

Optimized Generation of Entanglement by Real-Time Ordering of Swapping Operations

Ranjani G. Sundaram
Stony Brook University, NY

Himanshu Gupta
Stony Brook University, NY

Abstract—Long-distance quantum communication in quantum networks faces significant challenges due to the constraints imposed by the no-cloning theorem. Most existing quantum communication protocols rely on the a priori distribution of entanglement pairs (EPs), a process known to incur considerable latency due to its stochastic nature. In this work, we consider the problem of minimizing the latency of establishing an EP across a pair of nodes in a quantum network. While prior research has primarily focused on minimizing the expected generation latency by selecting *static* entanglement routes and/or swapping trees in advance, our approach considers a real-time adaptive strategy—wherein the order of entanglement-swapping operations (hence, the swapping tree used) is progressively determined at runtime based on the runtime success/failure of the stochastic events. In this context, we present a greedy algorithm that iteratively determines the best route and/or entanglement-swapping operation to perform at each stage based on the current network. We evaluate our schemes on randomly generated networks and observe a reduction in latency of up to 40% from the optimal offline approach.

I. Introduction

Quantum networks enable the construction of large, robust, and more capable quantum computing platforms by connecting smaller quantum computers. Quantum networks [17] also enable various important applications [9, 13]. Quantum network communication is challenging, as the physical transmission of quantum states across nodes can incur irreparable errors due to quantum no-cloning [7]. Thus, quantum protocols such as teleportation [1] or telestates [6] are used to facilitate quantum communication or computation across nodes in a quantum network; these protocols, however, require an a priori entangled pair (EP) shared across the end nodes. Our work focuses on efficiently generating EPs in a quantum network.

Generating EPs over long distances in a quantum network is challenging. In particular, coordinated entanglement swapping (e.g., DLCZ protocol [8]) using quantum repeaters can be used to establish long-distance entanglements from short-distance entanglements. However, due to the low probability of success of the underlying physical processes, EP generation can incur significant latency [15]; high generation latency of EPs can lead to high circuit execution times and even defeat their purpose due to decoherence. Thus, this paper considers the problem of generating EPs across remote nodes in a quantum network with minimal latency.

Adaptive Generation of EPs. Prior work on efficient EP generation [3, 4, 10, 11, 14, 16] almost exclusively consists of *offline* approaches that select *static* entanglement routes or swapping trees to minimize *expected* generation latency. These structures are constructed based on expected latencies

and, thus, may incur high latency at runtime depending on the actual outcomes of the underlying stochastic events, as illustrated later (Fig. 2). In contrast, a runtime-adaptive approach that selects the entanglement route and/or the order of entanglement-swapping operations based on the outcomes of the stochastic events *during the generation process* can potentially minimize the runtime generation latency. Moreover, such adaptive approaches can also be very effective in incorporating decoherence constraints due to the runtime availability of actual ages of the qubits. Based on the above insights, in this paper, we develop effective adaptive generation approaches to generate EPs with minimal latency in quantum networks. To the best of our knowledge, the only prior work that develops an adaptive EP generation approach is [12]—but they assume a *given* entanglement path, and the time and space complexity of their approach is exponential in the network size and thus feasibly only for very small (at most 5-10 nodes) networks and low decoherence times (see §VI).

Our Contributions. In the above context, we make the following contributions.

- We first consider a simplified version of the problem where entanglement-swapping operations can be performed along exactly one selected (entanglement) path in the network. In this setting, we design an adaptive algorithm that, at each stage, selects the best entanglement-swapping operation to be performed based on the current network state (available EPs and their ages).
- We also consider a generalized setting wherein the designed adaptive approach selects, at each stage, the best route and the entanglement-swapping operation to be performed based on the current network state.
- We conduct extensive evaluations of our algorithms on a wide array of quantum networks and observe a reduction in latency of up to 40% compared to prior static approaches.

II. Background

This section presents the relevant background related to the generation of EPs over remote nodes in a quantum network.

Quantum Network. We consider a quantum network as a graph over the set of quantum computers as network nodes. Each edge in the graph represents a (quantum and classical) communication link. We refer to these network edges as *network links* or just *links*. Our network model is similar to the one used in some of the recent works [10, 11] on efficient generation of EPs and graph states. In particular, each node has an atom-photon EP generator with latency




Fig. 1: Establish EPs over remote nodes.

(t_g) and probability of success (p_g). A node’s atom-photon generation capacity/rate is its aggregate capacity and may be split across its incident links. Each network link generates EPs (called *link-EPs*) shared over the adjacent nodes using the nodes’ atom-photon generators and an optical-BSM (Bell state measurement) device located in the middle of the link. The optical-BSM has a certain probability of success (p_{ob}), and each half-link (from either end-node) to the device has a probability of transmission success (p_e). We implicitly assume the atom-photon generation latency to include the times for photon transmission, optical-BSM latency, and classical acknowledgment; thus, the the generation latency of a link-EP is $t_g/(p_g^2 p_e^2 p_{ob})$ between two nodes. To facilitate atom-atom ES operations, each network node is also equipped with an atomic-BSM device with appropriate operation latency t_b and probability of success p_b .

Generating EPs over Remote Nodes. To distribute an EP over remote nodes, EPs are generated over larger and larger distances from shorter-distance EPs using a sequence of “entanglement-swapping” operations. An entanglement swapping (ES) operation can be looked up as being performed over three nodes (A, B, C) with two EPs over (A, B) and (B, C); the ES operation results in an EP over the nodes (A, C), by essentially teleporting the first qubit in node B (i.e., the qubit of the EP over (A, B)) to the node C using the second EP over (B, C). In general, an EP over a pair of remote nodes A and B can be generated by a sequence of ES operations over the EPs over adjacent nodes along a path from A to B .

Swapping Trees. In recent work, [11] introduces the concept of *swapping trees*, which are a characterization of the different orderings of ES operations that can be used to establish an EP over the endpoints of a network path. Given a path P with endpoints s and d , any complete binary tree over the ordered links in P gives a way to generate an EP over (s, d). Each vertex in the tree corresponds to a pair of network nodes in P , with each leaf representing a link (See Fig 1). The latency of a swapping tree is the *expected* time taken to establish the EP corresponding to the tree’s root node.

Decoherence Threshold Time τ . The fidelity of the quantum state decreases over time due to interaction with the environment; this is called decoherence. In order to limit the decoherence of qubits in an EP, we enforce an upper bound on the age (time spent in a quantum memory) of the qubits

involved. A qubit is said to satisfy the decoherence constraint if its age is below a decoherence threshold of τ .

III. Problem Formulation and Related Work

We start by formulating the real-time¹ entanglement generation problem addressed in this paper. We now consider entanglement generation over a single path in the network; we will consider the more general multi-path problem later in §V.

SP-SWO Problem (Single-Path Swap Ordering). Given a quantum network and a source-destination pair $\{(s, d)\}$, the SP-SWO problem is to determine a path P from s to d and a real-time (during EP generation) ordering of the entanglement-swapping (ES) operations so as to minimize the actual generation latency of the *target* EP (s, d) under the below node and decoherence constraints.

- 1) *Node Constraints.* For each node, the aggregate resources used (in generating the link-EPs) are less than the available resources.
- 2) *Decoherence Constraints.* The age of any qubit is less than the decoherence threshold τ .

Related Work In recent years, a significant amount of work has been done on efficiently generating remote EPs in a quantum network. Below, we discuss these works categorized as follows.

Selection of Entanglement Routes. Most of the work on efficient generation of EPs in quantum networks focuses on selecting one or more entanglement paths/routes [3, 4, 14, 16]. These works typically assume that the underlying processes, including link-EP generation and atomic BSMs, are synchronized; if all of these processes along the chosen entanglement route succeed, the end-to-end EP is generated. On failure of any of the underlying processes (atomic BSM or link-EP generation), all the generated link/subpath-EPs are discarded, and the whole process starts again from scratch. Such approaches require no quantum memories or require only memories with minimal storage time. In particular, [16] designs a Dijkstra-like algorithm to construct an optimal path between a pair of nodes when there are multiple channels between adjacent nodes; they generalize their techniques to also select multiple paths over multiple pairs of nodes. In addition, [3] designs a linear programming (LP) formulation based on multi-commodity flow to select routing paths for a given set of source-destination pairs; they map the operation-based fidelity constraint to the path length and use copies of nodes to model the fidelity constraint in the LP. Among earlier works, [14] proposes a greedy solution for grid networks, and [5] proposes virtual-path-based routing in ring/grid networks.

Constructing Swapping Trees. In networks with memories at nodes, a qubit of an EP may wait (in a quantum memory) for its counterpart to become available so that an ES operation can be performed. In such approaches, as mentioned in the previous section, the goal is to determine a swapping tree that essentially represents the *order* in which ES operations

¹We use real-time and runtime interchangeably, in this paper.

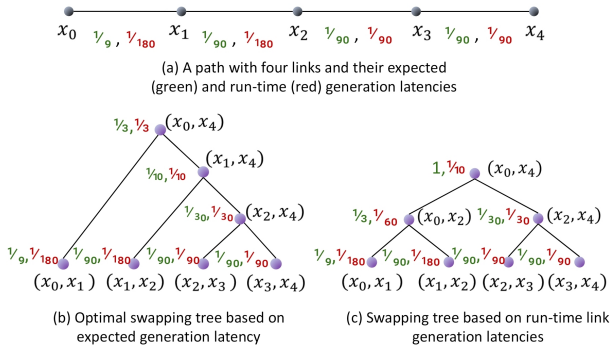


Fig. 2: (a). Path (x_1, \dots, x_4) with expected (green) and actual (at runtime, in red) generation latencies of the link-EPs. (b) Optimal swapping tree that minimizes the expected generation latency of the EP over (x_1, x_4) . Here, the green values are the expected generation latencies of the EP at each vertex, and the red values are the actual generation latencies. (c) Alternate swapping tree with sub-optimal expected generation latency but with lower runtime generation latency compared to the swapping tree in (b).

must be performed along the selected entanglement path. In particular, [2] formulated the entanglement generation rate on a given path between two nodes and then used the formulated path metric to select the optimal path exhaustively; their path metric is based on balanced swapping trees. In more recent work, [11] develops a dynamic programming approach to determine the swapping tree with minimum expected generation latency, under node and decoherence constraints, for generating an EP over a given source-destination pair.

Real-Time Determination of Swapping Order. The above works seek to minimize the expected latency of generating the target EP by selecting one or more *static* swapping trees up front. However, as discussed in the previous section, an optimal entanglement route or swapping tree may incur high generation latency at runtime depending on the outcome of the link-EP generation attempts and/or ES operations. Thus, an adaptive strategy that determines the (order of) ES operations in runtime can outperform fixed swapping trees (see Fig. 2). With the above insight, a recent work [12] considers the problem of minimizing the runtime latency of generating an EP over a pair of nodes using a *given* path. They formulate the problem as a Markov Decision Process wherein each state corresponds to each possible network state (available EPs with ages), and develop an optimal state transition policy. However, their algorithm takes exponential time and space—and is feasible only for very small networks (at most 5-10 nodes) and low decoherence times (see §VI). Our work is inspired by the above work.

IV. GREEDY-SP Algorithm

We describe our proposed GREEDY-SP algorithm for the SP-SWO problem. We start by describing a high-level idea of the algorithm.

Basic Idea. The basic idea behind GREEDY-SP is to first select a path of minimum expected generation latency over the (s, d) pair, and, then, at each event (as defined below), pick the ES X that minimizes the expected latency (after performing X)

of generating an EP across (s, d) . More elaborately, we first use the optimal dynamic programming algorithm of [11] to determine the swapping tree with the lowest expected latency of establishing an EP across (s, d) . We select the path P corresponding to the leaf nodes of this optimal swapping tree, as the entanglement path to generate the EP over (s, d) . Initially, we start generating link-EPs over network links on P . Then, whenever a (link or subpath) EP over this path P is generated or destroyed or a sufficient amount of time has passed, we make a determination to either perform an ES operation or *wait*. This determination is made by considering all possible ES operations available at the given point of time, and for each potential ES X , estimating the minimum latency of generating the target EP over (s, d) if the ES X is performed. Before presenting the pseudo-code, we discuss two concepts.

Active Links and Trigger Events. A link may be *active* or *inactive*; an active link continuously attempts to generate a link-EP over it. Our motivation for inactivating links, as needed, is two-fold: (i) Inactivating a link limits the memory requirements at each network node to at most two; (ii) Allowing a link to continue to generate link-EPs while a derived EP exists—leads to existence of multiple EPs over a pair of nodes with potentially different ages; handling such a generalized network state in our algorithm above is very challenging, and is deferred to our future work.

In our context, the *trigger events* are: (i) Generation of a new EP, (ii) Destruction of an EP (due to decoherence or ES failure), (iii) *Passage of “sufficient” time* after a trigger event. The motivation behind the last trigger event is to re-determine, based on the updated ages of available EPs, if some pair of available EPs should be swapped. We thus pick the “sufficient time” to be $50\mu\text{s}$, which is equal to half of the time it takes to *attempt* a link-EP generation. Thus, if $50\mu\text{s}$ passes without an event being triggered, then we trigger the *passage of sufficient time* event.

GREEDY-SP Algorithm Pseudo Code.

Input: (i) A quantum network, a pair of network nodes (s, d) .

Output: An EP over (s, d) with minimum generation latency.

Algorithm.

- 1) Determine the routing path P from the optimal swapping tree over (s, d) obtained using DP from [11] (see §IV-A).
- 2) Turn all links on P active.
- 3) Continuously (try to) generate a link-EP across all active links in P .
 - Once a link-EP E is generated over a link e , make e inactive. The link e is activated again when either E or an EP derived from E (via ESS) is destroyed (due to decoherence or ES failure).
- 4) On one of the trigger events (as defined above), do:
 - a) Call SWAP-OR-WAIT (see below) to determine which of the available pair of EPs to swap or *wait* (for the next trigger event).
 - b) Perform the ES operation or *wait* based on (a).

SWAP-OR-WAIT Subroutine.

Input: (i) Set of available EPs (with ages of their qubits), and (ii) Set of active links and the duration for which they have been generating EPs.

Output: Either a pair of available EPs that should be ES-ed next, or wait.

Algorithm.

- 1) Let $\mathcal{E} \leftarrow \{\text{Set of available EPs}\} \cup \{\text{EPs corresponding to active links}\}$
- 2) For each pair of EPs (E_1, E_2) from \mathcal{E} that can be swapped, do:
 - a) Using DP-WITH-AGES (§IV-A), determine the minimum expected latency L_{E_1, E_2} of generating the final EP over (s, d) , if (E_1, E_2) EPs are swapped.
- 3) Pick (E_1, E_2) corresponding to the lowest L_{E_1, E_2} . If both E_1 and E_2 are available, return (E_1, E_2) , else return wait.

A. DP-WITH-AGES: Estimating Minimum Latency

In [11], the authors develop a dynamic programming (DP) algorithm to determine an optimal (i.e., with minimum expected generation latency) swapping tree to generate an EP over a given source-destination pair in a quantum network. For our purposes of estimating the minimum latency, given a set of available EPs and the next ES operation, we need to modify the DP algorithm of [11] appropriately. We briefly describe the original DP algorithm of [11] and then discuss its modification.

DP Algorithm [11]. Consider a quantum network with nodes $V = \{1, 2, \dots, n\}$, and a pair (s, d) of source-destination nodes over which we want to generate an EP with minimum expected latency. The work in [11] develops a dynamic programming (DP) algorithm (referred to as DP, here) that constructs a swapping tree X over a path connecting s and d such that X incurs minimum generation latency. The DP algorithm is based on a recursive formula governing the generation latency of swapping trees under certain reasonable assumptions. In particular, let $T[i, j, h]$ represent the optimal expected latency of generating EP pairs over (i, j) using a swapping tree of height at most h . Then, we have:

1. $T[i, j, 0] = \frac{t_g}{p_g^2 p_{ob}}$ for adjacent nodes (i, j) .
2. $T[i, j, h] = \min(T[i, j, h-1], (\frac{3}{2}B + t_b)/p_b)$, where
$$B = \min_{k \in V} \max(T[i, k, h-1], T[k, j, h-1])$$

The above equations can be used to compute the optimal generation latency as well as the corresponding optimal swapping tree. The authors in [11] also incorporate decoherence and other constraints in the above algorithm. In particular, to incorporate decoherence constraints, [11] introduces four additional parameters to the above function $T[\cdot]$, corresponding to the depths of the four grandchildren of a swapping tree's root. We refer the reader to [11] for more details.

DP-WITH-AGES: Incorporating Ages of Available EPs.

The above-described DP starts with a “base state” wherein the network links have not yet started generating the corresponding link-EPs. In contrast, in our context, we need to estimate the minimum generation latency of the target EP (see Line 2(a) of Algorithm SWAP-OR-WAIT) when some of the EPs are already available—and more importantly, may be of non-zero ages. To estimate the optimal generation latency for such a generic network state (i.e., with already available EPs), we divide the entire possible age-range $[0, \tau]$ of a qubit into sufficiently many (we use 100) subranges and estimate the minimum generation latency of the target EP for each age-range recursively (using dynamic programming) as described below. In particular, let $T[i, j, h, a]$ represent the minimum expected latency of generating an EP pair over (i, j) using a swapping tree of height at most h such that the age of the target EP is within the age-interval a . Then, we have:

1. $T[i, j, 0, a] = 0$, if (i, j) is an available EP of age in interval a .
2. $T[i, j, 0, [0, x]] = \frac{t_g}{p_g^2 p_{ob}} - t$ for adjacent nodes (i, j)

Above, t is the time already spent generating a link-EP over (i, j) and $[0, x]$ is the first subrange of $[0, \tau]$.

3. $T[i, j, h, a] = \min(T[i, j, h-1, a], (\frac{3}{2}B + t_b)/p_b)$, where
$$B = \min_{k \in V, a_i \leq a, a_j \leq a} \max(T[i, k, h-1, a_i], T[k, j, h-1, a_j])$$

V. Multi-Path Generalization

In this section, we consider the generalization of the SP-SWO problem wherein we allow for multiple entanglement paths for generating EPs over a given (s, d) pair.

Motivation. In a large quantum network, there may be multiple paths available for independent generation of the desired remote EP, and the stochasticity of the processes imply that using multiple paths can result in lower expected latency (any one of the paths needs to succeed). Thus, many works have considered multi-paths [10, 11] for entanglement generation. In our context, we generalize the SP-SWO problem for multi-paths, i.e., for real-time determination of entanglement path and ES operations during the course of EP generation.

MP-SWO Problem (Multi-Path Swapping Order). Given a quantum network and a source-destination pair $\{(s, d)\}$, the MP-SWO problem is to determine the choice of entanglement path and ordering of ES operations (in real-time as the EPs are generated) so as to minimize the actual/real-time generation latency of an EP (s, d) under the decoherence and node constraints (as defined in SP-SWO problem formulation).

Multi-Path Optimal Static Approach. In recent work, [11] presents a linear programming (LP) based algorithm (referred to as LP, here) for constructing an “aggregated” structure of multiple swapping trees that yields an optimal generation rate of EP over a given source-destination pair. Their technique is based on constructing an appropriate hypergraph that has embedded in it all possible swapping trees in the given quantum network, and then developing a linear program to

determine the optimal hyperflow in the constructed hypergraph; the optimal hyperflow essentially represents the optimal “aggregated” tree structure for the generation of EPs over the given source-destination pair. We use the above approach as a baseline and develop a real-time ES-ordering approach as described below.

GREEDY-MP Algorithm for MP-SWO. We utilize the LP algorithm above first to obtain multiple paths (not necessarily disjoint) from s to d that can/should be used to establish EPs (s, d); these paths can be obtained by extracting a sufficient number of swapping trees from the LP “aggregated tree” solution. Then, we proceed similarly to the previous GREEDY-SP algorithm for the single path. In particular, as in GREEDY-SP, we activate all the links and continuously attempt to generate link-EPs over them. Then, on any trigger event (as defined before), we call the SWAP-OR-WAIT procedure on each of the paths *independently* to determine the best pair of EPs to swap (or wait) on each path. Since the paths may not be disjoint, the best pair of EPs chosen for a path may involve EPs that are part of multiple paths. Thus, we sort these best pairs from the paths based on the expected generation latencies of the corresponding paths and perform the swaps in that order (as long as the required EPs haven’t been consumed by a previous swap). We describe this formally in the below pseudo code. The algorithm terminates when an EP has been established across (s, d).

GREEDY-MP Psuedo-Code.

Input: (i) A quantum network, and a source-destination pair (s, d).

Output: An EP over the nodes (s, d) in minimum latency.

Algorithm.

- 1) Select a set of routing paths $\{P_i\}_{[k]}$ using LP, as described above.
- 2) Turn on all links in each P_i as active.
- 3) Continuously (try to) generate a link-EP across all active links in each P_i .
 - Once a link EP E is generated over a link e , make e inactive. The link e is activated again when either E or an EP derived from it (via ESS) is destroyed (due to decoherence or ES failure).
- 4) On one of the trigger events, do
 - a) For each path P_i independently, use SWAP-OR-WAIT (§IV) to determine the pair of EPs D_i to be swapped next. Let L_i be the expected generation latency of the target EP over path P_i after performing D_i . If D_i is null (i.e., when wait is returned), then let L_i be infinity.
 - b) Traverse the D_i ’s in ascending order of L_i ’s, and perform the ES corresponding to each D_i only if neither of the EPs in D_i have been consumed by a D_j .

VI. Evaluation

This section evaluates our algorithm on different network graphs with varying parameters.

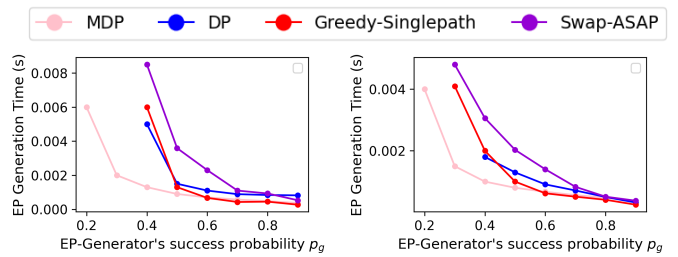


Fig. 3: SP-SWO (single-path) Problem. EP generation latency of various algorithms, for small networks (5 nodes) and small decoherence times (1.5×10^{-4} (left) and 3×10^{-4} (right) seconds), for varying node EP-generator’s success probabilities.

Algorithms Compared. For the SP-SWO problem, we compare four algorithms: (i) GREEDY-SP from §IV; (ii) Optimal dynamic programming algorithm (DP) from [11] which returns an optimal static swapping tree; (iii) SWAP-ASAP from [12], which performs an ES operation over EPs as soon as any are available; (iv) Markov Decision Process based (MDP) algorithm from [12]. For the multi-path MP-SWO problem, we compare two algorithms: (v) GREEDY-MP from §V; (vi) LP from [11] (see §V).

Implementation Details. Given a quantum network and a single (s, d) pair, DP picks a *static* swapping tree of minimum expected latency and performs swap operations in the order given by the tree. MDP maintains a state of all possible EPs over all possible ages, resulting in a state space of size $\mathcal{O}(t^n)$, where n is the number of nodes in the network and $t = 2\tau/3t_g$ is a discretized representation of the decoherence threshold. It then pre-computes a policy that, for each state, returns an action: *wait* or perform a certain swapping operation. Given a network path and an objective of establishing an EP across the end nodes, SWAP-ASAP continuously attempts link generation of active links. When a pair of EPs that can be swapped become available, SWAP-ASAP swaps them. Note that both MDP and SWAP-ASAP take a network path rather than a network graph as input; thus, as for the case of GREEDY-SP, they operate on the path corresponding to the optimal swapping tree of DP. In all approaches, the links become active or inactive as in our approaches GREEDY-SP and GREEDY-MP.

Generating Random Networks. We use a quantum network spread over an area of $100km \times 100km$. We use the Waxman model [18], which has been used to create Internet topologies. We vary the number of nodes in the network from 20 to 60, with 40 as the default value. We generate random networks by varying the following parameters.

- Number of nodes (default = 40)
- Swapping success probability p_{ob} (default = 0.5)
- Node EP-generator’s success probability p_g (default = 0.5)
- Decoherence threshold τ (default = 1.5s)
- Physical distance between the source-destination pair s and d (default = 20-50km).

In all graphs, each data point is the average generation latency of 10 runs of the corresponding algorithm on a set of input

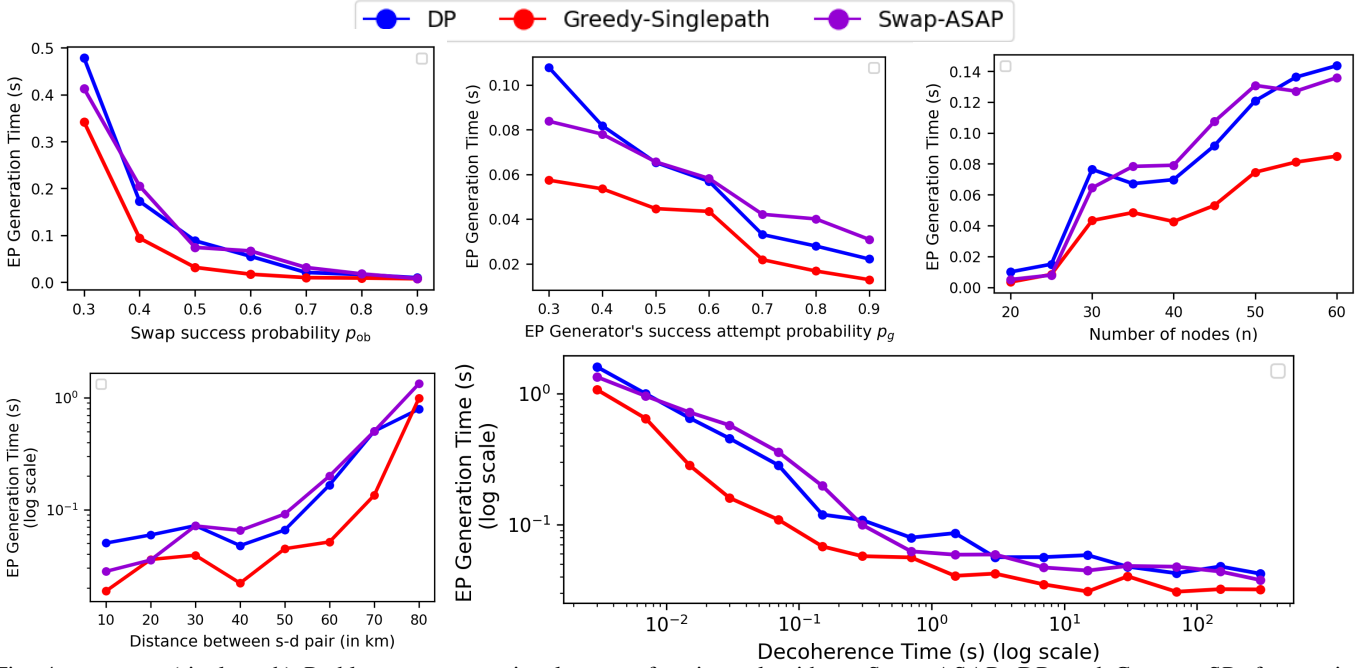


Fig. 4: SP-SWO (single-path) Problem. EP generation latency of various algorithms, SWAP-ASAP, DP, and GREEDY-SP, for varying parameters. MDP is not shown due to its prohibitive runtime for these parameter values (see text).




Fig. 5: MP-SWO (multi-path) Problem. EP generation latencies of GREEDY-MP and LP for varying parameters.

parameters.

Evaluation Results. For the single-path (SP-SWO) problem, we evaluate the four algorithms on inputs generated by varying the parameters described above. We note that MDP takes prohibitively long to run on paths with more than 7 nodes or settings with decoherence time more than 3×10^{-4} sec with paths of more than 5 nodes. Thus, we evaluate MDP on small networks with 5 nodes and two decoherence times, viz., 0.15ms and 0.3ms. See Figs 3a-3b. For such small networks, we observe that MDP performs the best, as expected, as it is the optimal algorithm. However, we note that GREEDY-SP also performs very close to the optimal MDP for $p_g \geq 0.6$ values. For larger networks (Figs 4a-4e), we observe that GREEDY-SP consistently performs better than DP and SWAP-ASAP for most instances; in particular, it achieves generation latencies of up to 40% lower than the optimal static approach (DP) and up to 45% lower than the naive adaptive approach SWAP-ASAP. Finally, for the multi-path (MP-SWO) problem, we compare the adaptive GREEDY-MP and the static-optimal LP approaches. See Figs. 5a-5b. We observe that GREEDY-MP performs 10-20% better than LP in most cases.

Runtime Overhead. The running time of the selection phase of our GREEDY-SP and GREEDY-MP algorithms is $100 \mu\text{s}$ and $700 \mu\text{s}$ for the default parameter values on an 8-core Intel Core i7-11700 system. Since the selection phase involves computing the expected latency for each possible choice independently—the running time can be easily parallelized and reduced to the order of a few μs . Thus, the running overhead of our approaches has minimal impact on the decoherence of qubits (since realistic values of decoherence times is of the order of several seconds) and generation latency (which is of the order of 10s-100s of milliseconds).

VII. Conclusion

In this work, we have developed an adaptive EP generation algorithm that chooses the entanglement route and/or swapping operations at EP-generation time, based on the network state at the time. In our future work, we seek to extend our work in the following ways: (i) Consider multiple source-destination pairs; (ii) Continuous generation of EPs with node-memory constraints; (iii) Consider discarding EPs that are likely to decohere before the target EP can be generated.

REFERENCES

- [1] Charles H Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K Wootters. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Physical review letters*, 70(13):1895, 1993.
- [2] Marcello Caleffi. Optimal routing for quantum networks. *IEEE Access*, 2017.
- [3] Kaushik Chakraborty, David Elkouss, Bruno Rijsman, and Stephanie Wehner. Entanglement distribution in a quantum network: A multicommodity flow-based approach. *IEEE Transactions on Quantum Engineering*, 1:1–21, 2020. URL: <https://api.semanticscholar.org/CorpusID:219124346>.
- [4] Kaushik Chakraborty, Filip Rozpdek, Axel Dahlberg, and Stephanie Wehner. Distributed routing in a quantum internet. *ArXiv*, abs/1907.11630, 2019.
- [5] Kaushik Chakraborty, Filip Rozpedek, Axel Dahlberg, and Stephanie Wehner. Distributed routing in a quantum internet, 2019.
- [6] Daniele Cuomo, Marcello Caleffi, Kevin Krsulich, Filippo Tramonto, Gabriele Agliardi, Enrico Prati, and Angela Sara Cacciapuoti. Optimized compiler for distributed quantum computing. *ACM Transactions on Quantum Computing*, 4(2), feb 2023.
- [7] DGBJ Dieks. Communication by epr devices. *Physics Letters A*, 92(6):271–272, 1982.
- [8] L-M Duan, Mikhail D Lukin, J Ignacio Cirac, and Peter Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–418, 2001.
- [9] Zachary Eldredge, Michael Foss-Feig, Jonathan A Gross, Steven L Rolston, and Alexey V Gorshkov. Optimal and secure measurement protocols for quantum sensor networks. *Physical Review A*, 97(4):042337, 2018.
- [10] Mohammad Ghaderibaneh, Himanshu Gupta, and CR Ramakrishnan. Generation and distribution of ghz states in quantum networks. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 1, pages 1120–1131. IEEE, 2023.
- [11] Mohammad Ghaderibaneh, Caitao Zhan, Himanshu Gupta, and C. Ramakrishnan. Efficient quantum network communication using optimized entanglement-swapping trees. 12 2021.
- [12] Álvaro G Iñesta, Gayane Vardoyan, Lara Scavuzzo, and Stephanie Wehner. Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs. *npj Quantum Information*, 9(1):46, 2023.
- [13] Peter Komar, Eric M Kessler, Michael Bishof, Liang Jiang, Anders S Sørensen, Jun Ye, and Mikhail D Lukin. A quantum network of clocks. *Nature Physics*, 10(8):582–587, 2014.
- [14] Mihir Pant, Hari Krovi, Don Towsley, Leandros Tassiulas, Liang Jiang, Prithwish Basu, Dirk Englund, and Saikat Guha. Routing entanglement in the quantum internet. *npj Quantum Information*, 08 2017.
- [15] Nicolas Sangouard, Christoph Simon, Hugues De Riedmatten, and Nicolas Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33, 2011.
- [16] Shouqian Shi and Chen Qian. Concurrent entanglement routing for quantum networks: Model and designs. SIGCOMM '20, New York, NY, USA, 2020. Association for Computing Machinery. URL: <https://doi.org/10.1145/3387514.3405853>.
- [17] Christoph Simon. Towards a global quantum network. *Nature Photonics*, 11(11):678–680, 2017.
- [18] Bernard M Waxman. Routing of multipoint connections. *IEEE journal on selected areas in communications*, 6(9):1617–1622, 1988.