obtain the performance metrics shown in this paper. In particular, to estimate the expected virtual neighborhood size v_i (3) and the expected virtual node degree k_i (4), we employ discrete-time network simulations that implement the CD protocol (Algorithm 1). We simulate the networks for ℓ time steps and verify (using Algorithm 2) that the performance metrics attain their steady state at the end of the simulation during a time window w , i.e., at times $t = t_{\ell-w}, t_{\ell-w+1}, \dots, t_{\ell-1}$. If the steady state is achieved, then we estimate the expected virtual neighborhood size and expected virtual node degree by

sampling the performance metrics at the final time step over many ($N = 10^4$) network simulations,

$$v_i \equiv \lim_{t \rightarrow \infty} \mathbb{E}[v_i(t)] \approx \bar{v}_{i,N}(t_{\ell-1}), \quad (\text{B1})$$

$$k_i \equiv \lim_{t \rightarrow \infty} \mathbb{E}[k_i(t)] \approx \bar{k}_{i,N}(t_{\ell-1}). \quad (\text{B2})$$

Here $\bar{v}_{i,N}(t_{\ell-1})$ and $\bar{k}_{i,N}(t_{\ell-1})$ are the virtual neighborhood size and virtual node degree of node i at time $t_{\ell-1}$ averaged over a sample with N realizations.

We now discuss the algorithm to find the steady-state expected value of a stochastic process given a sample of N realizations as introduced by Ref. [12, Appendix D, Algorithm 2]. Specifically, algorithm 2 determines whether the sample mean of the stochastic process $\{X(t)\}$ over N realizations attains a steady state, which is an adaptation of the steady-state algorithm of Ref. [12] to better to suit our needs. Specifically, we only employ steps 1–4 of the steady-state algorithm of Ref. [12] (further steps determined *when* the steady state starts). Furthermore, compared to the error ε employed by Ref. [12], we adjust the error $\varepsilon' = 3\varepsilon$ to ensure that the algorithm declares that the steady state has been reached once we are “close enough” to the steady-state value (see below for more details).

We now adjust the results found by Ref. [12] to conform to the choice of ε' in Algorithm 2. Let us consider a stochastic process $\{X(t)\}$. We assume that $\{X(t)\}$ has a constant, steady-state mean, $\lim_{t \rightarrow \infty} \mathbb{E}[X(t)] = X_\infty < \infty$, and a finite variance, $\sigma(t)^2 < \infty$. Observing the stochastic process over N realizations at times $t = \{t_0, t_1, \dots, t_{\ell-1}\}$, $t_0 < t_1 < \dots < t_{\ell-1}$, we denote the value taken in realization n as $x_n(t)$, where $a \leq x_n(t) \leq b$, with $a, b \in \mathbb{R}$. Then, we denote the sample average at time t over N realizations as

$$\bar{X}_N(t) = \frac{1}{N} \sum_{n=0}^{N-1} x_n(t). \quad (\text{B3})$$

Now let us assume that $\{X(t)\}$ has attained a steady state at some time $t = t_\alpha$. Then, considering $X(t_i)$ for all $i \geq \alpha$, we follow Ref. [12] and use the central limit theorem and the properties of a normal distribution (the probability that a normally distributed random variable takes a value more than six standard deviations from the mean value is approximately 2×10^{-9}), to conclude that

$$\Pr \left[\mathbb{E}[X(t_i)] \in \left(\bar{X}_N(t) - \frac{6\sigma(t)}{\sqrt{N}}, \bar{X}_N(t) + \frac{6\sigma(t)}{\sqrt{N}} \right) \right] \geq 1 - 2 \times 10^{-9} \approx 1. \quad (\text{B4})$$

Now, let us define an error $\varepsilon' = 3(b-a)/\sqrt{N}$ and consider a confidence interval for the steady-state sample average of $X(t_i)$,

$$\text{IC}_i = (\bar{X}_N(t_i) - \varepsilon', \bar{X}_N(t_i) + \varepsilon'). \quad (\text{B5})$$

Furthermore, we use that the standard deviation is bounded by $\sigma(t) \leq (b-a)/2$ such that $\varepsilon' = 3(b-a)/\sqrt{N} \geq 6\sigma(t)/\sqrt{N}$. Then it follows from probability (B4) that

$$\Pr[\mathbb{E}[X(t_i)] \in \text{IC}_i] \approx 1. \quad (\text{B6})$$

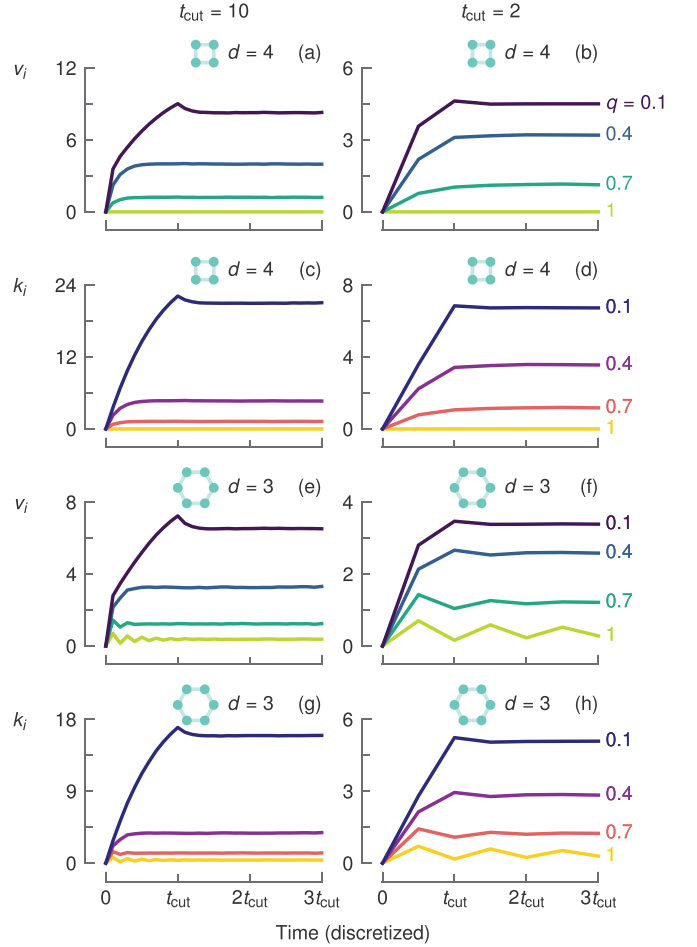


FIG. 7. Generally, the virtual neighborhood size v_i and the virtual node degree k_i quickly converge to a steady state. Performance metrics [(a), (c), (e), and (g)] v_i and [(b), (d), (f), and (h)] k_i as a function of time in infinite [(a)–(d)] square-lattice networks and [(e)–(h)] honeycomb-lattice networks. We vary [(a), (c), (e), (g)] the cutoff time $t_{\text{cut}} = 10$ time steps (coherence time $T = 45$ time steps) and [(b), (d), (f), (h)] $t_{\text{cut}} = 2$ time steps ($T = 9$ time steps) as well as the swap attempt probability $q = 1, 0.7, 0.4, 0.1$ (colored from light to dark). In networks with a square-lattice topology [(a)–(d)], the performance metrics reach their steady after approximately t_{cut} time steps. The situation is similar for infinite honeycomb-lattice networks [(e)–(h)], except that, when nodes attempt swaps often (large q), the performance metrics show periodic oscillations, i.e., the metrics do not attain a steady state. We consider an entanglement generation fidelity $F_{\text{new}} = 0.9$ (maximum swap distance $M = 3$). We assume that nodes generate entanglement and execute swaps deterministically ($p_{\text{gen}}, p_{\text{swap}} = 1$) and that nodes require a minimum link fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained using network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

Similarly, we consider the steady state $X(t_i)$ and $X(t_j)$ for all $i, j \geq \alpha$, and follow Ref. [12, Equation D.9] to find that

$$\Pr[\mathbb{E}[X(t_i)] \in \text{IC}_{ij}] \geq \frac{1 - 2 \times 10^{-9}}{2} + \frac{0.9973}{2} \approx 1 \quad (\text{B7})$$

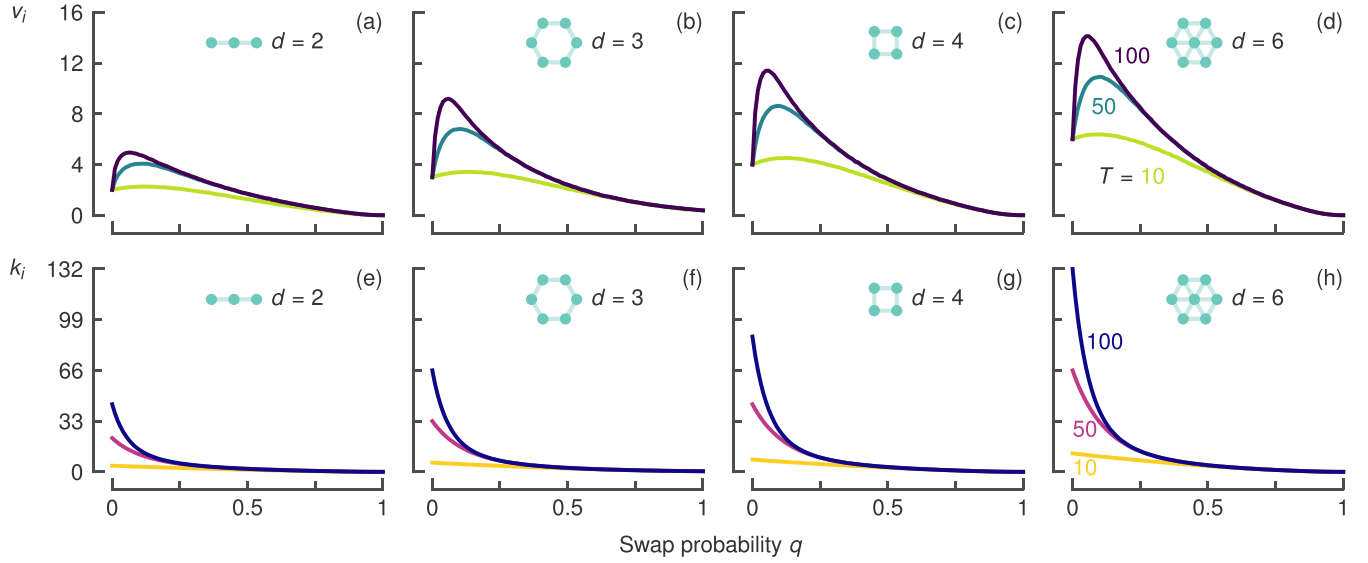


FIG. 8. The maximum virtual neighborhood size v_i and virtual node degree k_i increase with the physical node degree d . Performance metrics [(a)–(d)] v_i and [(e)–(h)] k_i as a function of the swap attempt probability q . We vary the coherence time $T = 10, 50, 100$ time steps (cutoff time $t_{\text{cut}} = 2, 11, 22$ time steps; colored from light to dark) in infinite [(a) and (e)] chains ($d = 2$), [(b) and (f)] honeycomb lattices ($d = 3$), [(c) and (g)] square lattices ($d = 4$) and [(d) and (h)] triangular lattices ($d = 6$). The bound on v_i and, consequently, v_i itself increases more than the physical node degree ratio $3/2$ when going from $d = 2$ to $d = 3$ (see Appendix B, Table II). The growth of v_i diminishes when the physical node degree increases further to $d = 4$ and $d = 6$. The number of links connected to node i , k_i , decreases monotonically from a maximum $k_i = d t_{\text{cut}}$ ($q = 0$; nodes only share entangled links with physical neighbors and generate entangled links with all their d physical neighbors) to $k_i = 0$ as $q = 1$ (nodes discard all links as nodes involve all of them in too many swaps). We consider an entanglement generation fidelity $F_{\text{new}} = 0.9$ (maximum swap distance $M = 3$). We assume that nodes generate entanglement and execute swaps deterministically ($p_{\text{gen}}, p_{\text{swap}} = 1$) and that nodes require a minimum link fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained using network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation. We omit the performance metrics for $d = 3$, $T = 10$ time steps, and $q = 0.95, 0.96, \dots, 1$ as those simulations did not attain a steady state; see the notes on the existence of a unique steady state in Appendix B.

for an interval of confidence ij

$$\text{IC}_{ij} = \left(\max(\bar{X}_N(t_i), \bar{X}_N(t_j)) - \varepsilon', \min(\bar{X}_N(t_i), \bar{X}_N(t_j)) + \varepsilon' \right). \quad (\text{B8})$$

This interval indicates the overlap in the intervals of confidence for the steady-state sample averages of $X(t_i)$ and $X(t_j)$ and has a size

$$\Delta_{ij} = 2\varepsilon' - |\bar{X}_N(t_i) - \bar{X}_N(t_j)|. \quad (\text{B9})$$

We used that $\varepsilon' = 3(b - a)/\sqrt{N} \geq 6\sigma(t)/\sqrt{N}$ and that the probability that a normally distributed random variable takes a value more than six respectively three standard deviations from the mean is approximately 2×10^{-9} and 0.0027.

For infinite quantum networks with a regular topology with a physical node degree d , the virtual neighborhood size v_i and the virtual node degree k_i are bounded by a function of the cutoff time t_{cut} and the maximum swap distance M (Table II) [51, see Appendix B for details]. We use these values as the upper bound b in calculating the error ε' . Note that, in finite networks, the boundary nodes have fewer than d physical neighbors such that also the upper bound b is smaller. We also note that the performance metrics are bounded below by $a = 0$.

We execute each network simulation for $3t_{\text{cut}}$ time steps and use a steady-state window $w = t_{\text{cut}}$ (except for the finite

chain networks, which we simulate for $6t_{\text{cut}}$ time steps). Visual inspection (see Fig. 7) showed that the performance metrics converged quickly to their steady states, with Algorithm 2 confirming that the performance metrics reached the steady state.

As noted previously by Ref. [12], the overlaps between intervals of confidence Δ_{ij} may be too small in specific scenarios, meaning that Algorithm 2 aborts, even when a closer visual inspection strongly indicates that the performance metrics attain some form of a steady state. For instance, the error is relatively small when the upper bound b is relatively low, e.g., when the cutoff time t_{cut} is short, or the maximum swap distance M is low. Then, using the original error $\varepsilon = (b - a)/\sqrt{N}$, the performance metrics [sometimes $v_i(t)$, other

TABLE II. Bounds on the virtual neighborhood size v_i and the virtual node degree k_i in infinite regular networks.

	v_i	k_i
$d = 2$	$2 \min(t_{\text{cut}}, M)$	$2t_{\text{cut}}$
$d = 3$	$3 \min(t_{\text{cut}}, \frac{1}{2}M(M+1))$	$3t_{\text{cut}}$
$d = 4$	$4 \min(t_{\text{cut}}, \frac{1}{2}M(M+1))$	$4t_{\text{cut}}$
$d = 6$	$6 \min(t_{\text{cut}}, \frac{1}{2}M(M+1))$	$6t_{\text{cut}}$

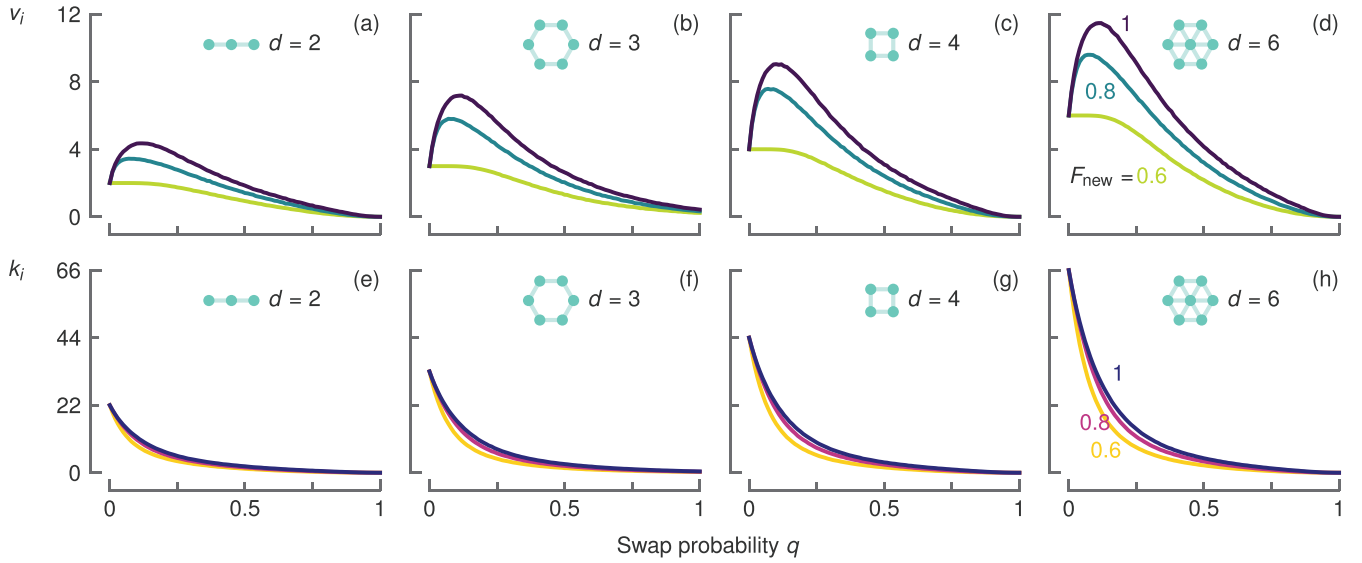


FIG. 9. The maximum virtual neighborhood size v_i and virtual node degree k_i increase with the physical node degree d . Performance metrics [(a)–(d)] v_i and [(e)–(h)] k_i as a function of the swap attempt probability q . We vary the entanglement generation fidelity $F_{\text{new}} = 0.6, 0.8, 1$ (maximum swap distance $M = 1, 2, 4$; colored from light to dark) in infinite [(a) and (e)] chains ($d = 2$), [(b) and (f)] honeycomb lattices ($d = 3$), [(c) and (g)] square lattices ($d = 4$), and [(d) and (h)] triangular lattices ($d = 6$). We consider a coherence time $T = 50$ time steps (cutoff time $t_{\text{cut}} = 11$ time steps). We assume that nodes generate entanglement and execute swaps deterministically ($p_{\text{gen}}, p_{\text{swap}} = 1$) and that nodes require a minimum link fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained using network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

times $k_i(t)$] did not attain a steady state for some values of q according to Algorithm 2. However, the resulting “unsteady-state” values were strongly in line with expectations compared to the steady-state values of nearby swap probabilities q . To prevent Algorithm 2 from aborting in such a situation, we redefine the error $\varepsilon' = 3(b - a)/\sqrt{N} = 3\varepsilon$ (which is similar to increasing the value of b as proposed by Ref. [12]).

We measure the error in the estimate of the expected steady-state values using the standard error $s_{\bar{X}} = s/\sqrt{N}$, where s is the sample (N realizations) standard deviation. The plots in this work show the data as $\bar{X} \pm 6s_{\bar{X}}$, providing a $1 - 2 \times 10^{-9} \approx 100\%$ interval of confidence. Even with such a large confidence interval, most error bands are on the order of or smaller than the plot line width.

1. Notes on the existence of a unique steady state

Reference [12] showed that there is a unique steady-state value for the expected number of virtual neighbors $v_i \equiv \lim_{t \rightarrow \infty} \mathbb{E}[v_i(t)]$ and expected virtual degree of any node $k_i \equiv \lim_{t \rightarrow \infty} \mathbb{E}[k_i(t)]$ when a quantum network is running the CD protocol of Algorithm 1. This proof is under the assumption that entanglement generation is probabilistic, $p_{\text{gen}} < 1$. However, we generally assume $p_{\text{gen}} = 1$ to simplify performance analysis. We now elaborate on the assumption that the unique steady state also exists when $p_{\text{gen}} = 1$.

The proof of Ref. [12] uses that the steady state is unique for an aperiodic, irreducible, positive recurrent Markov chain [52, Theorem 9.3.6]. In particular, they use the ages of all entangled links in the network to represent the state of the network s and the set of all possible states \mathcal{S} (both finite). Then

they show that the transition of a state $s(t)$, $t \in \mathcal{N}$ (discrete time steps) does not depend on past information,

$$\begin{aligned} \Pr[s(t+1) = \sigma \mid s(0), s(1), \dots, s(t)] \\ = \Pr[s(t+1) = \sigma \mid s(t)]. \end{aligned} \quad (\text{B10})$$

Hence, the state of the network can be modeled as a Markov chain. Then they show that this Markov chain is aperiodic, irreducible, and positive recurrent using that, for $p_{\text{gen}} < 1$, there is a nonzero probability of returning to the initial state (no links). They conclude that the limit $\lim_{t \rightarrow \infty} \Pr[s(t) = \sigma]$, $\forall \sigma \in \mathcal{S}$ is unique and exists. Last, they express v_i and k_i as a function of this limit to conclude that v_i and k_i also exist and are unique.

From their simulations, Ref. [12] also expects a unique steady state for $p_{\text{gen}} = 1$. However, they note that “the main difficulty in proving its existence is that the Markov chain is not always irreducible (the state with no links may not be reachable from some other states since links are generated at the maximum rate).” We observe almost indistinguishable steady-state behavior when $p_{\text{gen}} = 0.99$ compared to $p_{\text{gen}} = 1$ [Fig. 5(a)]. Additionally, we observe that the performance metrics appear to quickly converge to a steady state (Fig. 7).

We note that for regular networks and a nonzero probability of nodes attempting entanglement swaps ($q > 0$), there is a nonzero probability that all links are involved in too many swaps and discarded. That is, the state of the network returns to the initial state of no links (for $q = 0$ and $p_{\text{gen}} = 1$, the system deterministically reaches $v_i = d$, $k_i = d t_{\text{cut}}$). For example, starting from the initial (no links) state, regular networks with an even number of physical neighbors can pair

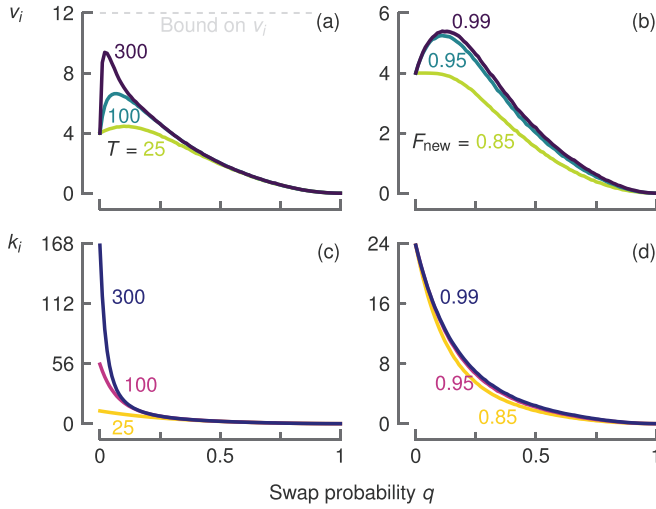


FIG. 10. Independent of the choice of network parameters, the optimal probability of attempting to swap q depends on the coherence time T and entanglement generation fidelity F_{new} . [(a) and (b)] Virtual neighborhood size v_i and [(c) and (d)] virtual node degree k_i of a node in an infinite square lattice ($d = 4$) as a function of q . We vary [(a) and (c)] $T = 25, 100, 300$ time steps (cutoff time $t_{\text{cut}} = 3, 14, 42$ time steps) and [(b) and (d)] $F_{\text{new}} = 0.85, 0.95, 0.99$ (maximum swap distance $M = 1, 2, 3$); both colored from light to dark. Similarly to Fig. 3 but for different network parameters, the optimal q decreases with longer T and increases for higher F_{new} , and k_i decreases monotonically as q increases. Note that for long T and decreasing q , the maximum v_i seems to approach the bound on v_i (a). For a relatively short $T = 75$ ($t_{\text{cut}} = 6$), the cutoff time quickly limits v_i for increasing F_{new} (b). We consider [(a) and (c)] $F_{\text{new}} = 0.99$ ($M = 2$) and [(b) and (d)] $T = 75$ time steps ($t_{\text{cut}} = 6$ time steps). We assume that nodes generate entanglement and execute swaps probabilistically ($p_{\text{gen}} = 0.99$, $p_{\text{swap}} = 0.5$) and that nodes require a minimum link fidelity $F_{\text{min}} = 0.8$. Results obtained using network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

the even number of generated links ($p_{\text{gen}} = 1$) to involve each link in too many swaps (resulting in $v_i, k_i = 0$ when $q \neq 0$). However, in a network with a honeycomb topology ($d = 3$), nodes generate three links in each time step, meaning that it is more challenging to match all the links in a way where all links are involved in too many swaps. This difficulty in matching links results in nonzero performance metrics in a honeycomb lattice as $q \rightarrow 1$ (Figs. 8 and 9; the performance metrics converge to zero for the other topologies).

On closer inspection (Fig. 7), we see that the performance metrics show periodic oscillations as $q \rightarrow 1$ in a honeycomb lattice network ($d = 3$). To illustrate how this happens, assume that nodes will always attempt swaps ($q = 1$) and that the performance metrics start without links. Then, nodes generate three links at the first time step and generally swap two links. This makes it challenging to involve each link in too many swaps (M), not removing all links at this first time step. Then, in the next time step, nodes again generate three links, meaning some nodes now have four links that nodes can all swap. At the end of this step, nodes have involved more links

in too many swaps, thus discarding them, resulting in lower performance metrics. This oscillatory behavior diminishes after some time. For longer cutoff times t_{cut} and simulation times (recall that we simulate the networks for $3t_{\text{cut}}$ time steps), there is “enough” time for this periodicity to vanish [Figs. 7(e) and 7(g)] and for Algorithm 2 to declare a steady state has been attained. However, for short t_{cut} , the periodicity is still strong after $3t_{\text{cut}}$ time steps and Algorithm 2 declares there is no steady state [as is the case for $q = 0.95, 0.96, \dots, 1$ in Figs. 8(b) and 8(f) with $t_{\text{cut}} = 2$; we omit those values in the plot for $d = 3$]. Additionally, the oscillatory behavior persists longer when t_{cut} is small, diminishing for (significantly) longer simulation times.

Last, we note that the proof by Ref. [12] assumes that the state space of the network is finite. However, Theorem 9.3.6 [52] also applies to Markov chains with an infinite state space, as is the case for infinite networks. The difficulty becomes showing that the chain is positive recurrent, i.e., for a state to have a finite mean return time (for finite-state space chains, it is sufficient to be irreducible in order to be positive recurrent [52, Theorem 9.3.5]).

APPENDIX C: EXTENDED NETWORK SIMULATIONS

In this Appendix, we provide additional data to the results presented in Sec. III. In particular, we present both the virtual neighborhood size v_i and virtual node degree k_i for all the infinite regular topologies ($d = 2, 3, 4, 6$) for varying coherence times T (Fig. 8) and for varying entanglement generation fidelity F_{new} (Fig. 9). Additionally, we present the performance metrics in an infinite square-lattice network for varying T and F_{new} with different network parameters (Fig. 10). Last, we present k_i in addition to v_i for finite chains and finite square lattices (Fig. 11).

For increasing physical node degrees d , the maximum virtual neighborhood size v_i also increases (Figs. 8 and 9). For example, increasing $d = 2$ to $d = 3$ increases the maximum value of v_i by more than the ratio of node degrees ($3/2$). We note that, for increasing maximum swap distance M , the bound on v_i grows quicker in networks with $d = 3$ than those with $d = 2$ (see Appendix B, Table II). For example, increasing $M = 1 \rightarrow M = 2$ (and assuming sufficiently large t_{cut} such that v_i is not bounded by the cutoff time), the bound on v_i increases from $2 \rightarrow 4$ when $d = 2$, and from $3 \rightarrow 9$ when $d = 3$. Increasing the physical node degree to $d = 4$ and $d = 6$ still increases the maximum v_i , but the increase relative to the ratio of physical node degrees (compared to $d = 2$) diminishes.

The performance metrics show the same qualitative behavior for different combinations of network parameters. For example, using probabilistic entanglement generation and execution of swaps ($p_{\text{gen}} = 0.99$, $p_{\text{swap}} = 0.5$; in contrast to deterministic generation and execution we previously assumed) and a higher minimum required fidelity ($F_{\text{min}} = 0.8$), (i) the virtual neighborhood size v_i still increases with longer coherence time T and better entanglement generation fidelity F_{new} , (ii) the optimal swap probability q —maximizing v_i —decreases for longer T and increases for higher F_{new} , and (iii) the virtual node degree k_i decreases monotonically for

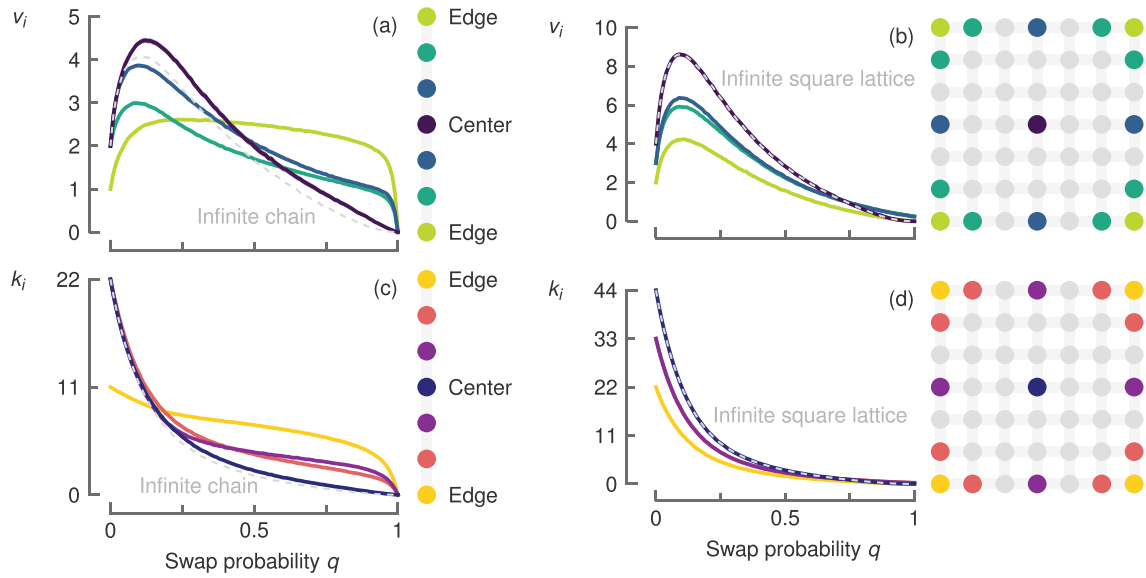


FIG. 11. Similarly to the virtual neighborhood size v_i , the behavior of the virtual node degree k_i as a function of the swap attempt probability q strongly depends on a node's location in a finite chain but, in a finite square lattice, it behaves qualitatively the same for all nodes. Performance metrics [(a) and (b)] v_i and [(c) and (d)] k_i as a function of q for nodes near and far from the network boundary (colored from light to dark) in a finite [(a) and (c)] chain ($d = 2$) and [(b) and (d)] square lattice ($d = 4$). We consider a coherence time $T = 50$ time steps (cutoff time $t_{\text{cut}} = 11$ time steps) and an entanglement generation fidelity $F_{\text{new}} = 0.9$ (maximum swap distance $M = 3$). We assume that nodes generate entanglement and execute swaps deterministically ($p_{\text{gen}}, p_{\text{swap}} = 1$) and that nodes require a minimum link fidelity $F_{\text{min}} = \frac{1}{2}$. Results obtained with network simulations and Monte Carlo sampling with $N = 10^4$ realizations per sample, presented with an error band of $\pm 6s/\sqrt{N}$ (generally smaller than the line width), where s is the sample standard deviation.

increasing q , going from $k_i = d t_{\text{cut}}$ ($q = 0$) to $k_i = 0$ ($q = 1$) (Fig. 10; for more details on the influence of probabilistic entanglement generation and swap execution on the behavior of the performance metrics, see Appendix A, Fig. 5). Additionally, we note that for long T (and associated long t_{cut}), the maximum v_i seems to approach the bound on v_i as nodes attempt very few swaps (low q). For a maximum swap distance $M = 2$, v_i is bounded by $4 \min(t_{\text{cut}}, \frac{1}{2}M(M+1)) = 4 \times \frac{1}{2} \times 2 \times 3 = 12$ (Appendix B, Table II). Furthermore, when T , t_{cut} are relatively short, increasing F_{new} results in a limited increase in the maximum v_i .

Similarly to the virtual neighborhood size v_i , the effect of network boundaries on the virtual node degree k_i depends on the network topology (Fig. 11). For the same reasons as explained in the main text for v_i , the behavior of k_i depends on the node's distance to the edge of a chain. In contrast, the behavior of k_i is qualitatively similar for all nodes in a finite square-lattice network—decreasing monotonically from $k_i = d t_{\text{cut}}$ ($q = 0$) to $k_i = 0$ ($q = 1$).

Last, for more data, we refer to the Jupyter Notebooks in this project's GitHub repository [41] and to the thesis [51] that uses the same CD protocol.

-
- [1] H. J. Kimble, The quantum internet, *Nature (London)* **453**, 1023 (2008).
 - [2] S. Wehner, D. Elkouss, and R. Hanson, Quantum internet: A vision for the road ahead, *Science* **362**, eaam9288 (2018).
 - [3] W. Kozłowski, S. Wehner, R. V. Meter, B. Rijsman, A. S. Cacciapuoti, M. Caleffi, and S. Nagayama, *Architectural Principles for a Quantum Internet*, RFC 9340 (RFC Editor, 2023).
 - [4] A. K. Ekert, Quantum cryptography based on Bell's theorem, *Phys. Rev. Lett.* **67**, 661 (1991).
 - [5] C. H. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, *Theor. Comput. Sci.* **560**, 7 (2014).
 - [6] A. Broadbent, J. Fitzsimons, and E. Kashefi, Universal blind quantum computation, in *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science* (IEEE, Los Alamitos, CA, 2009), pp. 517–526.
 - [7] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, Distributed quantum computation based on small quantum registers, *Phys. Rev. A* **76**, 062323 (2007).
 - [8] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links, *Phys. Rev. X* **4**, 041041 (2014).
 - [9] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, *Nature (London)* **402**, 390 (1999).
 - [10] J. I. Cirac, A. K. Ekert, S. F. Huelga, and C. Macchiavello, Distributed quantum computation over noisy channels, *Phys. Rev. A* **59**, 4249 (1999).
 - [11] J. Eisert, K. Jacobs, P. Papadopoulos, and M. B. Plenio, Optimal local implementation of nonlocal quantum gates, *Phys. Rev. A* **62**, 052317 (2000).

- [12] A. G. Iñesta and S. Wehner, Performance metrics for the continuous distribution of entanglement in multiuser quantum networks, *Phys. Rev. A* **108**, 052615 (2023).
- [13] K. Chakraborty, F. Rozpedek, A. Dahlberg, and S. Wehner, Distributed routing in a quantum internet, [arXiv:1907.11630](https://arxiv.org/abs/1907.11630).
- [14] G. Vardoyan, S. Guha, P. Nain, and D. Towsley, On the stochastic analysis of a quantum entanglement switch, *ACM SIGMETRICS Perf. Eval. Rev.* **47**, 27 (2019).
- [15] Á. G. Iñesta, G. Vardoyan, L. Scavuzzo, and S. Wehner, Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs, *npj Quantum Inf.* **9**, 46 (2023).
- [16] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, Routing entanglement in the quantum internet, *npj Quantum Inf.* **5**, 25 (2019).
- [17] W. Tittel, J. Brendel, H. Zbinden, and N. Gisin, Violation of Bell inequalities by photons more than 10 km apart, *Phys. Rev. Lett.* **81**, 3563 (1998).
- [18] C. Harney and S. Pirandola, Analytical methods for high-rate global quantum networks, *PRX Quantum* **3**, 010349 (2022).
- [19] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress *et al.*, Heralded entanglement between solid-state qubits separated by three metres, *Nature (London)* **497**, 86 (2013).
- [20] M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, “Event-ready-detectors” Bell experiment via entanglement swapping, *Phys. Rev. Lett.* **71**, 4287 (1993).
- [21] O. A. Collins, S. D. Jenkins, A. Kuzmich, and T. A. B. Kennedy, Multiplexed memory-insensitive quantum repeaters, *Phys. Rev. Lett.* **98**, 060502 (2007).
- [22] F. Rozpedek, K. Goodenough, J. Ribeiro, N. Kalb, V. C. Vivoli, A. Reiserer, R. Hanson, S. Wehner, and D. Elkouss, Parameter regimes for a single sequential quantum repeater, *Quantum Sci. Technol.* **3**, 034002 (2018).
- [23] K. Azuma, S. Bäuml, T. Coopmans, D. Elkouss, and B. Li, Tools for quantum network design, *AVS Quantum Sci.* **3**, 014101 (2021).
- [24] S. Khatri, C. T. Matyas, A. U. Siddiqui, and J. P. Dowling, Practical figures of merit and thresholds for entanglement distribution in quantum networks, *Phys. Rev. Res.* **1**, 023032 (2019).
- [25] E. Shchukin, F. Schmidt, and P. van Loock, Waiting time in quantum repeaters with probabilistic entanglement swapping, *Phys. Rev. A* **100**, 032322 (2019).
- [26] E. Shchukin and P. van Loock, Optimal entanglement swapping in quantum repeaters, *Phys. Rev. Lett.* **128**, 150502 (2022).
- [27] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. Vermeulen, R. N. Schouten, C. Abellán *et al.*, Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres, *Nature (London)* **526**, 682 (2015).
- [28] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Entanglement of single-atom quantum bits at a distance, *Nature (London)* **449**, 68 (2007).
- [29] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, High-rate, high-fidelity entanglement of qubits across an elementary quantum network, *Phys. Rev. Lett.* **124**, 110501 (2020).
- [30] R. J. Hughes, J. E. Nordholt, D. Derkacs, and C. G. Peterson, Practical free-space quantum key distribution over 10 km in daylight and at night, *New J. Phys.* **4**, 43 (2002).
- [31] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai *et al.*, Satellite-based entanglement distribution over 1200 kilometers, *Science* **356**, 1140 (2017).
- [32] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2010).
- [33] R. F. Werner, Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model, *Phys. Rev. A* **40**, 4277 (1989).
- [34] R. Jozsa, Fidelity for mixed quantum states, *J. Mod. Opt.* **41**, 2315 (1994).
- [35] W. J. Munro, K. Azuma, K. Tamaki, and K. Nemoto, Inside quantum repeaters, *IEEE J. Sel. Top. Quantum Electron.* **21**, 78 (2015).
- [36] P. C. Humphreys, N. Kalb, J. P. Morits, R. N. Schouten, R. F. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, Deterministic delivery of remote entanglement on a quantum network, *Nature (London)* **558**, 268 (2018).
- [37] M. Pompili, C. Delle Donne, I. te Raa, B. van der Vecht, M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk, S. L. Hermans, P. Pawełczak *et al.*, Experimental demonstration of entanglement delivery using a quantum network stack, *npj Quantum Inf.* **8**, 121 (2022).
- [38] W. Pfaff, B. J. Hensen, H. Bernien, S. B. van Dam, M. S. Blok, T. H. Taminiau, M. J. Tiggelman, R. N. Schouten, M. Markham, D. J. Twitchen *et al.*, Unconditional quantum teleportation between distant solid-state quantum bits, *Science* **345**, 532 (2014).
- [39] M. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. Itano, J. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri *et al.*, Deterministic quantum teleportation of atomic qubits, *Nature (London)* **429**, 737 (2004).
- [40] C. E. Bradley, J. Randall, M. H. Abobeih, R. C. Berrevoets, M. J. Degen, M. A. Bakker, M. Markham, D. J. Twitchen, and T. H. Taminiau, A ten-qubit solid-state spin register with quantum memory up to one minute, *Phys. Rev. X* **9**, 031045 (2019).
- [41] L. Talsma, A. G. Iñesta, and S. Wehner, Data underlying the publication: Continuously distributing entanglement in quantum networks with regular topologies (2024) Github repository, <https://github.com/lars-talsma/CD-regular-networks/>.
- [42] M. Pompili, S. L. Hermans, S. Baier, H. K. Beukers, P. C. Humphreys, R. N. Schouten, R. F. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse *et al.*, Realization of a multi-node quantum network of remote solid-state qubits, *Science* **372**, 259 (2021).
- [43] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpedek, M. Pompili, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, R. Hanson, and S. Wehner, A link layer protocol for quantum networks, in *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM’19)* (Association for Computing Machinery, New York, 2019), pp. 159–173.
- [44] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, Long-distance quantum communication with atomic ensembles and linear optics, *Nature (London)* **414**, 413 (2001).

- [45] F. Ewert and P. van Loock, 3/4-efficient Bell measurement with passive linear optics and unentangled ancillae, *Phys. Rev. Lett.* **113**, 140403 (2014).
- [46] M. J. Bayerbach, S. E. D'Aurelio, P. van Loock, and S. Barz, Bell-state measurement exceeding 50% success probability with linear optics, *Sci. Adv.* **9**, eadf4080 (2023).
- [47] A. Peres, Separability criterion for density matrices, *Phys. Rev. Lett.* **77**, 1413 (1996).
- [48] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Purification of noisy entanglement and faithful teleportation via noisy channels, *Phys. Rev. Lett.* **76**, 722 (1996).
- [49] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, Quantum privacy amplification and the security of quantum cryptography over noisy channels, *Phys. Rev. Lett.* **77**, 2818 (1996).
- [50] C. Bradley, S. de Bone, P. Möller, S. Baier, M. Degen, S. Loenen, H. Bartling, M. Markham, D. Twitchen, R. Hanson *et al.*, Robust quantum-network memory based on spin qubits in isotopically engineered diamond, *npj Quantum Inf.* **8**, 122 (2022).
- [51] L. Y. Talsma, Continuous distribution of entanglement in quantum networks with regular topologies, MSc thesis, Delft University of Technology, 2023.
- [52] P. Van Mieghem, *Performance Analysis of Complex Networks and Systems* (Cambridge University Press, Cambridge, UK, 2014).