

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/TQE.2020.DOI

# Efficient Quantum Network Communication using Optimized Entanglement-Swapping Trees

MOHAMMAD GHADERIBANEH, CAITAO ZHAN, HIMANSHU GUPTA, C. R. RAMAKRISHNAN

Department of Computer Science, Stony Brook University, Stony Brook, NY 11790, USA

Corresponding author: Mohammad Ghaderibaneh (email: mghaderibaneh@cs.stonybrook.edu).

This work was supported by NSF awards FET-2106447 and CNS-2128187, and a Cisco industry grant.

**ABSTRACT** Quantum network communication is challenging, as the No-cloning theorem in quantum regime makes many classical techniques inapplicable; in particular, direct transmission of qubit states over long distances is infeasible due to unrecoverable errors. For long-distance communication of unknown quantum states, the only viable communication approach (assuming local operations and classical communications) is teleportation of quantum states, which requires a prior distribution of entangled pairs (EPs) of qubits. Establishment of EPs across remote nodes can incur significant latency due to the low probability of success of the underlying physical processes. The focus of our work is to develop efficient techniques that minimize EP generation latency. Prior works have focused on selecting entanglement *paths*; in contrast, we select *entanglement swapping trees*—a more accurate representation of the entanglement generation structure. We develop a dynamic programming algorithm to select an optimal swapping-tree for a single pair of nodes, under the given capacity and fidelity constraints. For the general setting, we develop an efficient iterative algorithm to compute a set of swapping trees. We present simulation results which show that our solutions outperform the prior approaches by an order of magnitude and are viable for long-distance entanglement generation.

**INDEX TERMS** Quantum Communications, Quantum Networks

## I. INTRODUCTION

FUNDAMENTAL advances in physical sciences and engineering have led to the realization of working quantum computers (QCs) [1], [2]. However, there are significant limitations to the capacity of individual QC [3]. Quantum networks (QNs) enable the construction of large, robust, and more capable quantum computing platforms by connecting smaller QCs. Quantum networks [4] also enable various important applications [5]–[9]. However, quantum network communication is challenging — e.g., physical transmission of quantum states across nodes can incur irreparable communication errors, as the No-cloning Theorem [10] proscribes making independent copies of arbitrary qubits. At the same time, certain aspects unique to the quantum regime, such as entangled states, enables novel mechanisms for communication. In particular, teleportation [11] transfers quantum states with just classical communication, but requires an a priori establishment of entangled pairs (EPs). This paper presents techniques for efficient establishment of EPs in a network.

Establishment of EPs over long distances is challeng-

ing. Coordinated entanglement swapping (e.g. DLCZ protocol [12]) using quantum repeaters can be used to establish long-distance entanglements from short-distance entanglements. However, due to low probability of success of the underlying physical processes (short-distance entanglements and swapplings), EP generation can incur significant latency— of the order of 10s to 100s of seconds between nodes 100s of kms away [13]. Thus, we need to develop techniques that can facilitate fast generation of long-distance EPs. We employ two strategies to minimize generation latencies: (i) select optimal swapping **trees** (not, just paths as in prior works [14]–[17]) with a protocol that retains unused EPs; (ii) use multiple trees for each given node pair; this reduces effective latency by using all available network resources. In the above context, we address the following problems: (i) **QNR-SP Problem**: Given a single  $(s, d)$  pair, select a minimum-latency swapping tree under given constraints. (ii) **QNR Problem**: Given a set of source-destination  $(s, d)$  pairs, select a set of swapping trees for each pair with maximum aggregate EP generation rate, under fidelity and resource

constraints.


To the best of our knowledge, no prior work has addressed the problem of selecting an efficient swapping-tree for entanglement routing; they all consider selecting routing *paths* ([18] selects a path using a metric based on balanced trees; see §III-B). Almost all prior works have considered the “waitless” model, wherein all underlying physical processes much succeed *near-simultaneously* for an EP to be generated; this model incurs minimal decoherence, but yields very low EP generation rates. In contrast, we consider the “waiting” protocol, wherein, at each swap operation, the earlier arriving EP waits for a limited time for the other EP to be generated. Such an approach with efficient swapping trees yields high entanglement rates; the potential decoherence risk can be handled by discarding qubits that “age” beyond a certain threshold.

**Our Contributions.** We formulate the entanglement routing problem (§III) in QNs in terms of selecting optimal swapping *trees* in the “waiting” protocol, under fidelity constraints. In this context, we make the following contributions:

- 1) For the QNR-SP problem, we design an optimal algorithm with fidelity and resource constraints (§IV).
- 2) Though polynomial-time, the above optimal algorithm has high time complexity; we thus also design a near-linear time heuristic for the QNR-SP problem based on an appropriate metric which essentially restricts the solutions to balanced swapping trees (§V).
- 3) For the general QNR problem, we design an efficient iterative augmenting-tree algorithm (§VI), and show its effectiveness w.r.t. an optimal LP solution based on hypergraph-flows.
- 4) We conduct extensive evaluations (§VII) using Net-Squid simulator, and show that our solutions outperform the prior approaches by an order of magnitude, while incurring little fidelity degradation. We also show that our schemes can generate high-fidelity EPs over nodes 500-1000kms away.

## II. QC BACKGROUND

**Qubit States.** Quantum computation manipulates *qubits* analogous to how classical computation manipulates *bits*. At any given time, a bit may be in one of two states, traditionally represented by 0 and 1. A quantum state represented by a *qubit* is a *superposition* of classical states, and is usually written as  $\alpha_0|0\rangle + \alpha_1|1\rangle$ , where  $\alpha_0$  and  $\alpha_1$  are *amplitudes* represented by complex numbers and such that  $|\alpha_0|^2 + |\alpha_1|^2 = 1$ . Here,  $|0\rangle$  and  $|1\rangle$  are the standard (orthonormal) *basis* states; concretely, they may represent physical properties such as spin (down/up), polarization, charge direction, etc. When a qubit such as above is *measured*, it collapses to a  $|0\rangle$  state with a probability of  $|\alpha_0|^2$  and to a  $|1\rangle$  state with a probability of  $|\alpha_1|^2$ . In general, a state of an  $n$  qubit system can be represented as  $\sum_{i=0}^{2^n-1} \alpha_i|i\rangle$  where “ $i$ ” in  $|i\rangle$  is  $i$ ’s bit representation.



**FIGURE 1.** (a) Teleportation of  $|q\rangle$  from  $A$  to  $B$ , while consuming an entangled pair  $(e_1, e_2)$ . (b) Entanglement swapping over the triplet of nodes  $(A, B, C)$ , which results in  $A$ ’s qubit entangled with  $C$ ’s qubit. This can be viewed as a teleportation of  $e_2$  from node  $B$  to  $C$ .

**Entanglement.** Entangled pure<sup>1</sup> states are multi-qubit states that cannot be “factorized” into independent single-qubit states. E.g., the 2-qubit state  $\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$ ; this particular system is a *maximally-entangled* state. We refer to maximally-entangled pairs of qubits as EPs. The surprising aspect of entangled states is that the combined system continues to stay entangled, even when the individual qubits are physically separated by large distances. This facilitates many applications, e.g., teleportation of qubit states by local operations and classical information exchange, as described next.

**Teleportation.** Direct transmission of quantum data is subject to unrecoverable errors, as classical procedures such as amplified signals or re-transmission cannot be applied due to quantum no-cloning [10], [20].<sup>2</sup> An alternative mechanism for quantum communication is *teleportation*, Fig. 1 (a), where a qubit  $q$  from a node  $A$  is recreated in another node  $B$  (*while “destroying” the original qubit  $q$* ) using only classical communication. However, this process requires that an EP already established over the nodes  $A$  and  $B$ . Teleportation can thus be used to reliably transfer quantum information. At a high-level, the process of teleporting an arbitrary qubit, say qubit  $q$ , from node  $A$  to node  $B$  can be summarized as follows:

- 1) an EP pair  $(e_1, e_2)$  is generated over  $A$  and  $B$ , with  $e_1$  stored at  $A$  and  $e_2$  stored at  $B$ ;
- 2) at  $A$ , a *Bell-state measurement* (BSM) operation over  $e_1$  and  $q$  is performed, and the 2 classical bits measurement output  $(c_1, c_2)$  is sent to  $B$  through the classical communication channel; at this point, the qubits  $q$  and  $e_1$  at  $A$  are destroyed.
- 3) manipulating the EP-pair qubit  $e_2$  at  $B$  based on received  $(c_1, c_2)$  changes its state to  $q$ ’s initial state.

Depending on the physical realization of qubits and the BSM operation, it may not always be possible to successfully generate the 2 classical bits, as the BSM operation is stochastic.

**Entanglement Swapping (ES).** Entanglement swapping is an application of teleportation to generate EPs over remote nodes. See Fig. 1 (b). If  $A$  and  $B$  share an EP and  $B$  teleports

<sup>1</sup>In this work, we largely deal with only pure qubit states. Entanglement of general mixed states is defined in terms of separation of density matrices [19].

<sup>2</sup>Quantum error correction mechanisms [21], [22] can be used to mitigate the transmission errors, but their implementation is very challenging and is not expected to be used until later generations of quantum networks.

its qubit to  $C$ , then  $A$  and  $C$  end up sharing an EP. More elaborately, let us assume that  $A$  and  $B$  share an EP, and  $B$  and  $C$  share a separate EP. Now,  $B$  performs a BSM on its two qubits and communicates the result to  $C$  (teleporting its qubit that is entangled with  $A$  to  $C$ ). When  $C$  finishes the protocol, it has a qubit that is entangled with  $A$ 's qubit. Thus, an entanglement swapping (ES) operation can be looked up as being performed over a triplet of nodes  $(A, B, C)$  with EP available at the two pairs of adjacent nodes  $(A, B)$  and  $(B, C)$ ; it results in an EP over the pair of nodes  $(A, C)$ .

**Fidelity: Decoherence and Operations-Driven.** Fidelity is a measure of how close a realized state is to the ideal. Fidelity of qubit decreases with time, due to interaction with the environment, as well as gate operations (e.g., in ES). Time-driven fidelity degradation is called *decoherence*. To bound decoherence, we limit the aggregate time a qubit spends in a quantum memory before being consumed. With regards to operation-driven fidelity degradation, Briegel et al. [23] give an expression that relates the fidelity of an EP generated by ES to the fidelities of the operands, in terms of the noise introduced by swap operations and the number of link EPs used. The order of the swap operations (i.e., the structure of the swapping tree) does not affect the fidelity. Thus, the operation-driven fidelity degradation of the final EP generated by a swapping-tree  $T$  can be controlled by limiting the number of leaves of  $T$ , assuming that the link EPs have uniform fidelity (as in [15]).

Entanglement Purification [23, e.g.] and Quantum Error Correction [24, e.g.] have been widely used to combat fidelity degradation. Our work focuses on optimally scheduling ES operations with constraints on fidelity degradation, without purification or error correction.

**Quantum Memories.** Multiple quantum memories have been recently proposed to bring quantum networks into realization. Types of quantum memories that support BSM measurements and gate unitary operations, and probably have a long decoherence time can be used in quantum communications. Most of them are matter-based which have photonic interface to produce matter-matter entanglement over two neighboring nodes (see below). At a high-level, there are three different quantum memory platforms: quantum dots, trapped atoms or ions, and colour centers in diamond. Each has its own physical characteristics and applications. While quantum dots have the ability to process quantum information very fast, they exhibit a very low decoherence time among others [25], [26]. To overcome the low efficiency of single atom-photon coupling process, atomic ensemble schemes have been proposed [12] where along with dynamic decoupling and cooling techniques, decoherence times of a few seconds have been achieved [27]–[29]. For trapped ion memories, decoherence times from several minutes to few hours have been demonstrated [30], [31]. To further increase the entanglement generation rate, [32] proposes a way to use a single silicon–vacancy (SiV) colour center in diamond to perform asynchronous photonic BSM at the node located in

the middle of two adjacent quantum nodes.

#### A. GENERATING ENTANGLEMENT PAIRS (EPs)

As described above, teleportation, which is the only viable means of transferring quantum states over long distances, requires an a priori distribution of EPs. Thus, we need efficient mechanisms to establish EPs across remote QN nodes; this is the goal of our work. Below, we start with describing how EPs are generated between adjacent (i.e., one-hop away) nodes, and then discuss how EPs across a pair of remote nodes can be established via ESs.


**Generating EP over Adjacent Nodes.** The physical realization of qubits determines the technique used for sharing EPs between adjacent nodes. The *heralded entanglement* process [14], [18] to generate an atom-atom EP between adjacent nodes  $A$  and  $B$  is as follows:

- 1) Generate an entangled pair of atom and a telecom-wavelength photon at node  $A$  and  $B$ . Qubits at each node are generally realized in an atomic form for longer-term storage, while photonic qubits are used for transmission.
- 2) Once an atom-photon entanglement is locally generated at each node (at the same time), the telecom-photons are then transmitted over an optical fiber to a photon-photon/optical BSM device  $C$  located in the middle of  $A$  and  $B$  so that the photons arrive at  $C$  at the same time.
- 3) The device  $C$  performs a BSM over the photons, and transmits the classical result to  $A$  or  $B$  to complete ES.

Other entanglement generation processes have been proposed [33]; our techniques themselves are independent of how the link EP are generated.

**Generating EP between Remote Nodes.** Now, EP between non-adjacent nodes connected by a path in the network can be established by performing a sequence of ESs at intermediate nodes; this requires an a priori EP over each of the adjacent pairs of nodes in the path. For example, consider a path of nodes  $x_0, x_1, x_2, x_3, x_4, x_5$ , with an EP between every pair of adjacent nodes  $(x_i, x_{i+1})$ . Thus, each node  $x_i$  ( $1 \leq i \leq 4$ ) has two qubits, one of which is entangled with  $x_{i-1}$  and the other with  $x_{i+1}$ . Nodes  $x_0$  and  $x_5$  have only one qubit each. To establish an EP between  $x_0$  and  $x_5$ , we can perform a sequence of entanglement swappings (ESs) as shown in Fig. 2. Similarly, the sequence of ES over the following triplets would also work:  $(x_2, x_3, x_4)$ ,  $(x_2, x_4, x_5)$ ,  $(x_0, x_1, x_2)$ ,  $(x_0, x_2, x_5)$ .

**Swapping Trees.** In general, given a path  $P = s \rightsquigarrow d$  from  $s$  to  $d$ , any complete binary tree (called a *swapping 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 in  $P$ . Every pair of siblings  $(A, B)$  and  $(B, C)$  perform an ES over  $(A, B, C)$  to yield an EP over  $(A, C)$ —their parent. See Fig. 2. Note that subtrees of a swapping tree execute in parallel. Different swapping trees over the same path  $P$  can



**FIGURE 2.** A swapping tree over a path. The leaves of the tree are the path-links, which generate link-EPs continuously.

have different performance characteristics, as discussed later (see Fig. 4).

**Expected Generation Latency/Rate of EPs.** In general, our goal is to continuously generate EPs at some rate using a swapping tree, using continuously generated EPs at the leaves. The stochastic nature of ES operations means that an EP at the tree’s root will be successfully generated only after many failed attempts and hence significant latency. We refer to this latency as the *generation latency* of the EP at the root, and in short, just the generation latency of the tree. EP generation rate is the inverse of its generation latency. Whenever we refer to generation latency/rate, we implicitly mean *expected* generation latency/rate.

**Two Generation Protocols:** `WaitLess` and `Waiting`  
When a swapping tree is used to (continuously) generate EPs, there are two fundamentally different generation protocols [13], [34].

- **WaitLess Protocol.** In this model, all the underlying processes, including link EP generations and atomic BSMs are synchronized. If all of them succeed then the end-to-end EP is generated. If *any* of the underlying processes fail, then all the generated EPs are discarded and the whole process starts again from scratch (from generation of EP at links). In the `WaitLess` protocol, all swapping trees over a given path  $P$  incur the same generation latency—thus, here, the goal is to select an optimal path  $P$  (as in [14], [15]).
- **Waiting Protocol.** In `Waiting` protocol, 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. Using such storage, we preclude discarding successfully generated EPs, and thus, reduce the overall latency in generation of a root-level EP. E.g., let  $(A, B)$  and  $(B, C)$  be two siblings in a swapping tree and EP for  $(A, B)$  is generated first. Then, EP  $(A, B)$  may wait for the EP  $(B, C)$  to be successfully generated. Once the EP  $(B, C)$  is generated, the ES operation is done over the triplet  $(A, B, C)$  to generate the EP  $(A, C)$ . If the EP  $(A, B)$  waits beyond a certain threshold, then it may decohere.

**Hardware Requirement Differences.** `WaitLess` protocols can generate EPs without quantum memories in a relay fashion if the EP generation among adjacent nodes can be

tightly synchronized. In contrast, `Waiting` protocols benefit from memories with good coherence times (§VII).

**Why `Waiting`’s Entanglement Generation Rate is Never Worse.** The focus of the `WaitLess` protocol is to avoid qubit decoherence due to storage. But it results in very low generation rates due to a very-low probability of *all* the underlying processes succeeding at the same time. However, since qubit coherence times are typically higher than the link-generation latencies<sup>3</sup>, an appropriately designed `Waiting` protocol will always yield better generation rates *without significantly compromising the fidelity* (see Theorem 1). The key is to bound the waiting time to limit decoherence as desired; e.g., in our protocol, we restrict to trees with high expected fidelities (§III), and discard qubits that “age” beyond a threshold (§IV-B). Both protocols use the same number of quantum memories (2 per node), though the `Waiting` protocols will benefit from low-decoherence memories; other hardware requirements also remain the same.

**Theorem 1:** Consider a quantum network, a path  $P$ , a swapping-tree  $T$  over  $P$ , a `WaitLess` protocol  $X$ , and a `Waiting` protocol  $Y$ . Protocol  $Y$  discards qubits that age (stay in memory) beyond a certain threshold  $\tau$  (presumably, equal to the coherence time). We claim that  $Y$ ’s EP generation rate will at least be that of  $X$ , irrespective of  $\tau$  and  $T$  (as long as it is over  $P$ ), while ensuring that EPs generated by  $Y$  are formed by non-decohered qubits and the operation-driven fidelity degradation of  $Y$  EPs is same as  $X$ . ■

The above theorem suggests that `Waiting` approach is always a better performing approach, irrespective of the decoherence time/limitations. See proof in Appendix B.

### III. MODEL, PROBLEM, AND RELATED WORKS

In this section, we discuss our network model, formulate the problem addressed, and discuss related work.

**Network Model.** We denote a quantum network (QN) with a graph  $G = (V, E)$ , with  $V = \{v_1, v_2, \dots, v_n\}$  and  $E = \{(v_i, v_j)\}$  denoting the set of nodes and links respectively. Pairs of nodes connected by a link are defined as *adjacent* nodes. We follow the network model in [18] closely. Thus, each node has an atom-photon EP generator with generation latency ( $t_g$ ) and probability of success ( $p_g$ ). Generation latency is the time between successive attempts by the node to excite the atom to generate an atom-photon EP; this implicitly includes the times for photon transmission, optical-BSM latency, and classical acknowledgement. *For clarity of presentation* and without loss of generality, we assume homogeneous network nodes with same parameter values. The generation rate is the inverse of generation latency, as before. A node’s atom-photon generation capacity/rate is its aggregate capacity, and may be split across its incident links (i.e., in generation of EPs over its incident links/nodes).

<sup>3</sup>Link generation latencies for 5 to 100km links range from about 3 to 350 milliseconds for typical network parameters [18], while coherence times of few seconds is very realistic (coherence times of several seconds [35], [36] have been shown long ago, and more recently, even coherence times of several minutes [37], [38] to a few hours [39], [40] have been demonstrated.




FIGURE 3. Key notations used.

Each node is also equipped with a certain number of atomic memories to store the qubits of the atom-atom EPs.

A network link is a quantum channel (e.g., using an optical fiber or a free-space link), and, in our context, is used only for establishment of link EP. In particular, a link  $e = (A, B)$  is used to transmit telecom-photons from  $A$  and  $B$  to the photon-photon BSM device in the middle of  $e$ . Thus, each link is composed of two half-links with a probability of transmission success ( $p_e$ ) that decreases exponentially with the link distance (see §VII). The optical-BSM operation has a certain probability of success ( $p_{0b}$ ). To facilitate atom-atom ES operations, each network node is also equipped with an atomic-BSM device with an operation latency ( $t_b$ ) and probability of success ( $p_b$ ). Finally, there is an independent classical network with a transmission latency ( $t_c$ ); we assume classical transmission always succeeds.

**Single vs. Multiple Links Between Nodes.** For our techniques multiple links between a pair of adjacent nodes can be replaced by a single link of aggregated rate/capacity. Hence we assume only a single link between every pair of nodes. However, distinct multiple links between nodes have been used creatively in [14] (which refers to them as multiple channels); thus, we will discuss multiple links further in §VII when we evaluate various techniques. We note that the all-photonic protocol in [41] is essentially a more sophisticated version of the multi-link `WaitLess` protocol in [14] to further minimize memory requirements, but it uses multipartite cluster states which are challenging to create. In either case, in terms of selection of paths/trees, the path-selection techniques from [14] should also apply to the all-photonic protocol with certain modifications to account for how the cluster states are generated.

**EP Generation Latency of a Swapping Tree.** Given a swapping tree and EP generation rates at the leaves (network links), we wish to estimate the generation latency of the EPs over the remote pair corresponding to the tree's root with the `Waiting` protocol. Below, we develop a recursive equation. Consider a node  $(A, C)$  in the tree, with  $(A, B)$  and  $(B, C)$  as its two children. Let  $T_{AB}$ ,  $T_{BC}$ , and  $T_{AC}$  be the corresponding (expected) generation latencies of the EPs over the three pairs of nodes. Below, we derive an expression for  $T_{AC}$  in terms of  $T_{AB}$  and  $T_{BC}$ ; this expression will be sufficient to determine the expected latency of the overall swapping tree by applying the expression iteratively. We start with an observation.

**Observation 1:** If two EP arrival processes  $X_1$  and  $X_2$  are




FIGURE 4. Consider the path in (a). The imbalanced tree of (b) has a higher EP generation rate than that of the balanced tree of (c). Here, the numbers represent the EP generation rates over adjacent links or node-pairs.

exponentially distributed with a mean inter-arrival latency of  $\lambda$  each, then the expected inter-arrival latency of  $\max(X, Y)$  is  $(3/2)\lambda$ .  $\square$

From above, if assume  $T_{AB}$  and  $T_{BC}$  to be exponentially distributed with the same expected generation latency of  $T$ , then the expected latency of both EPs arriving is  $(3/2)T$ . Thus, we have:

$$T_{AC} = (\frac{3}{2}T + t_b + t_c)/p_b, \quad (1)$$

**Remarks.** We make the following remarks regarding the above expression. First, when  $T_{AB} \neq T_{BC}$ , we are able to only derive an upper-bound on  $T_{AC}$  which is given by the above equation but with  $T$  replaced by  $\max(T_{AB}, T_{BC})$ .<sup>4</sup> However, in our methods, the above assumption of  $T_{AB} = T_{BC}$  will hold as we would only be considering “throttled” trees to save on underlying network resources (see §IV). Second, our motivation for the exponential distribution assumption stems from the fact that the EP generation latency at the link level is certainly exponentially distributed if we assume the underlying probabilistic events to have a Poisson distribution. Third, note that the resulting distribution is not exponential. Despite this, we apply the above equation recursively to compute the tree's generation latency. However, in our evaluations, we observe the validity of this approximation since our analysis matches closely with the simulation results. Finally, Eqn. (1) is conservative in the sense that each round of an EP generation of any subtree's root starts from scratch (i.e., with no link EPs from prior round) and ends with either a EP generation at the *whole swapping tree*'s root or an atomic-BSM failure at the subtree's root. We do not “pipeline” any operations across rounds within a subtree, which may lower latency; this is beyond this work's scope.

#### A. PROBLEM FORMULATION

We now formulate the central problem of selecting *multiple* swapping trees for each given source-destination pair. Selection of multiple routes is a well-established strategy [14]–[16] to maximize entanglement rates.

**Quantum Network Routing (QNR) Problem.** Given a quantum network and a set of source-destination pairs  $\{(s_i, d_i)\}$ ,

<sup>4</sup>The 3-over-2 formula as an upper bound has also been corroborated in a recent work [42] which derives analytical bounds on EP latency times in more general contexts.

the QNR problem is to determine a set  $\mathcal{T}_i$  of swapping trees for each pair  $(s_i, d_i)$  such that the sum of the EP rates of all the trees in  $\bigcup_i \mathcal{T}_i$  is maximized under the following constraints:

- 1) *Node Constraints.* For each node, the aggregate resources used by  $\bigcup_i \mathcal{T}_i$  is less than the available resources; we formulate this formally below.
- 2) *Fidelity Constraints.* Each swapping tree in  $\bigcup_i \mathcal{T}_i$  satisfies the following: (a) Number of leaves is less than a given threshold  $\tau_l$ ; this is to limit fidelity degradation due to gate operations. (b) Total memory storage time of any qubit is less<sup>5</sup> than a given *decoherence threshold*  $\tau_d$ .

Informally, the swapping-trees may also satisfy some fairness constraint across the given source-destination pairs. A special case of the above QNR problem is to select a single tree for a source-destination pair; we address this in the next section.

**Formulating Node Constraints.** Consider a swapping tree  $\mathcal{T} \in \bigcup_i \mathcal{T}_i$  over a path  $P$ . For each link  $e \in P$ , let  $R(e, \mathcal{T})$  be the EP rate being used by  $\mathcal{T}$  over the link  $e$  in  $P$ . Let us define  $R_e = \sum_{\mathcal{T}} R(e, \mathcal{T})$ , and let  $E(i)$  be the set of edges incident on  $i$ . Then, the node capacity constraint is formulated as follows.

$$1/t_g \geq \sum_{e \in E(i)} R_e / (p_g^2 p_e^2 p_{ob}) \quad \forall i \in V. \quad (2)$$

The above comes from the fact that to generate a single link EP over  $e$ , each end-node of  $e$  needs to generate  $1/(p_g^2 p_e^2 p_{ob})$  photons successfully, since each photon (from each end-node) has a generation success of  $p_g$  and a transmission success rate of  $p_e$ , and the optical-BSM's success probability over the two successfully arriving photons is  $p_{ob}$ . Note that  $1/t_g$  is a node's total generation capacity. Also, the memory constraint is that for any node  $i$ , the memory available in  $i$  should be more than  $2x + y$  where  $x$  is the number of swapping trees that use  $i$  as an intermediate node and  $y$  is the number of trees that use  $i$  as an end node.

## B. RELATED WORKS

There have been a few works in the recent years that have addressed generating long-distance EPs efficiently. All of these works have focused on selecting an efficient routing path for the swapping process ([18] also selects a path, but using a metric based on balanced trees). In addition, all except [18] have looked at the WaitLess protocol of generating the EPs. Recall that in the WaitLess model, selection of paths suffice, while in the Waiting model, one needs to consider selection of efficient swapping trees with high fidelity. Selection of optimal swapping trees is a fundamentally more challenging problem than selection of

<sup>5</sup>We note that, in our context, the storage time as well as the memory coherence time are statistical quantities due to the underlying statistical mechanisms. However, for the purposes of *selecting* a swapping tree, we use a fixed decoherence threshold  $\tau_d$  value equal to the mean of the distribution of the coherence time (recent work [43] computes optimal cut-offs/thresholds, and their techniques can be used to pick  $\tau_d$ ). When simulating a selected tree for generation of EPs, we can implement coherence time as a statistical measure.

paths—and has not been addressed before, to the best of our knowledge. We start with discussing how the WaitLess model works.

**WaitLess Approaches.** The most recent works to address the above problem are [14] and [15], both of which consider the WaitLess model. In particular, Shi and Qian [14] design a Dijkstra-like algorithm to construct an optimal path between a pair of nodes, when there are multiple links (channels) between adjacent nodes. Then, they use the algorithm iteratively to select multiple paths over multiple pairs of nodes. Chakraborty *et al.* [15] design a multi-commodity-flow like LP formulation to select routing paths for a set of source destination pairs. They map the operation-based fidelity constraint to the path length (as in [23]), and use node copies to model the constraint in the LP. However, they explicitly assume that the link EP generation is deterministic—i.e., always succeeds. Among earlier relevant works, [16] proposes a greedy solution for grid networks, and [17] proposes virtual-path based routing in ring/grid networks.

**Waiting Approach.** Due to photon loss, establishing long-distance entanglement between remote nodes at  $L$  distance by *direct* transmission yields EP rates that decay exponentially with  $L$ . DLCZ protocol [12], [13] broke this exponential barrier using  $2^k$  equidistant intermediate nodes to perform entanglement-swapping operations, implicitly over a balanced binary tree, with a Waiting protocol; this makes the EP generation rate decay only polynomially in  $L$ . More recently, Caleffi [18] formulated the entanglement generation rate on a given path between two nodes, under the more realistic condition where the intermediate nodes in the path may not all be equidistant, but still considered only balanced trees. Their path-based metric was then used to select the optimal path by enumerating over the exponentially many paths in the network.

**Our Approach (vs. [18]).** Though [18] considers only balanced trees, its brute-force algorithm is literally impossible to run for networks more than a few tens of nodes (§VII). In our work, we observe that a path has many swapping trees, and, in general, imbalanced trees may even be better; see Figure 4. Thus, we design a polynomial-time dynamic programming (DP) algorithm that delivers an *optimal* high-fidelity swapping-tree; our DP approach effectively considers all possible swapping trees, not just balanced ones (note that, even over a single path, there are exponentially many trees). Incorporation of fidelity (including decoherence) in our DP approach requires non-trivial observation and analysis (§IV-B). Our Balanced-Tree Heuristic (§V) is closer to [18]'s work, in that both consider only balanced trees; however, we use a heuristic metric that facilitates a polynomial-time Dijkstra-like heuristic to select the optimal path, while their recursive metric<sup>6</sup> (albeit more accurate than

<sup>6</sup>We note that their formula (Eqn. 10 in [18]) is incorrect as it either ignores the 3/2 factor or assumes the EP generations to be synchronized across all links. In addition, their expression for "qubit age" ignores the "waiting for ES" time completely.

ours) is not amenable to an efficient (polynomial-time) search algorithm.

**Other Works.** In [44], Jiang et al. address a related problem; given a path with uniform link-lengths, they give an algorithm for selecting an optimal sequence of swapping and purification operations to produce an EP with fidelity constraints. In other recent works, Dahlberg et al [45] design physical and link layer protocols of a quantum network stack, and [46] proposes a data plane protocol to generate EPs within decoherence thresholds along a *given* routing path. More recently, Bugalho et al. [47] propose an algorithm to efficiently distribute multipartite entanglement across over than two nodes.

#### IV. OPTIMAL ALGORITHM FOR SINGLE TREE

In this section, we consider a special case of the QNR problem, viz., the case wherein there is a single source-destination  $(s, d)$  pair and the goal is to select a *single* swapping tree for the  $(s, d)$  pair. For this special case, we design an optimal algorithm based on dynamic programming. This optimal algorithm can be used iteratively to develop an efficient heuristic for the general QNR problem, as in §VI.

**QNR Single Path (QNR-SP) Problem.** Given a quantum network and a source-destination pair  $(s, d)$ , the QNR-SP problem is to determine a single swapping tree that maximizes the expected generation rate (i.e., minimizes the expected generation latency) of EPs over  $(s, d)$ , under the capacity and fidelity constraints.

For homogeneous nodes and link parameters, it is easy to see that the best swapping-tree is the balanced or almost-balanced tree over the shortest path. We note that QNR-SP is not a special case of QNR in the formal sense; e.g., the LP algorithm (§A) for QNR cannot be used for the QNR-SP problem, due to the single tree requirement (LP may produce multiple trees). As described in §III-B, the QNR-SP problem has been addressed before in [14], [18] under different models.

#### A. DYNAMIC PROGRAMMING (DP) FORMULATION

First, we note that a Dijkstra-like shortest path approach which builds a shortest-path tree greedily doesn't work for the QNR-SP problem—mainly, because the task is to find an optimal *tree* rather than an optimal path. As noted before, a routing path can have exponentially many swapping trees over it, with different generation latency. The recursive expression for computing the generation latency given in §III suggests that a dynamic programming (DP) approach, similar to the Bellman-Ford or Floyd-Warshall's classical algorithms for shortest paths, may be applicable for the QNR-SP problem. However, we need to “combine” trees rather than paths in the recursive step of a DP approach. Consequently, we were unable to design a DP approach based on the Floyd-Warshall's approach, but, are able to extend the Bellman-Ford approach for the QNR-SP problem after addressing a few challenges discussed below.

**DP Formulation.** We start with designing a DP algorithm without worrying about the decoherence constraint; we incorporate the decoherence constraint in the next subsection. Given a network, let  $T[i, j, h]$  be the optimal expected latency of generating EP pairs over  $(i, j)$  using a swapping tree of height at most  $h$ . Note that  $T[i, j, 0]$  for adjacent nodes  $(i, j)$  can be given by  $\frac{t_g}{p_g^2 p_e^2 p_{ob}}$ . Now, based on Eqn. (1), we start with the following equation for computing  $T[i, j, h]$  in terms of smaller-height swapping trees.

$$T[i, j, h] = \min (T[i, j, h - 1], (\frac{3}{2}B + t_c + t_b)/p_b) \quad (3)$$

where :

$$B = \min_{k \in V} \max (T[i, k, h - 1], T[k, j, h - 1])$$

However, there are three issues that need to be addressed before the above formulation can be turned into a viable algorithm. We address these in the below three paragraphs.

**(1) The 3/2 Factor; Throttled Trees.** As mentioned in §III, the 3/2 factor is an accurate estimate if the corresponding  $T$ 's are equal. However, in the above equation,  $T[i, k, h - 1]$  and  $T[k, j, h - 1]$  may not be equal. In our overall methodology, to conserve node and link resources, we post-process or “throttle” the swapping-tree obtained from the DP algorithm by increasing the generation latencies of some of the **non-root** nodes such that (i) the latencies of siblings are equalized, and (ii) the parents latency is related to the children's latency by Eqn. (1). We refer to this post-processing as *throttling*, and a tree that satisfies the above conditions as a *throttled* tree. Note that throttling does not alter the generation latency of the root and thus the overall tree; we prove the optimality of the overall algorithm formally in Theorem 2. Below, we motivate throttling, and describe how it is achieved.

**Justification.** In a given swapping tree, consider a pair of siblings  $x$  and  $y$  that have unequal generation latencies/rates. Let  $x$  be the one with a lower latency (higher rate). Then,  $x$  will likely have to discard many EPs while waiting for an EP from  $y$ . To minimize this discarding of EPs from  $x$  and to conserve underlying network resources so that they can be used in other swapping trees (in a general QNR solution), we “throttle”, increase (decrease) the generation latency (rate) of the sibling  $x$  to match that of  $y$ .

**Throttling Process.** Consider a pair of siblings  $x$  and  $y$  in the tree; let their parent be  $z$ . Let  $T_x, T_y$ , and  $T_z$  be their current generation latencies, such that  $T_z = (\frac{3}{2} \max(T_x, T_y) + t_c + t_b)/p_b$ . There are two potential steps: (i) If the parent's latency is to be kept unchanged, but  $T_x < T_y$  then  $T_x$  is increased to  $T_y$  which, thus, makes the above equation valid. (ii) If the parent's latency  $T_z$  is increased to  $T$  (by the above first step, with  $z$  as a sibling), then we increase the latencies of both  $x$  and  $y$  to  $2/3(Tp_b - t_c - t_b)$ . It is easy to see that applying the two steps iteratively from the root to the leaves, yields a throttled tree, as defined above.

**(2) Capacity Violation at Node  $k$ .** Note that the middle/common node  $k$  in Eqn. (3) may violate (node) capacity

constraints in the merged tree corresponding to  $T[i, k, h]$ , as it may use its full capacity in the trees corresponding to  $T[i, k, h - 1]$  and  $T[k, j, h - 1]$ . We address the above by adding two additional parameters to the sub-problems function  $T$ , corresponding to “usage percentage” of the end nodes. In particular, we define  $T[i, j, h, u_i, u_j]$  as the optimal latency of a swapping tree of height at most  $h$ , under the constraint that the end nodes  $i$  and  $j$  use at most  $u_i$  and  $u_j$  percentage of the respective node generation capacities; here,  $u_i$  and  $u_j$  can be positive integers between 1 and 100. The base case  $T[i, j, 0, u_i, u_j]$  for adjacent nodes  $(i, j)$  is given by  $\frac{t_g \min(u_i, u_j)}{p_g^2 p_e^2 p_{ob}}$ . Eqn. (3) is modified as follows to accommodate the additional usage parameters.

$$T[i, j, h, u_i, u_j] = \min (T[i, j, h - 1, u_i, u_j], \quad (4) \\ (\frac{3}{2}B + t_c + t_b)/p_b)$$

where :


$$B = \min_{k, u+u'=100} \max (T[i, k, h - 1, u_i, u], \\ T[k, j, h - 1, u', u_j])$$

(3) Ensuring Disjoint Subtrees. Note that Eqn. (3) implicitly assumes that the swapping trees corresponding to the latency values  $T[i, k, h - 1]$  and  $T[k, j, h - 1]$  are over disjoint paths, i.e., there is no node  $v$  such that both the paths contain  $v$ . If there is a common node  $v$ , then the combined tree corresponding to  $[i, j, h]$  may violate the node capacity constraints at  $v$ . This issue also arises in the classical Bellman Ford’s or Floyd-Warshall’s algorithms for shortest weighted paths, but is harmless with the assumption of positive-weighted cycles. We resolve the issue similarly here via the below lemma (see Appendix C for the proof).

**Lemma 1:** Consider two swapping trees  $\mathcal{T}_{ik}$  and  $\mathcal{T}_{kj}$  each of height at most  $h - 1$  over paths  $P_1 : i \rightsquigarrow v \rightsquigarrow k$  and  $P_2 : k \rightsquigarrow v \rightsquigarrow j$ , each of which contains a common node  $v \neq k$ . Then there exists two swapping trees  $\mathcal{T}_{iv}$  and  $\mathcal{T}_{vj}$  each of height at most  $h - 1$  over paths  $P'_1 : i \rightsquigarrow v$  and  $P'_2 : v \rightsquigarrow j$  such that: (i)  $P'_1$  is a subset of  $P_1$ , and  $P'_2$  is a subset of  $P_2$ , and (ii) generation latency of  $\mathcal{T}_{iv}$  is no greater than that of  $\mathcal{T}_{ik}$ , and generation latency of  $\mathcal{T}_{vj}$  is no greater than that of  $\mathcal{T}_{kj}$ . ■

Lemma 1 implies that if the swapping trees  $\mathcal{T}_{ik}$  and  $\mathcal{T}_{kj}$  corresponding to the latency values  $T[i, k, h - 1]$  and  $T[k, j, h - 1]$  have a common node, then there exist swapping trees of equal or better latency without any common nodes and these trees can be used to build a lower-latency tree over  $(i, j)$ .

**Overall DP Algorithm and Optimality.** Our DP-based algorithm for the QNR-SP problem for a given  $(s, d)$  pair is as follows. We use a DP formulation based on Eqn. (4) and the corresponding base case values to compute optimal generation latency  $T[s, d, h, 100, 100]$  and the corresponding swapping tree  $\mathcal{T}$ . Then, we throttle the tree  $\mathcal{T}$  as described in paragraph (1) above. The below theorem (see Appendix D for proof) states that the throttled tree thus obtained has the



**FIGURE 5.** Qubit parameters in a swapping tree used to compute the *age* of a qubit  $q$  at a leaf node  $l(q)$ . Here,  $l(q)$  is the left-most leaf of the subtree  $\mathcal{T}(q)$ .

optimal (minimum) expected generation latency among all throttled trees.

**Theorem 2:** The above described DP-based algorithm returns a throttled swapping tree over  $(s, d)$  with minimum expected generation latency (maximum expected generation rate) among all throttled trees over the given  $(s, d)$  pair. ■

## B. INCORPORATING FIDELITY CONSTRAINTS

Till now, we have ignored the fidelity constraints. We incorporate them in this section, by extending our DP formulation from the previous section. Limiting the decoherence, i.e., the qubit storage time, is challenging and is addressed first below. Limiting the number of leaves of a swapping tree is relatively easier, and is discussed next. We start with a definition.

**Definition 1:** (Qubit/Tree Age.) Given a swapping tree, the total time spent by a qubit in a swapping tree is the time spent from its “birth” via an atom-photon EP generation at a node till its consumption in a swapping operation or in generation of the tree’s root EP. We refer to this as a qubit’s *age*. The maximum age over all qubits in a swapping tree is called the tree’s (expected) *age*. □

**Estimating Qubit Age in a Swapping Tree.** Consider a throttled swapping tree  $\mathcal{T}$ , with a generation latency of  $T$ . Consider two siblings  $(A, B)$  and  $(B, C)$  at a depth<sup>7</sup> of  $i$  ( $i > 0$ ) from  $\mathcal{T}$ ’s root. If we ignore the  $t_c$  and  $t_b$  terms in (1), then the expected generation latency  $T(i)$  of both  $(A, B)$  and  $(B, C)$  being at depth  $i$  is given by:  $T(i) = \frac{T}{2}(\frac{2}{3}p_b)^i$ . Also, note that only *one* of the EPs  $(A, B)$  or  $(B, C)$  waits for  $T(i)$  time on an average. Thus, the expected waiting times for each of the four<sup>8</sup> qubits is  $T(i)/2$ .

Based on the above, we can now easily estimate the total waiting by a qubit  $q$  (referred to as  $q$ ’s *age*) before it is destroyed in a swapping operation. Let  $l(q)$  be the leaf, i.e. the link EP, of  $\mathcal{T}$  that contains the qubit  $q$ . Let  $\mathcal{T}(q)$  be the maximal subtree in  $\mathcal{T}$  such that  $l(q)$  is either its right-most or left-most leaf. Note that  $\mathcal{T}(q)$  is well defined for a tree  $\mathcal{T}$  and a qubit  $q$ . Let  $d(q)$  be the depth of the root of  $\mathcal{T}(q)$  in  $\mathcal{T}$ , and let  $d'(q)$  be the depth of  $l(q)$  in the subtree  $\mathcal{T}(q)$ . See Fig. 5. The expected age  $A(q)$  of  $q$  can be estimated as follows.

<sup>7</sup>Defined as the distance of a node from the root; depth of the root is 0.

<sup>8</sup>Note that qubit  $B$  in  $(A, B)$  is different from that in  $(B, C)$ .

Note that *age* of  $q$  is the total waiting by  $q$  at each of  $l(q)$ 's ancestors in  $\mathcal{T}(q)$ ; also note that at  $\mathcal{T}(q)$ 's root, the qubit  $q$  is destroyed, and hence,  $q$  does not age at any ancestor of  $T(q)$ 's root. It is easy to see that the expected age  $A(q)$  is:

$$A(q) = \left( \sum_{i=d(q)}^{d(q)+d'(q)} T(i)/2 \right) + (t_{ob} + t_p)$$

Above, the last term is the time spent by  $q$  waiting for its link EP to be established and is given by sum of optical-BSM ( $t_{ob}$ ) and photon transmission latency ( $t_p$ ). Note that the actual age of a qubit  $q$  is some distribution with the above mean. We observe the following.

*Observation 2:* Given a swapping tree  $\mathcal{T}$ , let  $\mathcal{T}_l$  and  $\mathcal{T}_r$  be its left and right children. If the atomic BSM probability  $p_b$  is  $\leq 75\%$ , then the expected age of the right-most or left-most descendant of either  $\mathcal{T}_l$  or  $\mathcal{T}_r$  is greater than the expected age of any other qubit in the tree.  $\square$

**DP Formulation with Decoherence/Age Constraint.** If we assume the atomic BSM probability  $p_b \leq 75\%$ , then we can design a DP algorithm for the QNR-SP problem with the decoherence constraint, as follows. Let  $T[i, j, h, h_{ll}, h_{lr}, h_{rl}, h_{rr}, u_i, u_j]$  be the optimal latency from a swapping tree of height at most  $h$ , whose root's left (right) child's left-most and right-most descendants are at depths of (exactly)  $h_{ll}$  and  $h_{lr}$  ( $h_{rl}$  and  $h_{rr}$ ), each of which is upper bounded by  $h$ . Here,  $u_i$  and  $u_j$  parameters are as before. Note that  $T[i, j, 1, 0, 0, 0, 0, u_i, u_j] = \frac{t_g \min(u_i, u_j)}{p_g^2 p_e^2 p_{ob}}$ . We have:

$$\begin{aligned} T[i, j, h, h_{ll}, h_{lr}, h_{rl}, h_{rr}, u_i, u_j] \\ = \min (T[i, j, h-1, h_{ll}, h_{lr}, h_{rl}, h_{rr}, u_i, u_j], \\ (\frac{3}{2}B + t_c + t_b)/p_b) \end{aligned} \quad (5)$$

where :

$$\begin{aligned} B = \min_{k, g'_i s, u+u'=100} \max (T[i, k, h-1, h_{ll}-1, g_1, g_2 \\ h_{lr}-1, u_i, u], \\ T[k, j, h-1, h_{rl}-1, g_3, g_4 \\ h_{rr}-1, u', u_j]) \end{aligned}$$

The above formulation will give us the optimal latency swapping-tree for each combination of  $(h_{ll}, h_{lr}, h_{rl}, h_{rr})$ . We remove the trees that violate the decoherence constraint, and pick the minimum-latency tree from the remaining. This gives us a swapping tree with optimal latency under the decoherence constraint. The proof of optimality easily follows.

**Constraint on Number of Leaves.** Limiting the number of leaves to  $\tau_l$  can be easily done by adding another parameter for number of leaves in the  $T$  array/function above. This adds another factor of  $O(n^2)$  to the time complexity, as we need to check for all combination of number of leaves in the two subtrees. To optimize, we can now replace the height parameter, but keeping the height parameters aids in parallelism, as described below.

**Time Complexity; DP-OPT and DP-Approx Algorithms**  
Note that, in (5) above, we can pre-compute  $\min_{g_1, g_2} T[\dots, g_1, g_2, \dots]$  and similarly  $\min_{g_3, g_4} T[\dots, g_3, g_4, \dots]$  before computing  $B$ . With this, the time complexity of the DP formulation becomes  $O(n^9)$ , which can be further reduced  $O(n^5(\log n)^4)$  if we assume height of a tree to be at most  $(c \log n)$  for some constant  $c$ . For a real-time routing application, the above time complexity is still high—as the algorithm can take a few minutes on a single core. However, as the algorithm lends to obvious parallelism, it can be executed in as little as  $O((\log n)^2)$  time with sufficiently many cores, using the height parameter sequentially. We can also reduce the sequential time complexity to  $O(n^5)$ , by approximating the maximum qubit's age in a tree to the generation latency of the tree, which is at most  $3/(2p_b)$  the actual value. Note that maximum age of a qubit is at least  $2Tp_b/3$  and at most  $T$ , where  $T$  is the generation latency of the tree. Finally, we can make the algorithm more efficient by assuming the usage parameter values to be 50%.<sup>9</sup> We refer to the  $O(n^5)$  algorithm with the above assumptions as DP-Approx, and the  $O(n^5(\log n)^4)$  algorithm based on (5) as DP-OPT. Both algorithms use throttling after the DP formulation.

## V. Balanced-Tree HEURISTIC FOR QNR-SP

The DP-based algorithms presented in §IV for the QNR-SP problem have high time complexity, and thus, may not be practical for real-time route finding in large networks. In this section, we develop an almost-linear time heuristic for the QNR-SP problem, based on the classic Dijkstra shortest path algorithm; the designed heuristic performs close to the DP-based algorithms in our empirical studies.

**Basic Idea.** The main reason for the high-complexity of our DP-based algorithms in §IV is that the goal of the QNR-SP problem is to select an optimal swapping *tree* rather than a path. One way to circumvent this challenge efficiently while still selecting near-optimal swapping tree, is to restrict ourselves to only "balanced" swapping trees. This restriction allows us to think in terms of selection of paths—rather than trees—since each path has a unique<sup>10</sup> balanced swapping tree. We can then develop an appropriate path metric based on above, and design a Dijkstra-like algorithm to select an  $(s, d)$  path that has the optimal metric value. We note that Caleffi [18] also proposed a path metric based on balanced swapping trees, but their metric, though accurate, only had a recursive formulation without a closed-form expression—and hence, was ultimately not useful in designing an efficient algorithm. In contrast, we develop an approximate metric with a closed-form expression, based on the "bottleneck" link, as follows.

**Path Metric  $M$ .** Consider a path  $P = (s, x_1, x_2, \dots, x_n, d)$  from  $s$  to  $d$ , with links  $(s, x_1), (x_1, x_2), \dots, (x_n, d)$  with

<sup>9</sup>This also enforces  $s$  and  $d$  to use only 50% capacity; this can be resolved by doubling  $s$  and  $d$  capacity a priori.

<sup>10</sup>In fact, there can be multiple balanced trees over a path whose length is not a power of 2, but, since they differ minimally in our context, we can pick a unique way of constructing a balanced tree over a path.

given EP latencies. We define the path metric for path  $P$ ,  $M(P)$ , as the EP generation latency of a balanced swapping over  $P$ , which can be estimated as follows. Let  $L$  be the link in  $P$  with maximum generation latency. If  $L$ 's depth (distance from the root) is the maximum in a throttled swapping tree, then we can easily determine the accurate generation latency of the tree. However, in general,  $L$  may not have the maximum depth, in which case we can still estimate the tree's latency approximately, if the tree is balanced, as follows. In balanced swapping trees, *assuming* the maximum latency link  $L$  to be at the maximum depth, gives us a constant-factor approximation of the tree's generation latency. Thus, let us assume  $L$  to be at the maximum depth of a balanced tree over  $P$ ; this maximum depth is  $d = \lceil(\log_2 |P|)\rceil$ . Let the generation latency of  $L$  be  $T_L$ . If we ignore the  $t_b + t_c$  term in Eqn. (1), then, the generation latency of a throttled swapping tree can be easily estimated to  $T(\frac{3}{2p_b})^d$ . The term  $t_b + t_c$  can also be incorporated as follows. Let  $T(i)$  denote the expected latency of the ancestor of  $L$  at a distance  $i$  from  $L$ . Then, we get the recursive equation:  $T(i) = (\frac{3}{2}T(i-1) + t_b + t_c)/p_b$ . Then, the path metric value  $M(P)$  for path  $P$  is given by  $T(d)$ , the generation latency of the tree's root at a distance of  $d$  from  $L$ , and is equal to:

$$M(P) = T(d) = \bar{p}^d T_L + [(\bar{p} - 1)/(\bar{p}^d - 1)](t_b + t_c)/p_b$$

where  $\bar{p} = 3/(2p_b)$  and  $d = \lceil(\log_2 |P|)\rceil$ . The above is a  $(1+3/(2p_b))$ -factor approximation latency of a balanced and throttled swapping tree over  $P$ ; this can be shown easily using analysis from §IV-B.

**Optimal Balanced-Tree Selection.** The above path-metric  $M()$  is a monotonically increasing function over paths, i.e., if a path  $P_1$  is a sub-sequence of another path  $P_2$ , then  $M(P_1) \leq M(P_2)$ . Thus, we can tailor the classical Dijkstra's shortest path algorithm to select a  $(s, d)$  path with minimum  $M(P)$  value, using the link's EP generation latencies as their weights. We refer to this algorithm as *Balanced-Tree*, and it can be implemented with a time complexity of  $O(m + n \log n)$  using Fibonacci heaps, where  $m$  is the number of edges and  $n$  is the number nodes in the network.

**Incorporating Fidelity Constraints.** Fidelity constraints in our path-metric based setting can be handled by essentially computing the optimal path for each path-length (number of hops in the path) up to  $\tau_l$ , and then pick the best path among them that satisfies the fidelity constraints. This obviously limits the number of leaves to  $\tau_l$  and addresses the operations-based fidelity degradation. The above also address the decoherence/age constraint, since it is easy to see (from analysis in §IV-B) that the age of a balanced swapping tree can be very closely approximated in terms of the latency and the number of leaves. Now, to compute the optimal path for each path-length, we can use a simple dynamic programming approach that run in  $O(m\tau_l)$  time where  $m$  is the number of edges and  $\tau_l$  is the constraint on number of leaves.

## VI. ITER: ITERATIVE QNR HEURISTIC

The general QNR problem can be formulated in terms of hypergraph flows and solved using LP (see Appendix A). Although polynomial-time and provably optimal, the LP-based approach has a very high time-complexity for it to be practically useful. Here, we develop an efficient heuristic for the QNR problem by iteratively using an QNR-SP algorithm.

**ITER Heuristic.** To solve the QNR problem efficiently, we apply the efficient DP-Approx algorithm iteratively—finding an efficient swapping tree in each iteration for one of the  $(s_i, d_i)$  pairs. The proposed algorithm is similar to the classical Ford-Fulkerson augmenting path algorithm for the max network flow problem at a high level, with some low level and theoretical differences as discussed below. The iterative-DP-Approx algorithm for the QNR problem consists of the following steps:


- 1) Given a network, we compute maximum EP generation rates for each network link using Eqn. (6). Use these as weights on the link.
- 2) For each  $(s_i, d_i)$  pair, use DP-Approx algorithm to find the optimal path  $P_i$ , under the capacity and fidelity constraints. Consider a *throttled* and balanced swapping tree  $\mathcal{T}_i$  over  $P_i$ . Let  $\mathcal{T}^*$  be the swapping tree with highest generation rate; if this rate is below a certain threshold, then quit.
- 3) Construct a residual network graph by subtracting the resources used by  $\mathcal{T}^*$ , using Eqn. (7).
- 4) Go to step (1).

Before we present the expressions required above, we would like to point out key differences of our context with the classic network flow setting. Even though we are augmenting our solution one *path* at a time, the network resources are fundamentally being used by swapping trees created over these paths. These path-flows don't really have a direction of flow, but we can assign them a symbolic direction from source to the direction. Even with these symbolic directions, the flows in opposite directions over any edge  $k$  do not "cancel" each other as in the classical network flow. Moreover, flow conservation law doesn't hold in our context (e.g., even a path may not use same link rates on all links, due to them being at different depths of the tree), and thus, the max-flow min-cut theorem doesn't hold. Thus, ITER may not give an optimal solution, even for a single  $(s, d)$  pair.

**Link EP Generation Rate/Latency.** Consider a pair of network node  $i$  and  $j$  with corresponding *current (residual)* values of node latencies as  $t_g(i)$  and  $t_g(j)$ . Assuming  $p_g$  values to be same for both nodes, the minimum EP link rate for  $(i, j)$  is then given by

$$\min(1/t_g(i), 1/t_g(j))p_g^2 p_e^2 p_{ob}. \quad (6)$$

**Residual Node Capacities.** Let  $P$  be a path added by ITER, at some earlier stage, and let  $\mathcal{T}$  be the corresponding throttled swapping tree over  $P$ . As in §III, let  $R(e, \mathcal{T})$  be the EP generation rate being used by  $\mathcal{T}$  over a link  $e \in P$ ,



**FIGURE 6.** Swapping Tree Protocol Illustration. The shown tree is not a swapping-tree, but rather a certain hierarchy of nodes to illustrate the BSM operation in the swapping-tree protocol. A link-layer protocol continuously generates EPs over links  $(x_0, x_2)$  and  $(x_2, x_4)$ . On receiving EP on links on either side,  $x_1$  ( $x_3$ ) attempts a BSM operation on the stored qubit atoms. If the BSM succeeds,  $x_1$  ( $x_3$ ) sends two classical bits (solid green arrows) to  $x_2$  ( $x_4$ ) for desired manipulation/correction after which  $x_2$  ( $x_4$ ) sends an ACK (dashed green arrows) to the other end-node  $x_0$  ( $x_2$ ) to complete the EP generation. If BSM at  $x_1$  and  $x_3$  are both successful, then  $x_2$  attempts the BSM as above. If a BSM at say  $x_1$  fails, that  $x_1$  failure signals (red arrows) to all the descendant nodes of the subtree rooted at  $x_1$  so that they can start accept new EPs from the link layer protocol. Note that here node  $x_2$  plays multiple roles and hence appears at multiple places in the figure.

$R_e = \sum_{\mathcal{T}} R(e, \mathcal{T})$ , and  $E(i)$  be the edges incident on  $i$ . Then, the residual node rates can be calculated similar to Eqn. (2) as follows. Below,  $t_g'(i)$  is the *original* value.

$$1/t_g(i) = 1/t_g'(i) - \sum_{e \in E(i)} R_e / (p_g^2 p_e^2 p_{ob}) \quad \forall i \in V. \quad (7)$$

The residual memory capacity is easy to compute—each path/tree uses 2 memory units for each intermediate node, and 1 memory unit for the end nodes.

## VII. EVALUATIONS

The goal of our evaluations is to compare the EP generation rates, evaluate the fidelity of generated EPs, and validate our analytical models. We implement the various schemes over a discrete event simulator for QNs called NetSquid [48]. The NetSquid simulator accurately models various QN components/aspects, and in particular, we are able to define various QN components and simulate swapping-trees protocols by by implementing gate operations in entanglement swapping.

**Swapping Tree Protocol.** Our algorithms compute swapping tree(s), and we need a way to implement them on a network. We build our protocol on top of the link-layer of [45], which is delegated with the task of continuously generating EPs on a link at a desired rate (as per the swapping tree specifications). Note that a link  $(a, b)$  may be in multiple swapping trees, and hence, may need to handle multiple link-layer requests at the same time; we implement such link-layer requests by creating independent atom-photon generators at  $a$  and  $b$ , with one pair of synchronized generators for each link-layer

request. As the links generate continuous EPs at desired rates, we need a protocol to swap the EPs. Omitting the tedious bookkeeping details, the key aspect of the protocol is that swap operation is done only when both the appropriate EP pairs have arrived. We implement all the gate operations (including, atomic and optical BSMs) within NetSquid to keep track of the fidelity of the qubits. On BSM success, the swapping node transmits classical bits to the end node which manipulates its qubit, and send the final ack to the other end node. On BSM failure, a classical ack is send to all descendant link leaves, so that they can now start accepting new link EPs; note that in our protocol, a link  $l$  does not accept any more EPs, while its ancestor is waiting for its sibling's EP. See Fig. 6

**Simulation Setting.** We use a similar setting as in the recent work [14]. By default, we use a network spread over an area of  $100\text{km} \times 100\text{km}$ . We use the Waxman model [49], used to create Internet topologies, to randomly distribute the nodes and create links; we use the maximum link distance to be 10km. We vary the number of nodes from 25 to 500, with 100 as the default value. We choose the two parameters in the Waxman model to maintain the number of links to 3% of the complete graph (to ensure an average degree of 3 to 15 nodes). For the QNR-SP problem, we pick  $(s, d)$  pairs within a certain range of distance, with the default being 30-40 kms; for the QNR problem, we extend this range to 10-70 kms.

**Parameter Values.** We use parameter values mostly similar to the ones used in [18] corresponding to a single-atom based quantum memory platform, and vary some of them. In particular, we use atomic-BSM probability of success ( $p_b$ ) to be 0.4 and latency ( $t_b$ ) to be  $10 \mu\text{secs}$ ; in some plots, we vary  $p_b$  from 0.2 to 0.6. The optical-BSM probability of success ( $p_{ob}$ ) is half of  $p_b$ . We use atom-photon generation times ( $t_g$ ) and probability of success ( $p_g$ ) as  $50 \mu\text{sec}$  and 0.33 respectively. Finally, we use photon transmission success probability as  $e^{-d/(2L)}$  [18] where  $L$  is the channel attenuation length (chosen as 20km for an optical fiber) and  $d$  is the distance between the nodes. Each node's memory size is randomly chosen within a range of 15 to 20 units. Fidelity is modeled in NetSquid using two parameter values, viz., depolarization (for decoherence) and dephasing (for operations-driven) rates. We choose a decoherence time of two seconds based on achievable values with single-atom memory platforms [50]; note that decoherence times of even several minutes [37], [38] to hours [39], [40] has been demonstrated for other applicable memory platforms. Accordingly, we choose a depolarization rate of 0.01 such that the fidelity after a second is 90%. Similarly, we choose a dephasing rate of 1000 which corresponds to a link EP fidelity of 99.5% [15].

**Algorithms and Performance Metrics.** To compare our techniques with prior approaches, we implement most recently proposed approaches, viz., (i) the WaitLess-based linear programming (LP) approach from [15] (called Delft-LP here), (ii) Q-Cast approach from [14] which is WaitLess-based but uses multiple links and requires




FIGURE 7. QNR-SP Problem: EP Generation Rates for varying parameters.




FIGURE 8. QNR Problem: EP Generation Rates for varying parameters.




FIGURE 9. Compare the performance with Caleffi in a (a) low density network and a (b) high density network.


memories. The Waiting-based algorithm by Caleffi [18] uses an exponential-time approach, and is thus compared only for small networks. The [16] and [17] approaches are not compared as they were found to be inferior to Q-Cast.

For all algorithm except for Q-Cast, we use only one link between adjacent nodes, since only Q-Cast takes advantage of multiple links in a creative way. In particular, for Q-Cast, we use  $W = 1, 5$ , or  $10$  sub-links ([14] calls them channels) on each link, with the node and link "capacity" divided equally among the them. We note that in Q-Cast each node requires  $2W$  memories (2 for each sub-link) with sufficient coherence time to allow for the entire swapping operation over the path to be completed. The Delft-LP approach explicitly assumes the generation of link EPs is deterministic, i.e., the value  $p_g^2 p_e^2 p_{ob}$  is 1, and does not model node generation rates. We address these differences by extending their LP formulation: (i) We add a constraint on


node generation rates, and (ii) add a  $p_g^2 p_e(i, j)^2 p_{ob}$  factor to each link  $(i, j)$  in any path extracted from their LP solution.

Among our schemes, we use DP-OPT, DP-Approx and Balanced-Tree (see §IV-B) for the QNR-SP problem, and LP (Appendix A) and ITER schemes for the QNR problem. For ITER, we use three schemes: ITER-DPA, ITER-Bal and ITER-SP, which iterate over DP-Approx, Balanced-Tree and SP respectively. To be comprehensive, we also implement a simple SP algorithm which picks a balanced swapping tree over the shortest path (minimum number of links). We compare the schemes largely in terms of EP generation rates; we also compare the execution times and EPs fidelity.

**Comparison with [18] for QNR-SP Problem.** Note that [18] gives only an QNR-SP algorithm referred to as Caleffi; it takes exponential-time making it infeasible to run for network sizes much larger than 15-20. In particular, for network sizes 17-20, it takes several hours, and our preliminary analysis suggests that it will take of the order of  $10^{40}$  years on our 100-node network. See Appendix E. Thus, we use a small network of 15 nodes over a  $25\text{km} \times 25\text{km}$  area; we consider average node degrees of 3 or 6. See Fig.9. We see that DP-OPT outperforms Caleffi by 10% on a average for the sparser graph and minimally for the denser graph. However, for some instances, DP-OPT outperformed Caleffi by as much as 300% (see Appendix F). We see that DP-Approx performs similar to DP-OPT, while Balanced-Tree is outperformed slightly by Caleffi; however, for this small network, since the DP-OPT and DP-Approx algorithms only take 10-100s of msec (Appendix E), Balanced-Tree need not be used in practice.



**FIGURE 10.** EP generation over linear paths. (a) EP Rates, and (b) Fidelity, over linear paths with varying link lengths. (c) Maximum reachable distance with links of 30-35m lengths. (d) EP generation rates over linear paths with 10-50 km links to demonstrate impact of varying link lengths.



**FIGURE 11.** Comparing with Analytical Results. (a) Analytical vs. Simulation Results. (b) Throttled vs. Non-Throttled Trees (QNR-SP). (c) Throttled vs. Non-Throttled Trees (QNR). (d) Fairness Measure.

**QNR-SP Problem (Single Tree) Results.** We start with comparing various schemes for the QNR-SP problem, in terms of EP generation rate. We compare DP-Approx, DP-OPT, Balanced-Tree, SP, and Q-Cast; note that the LP schemes can't be used to select a *single* tree, as they turn into ILPs. See Fig. 7, where we plot the EP generation rate for various schemes for varying number of nodes,  $(s, d)$  distance,  $p_b$ , and network link density. We observe that DP-Approx and DP-OPT perform very closely, with the Balanced-Tree heuristic performing close to them; all these three schemes outperform the Q-Cast schemes (for  $W = 5, 10$  sub-links) by an order of magnitude. We don't plot Q-Cast for  $W = 1$  sub-links, as it performs much worse (less than  $10^{-3}$  EP/sec). We note that Q-Cast's EP rates here are much lower than the ones published in [14], because [14] uses link EP success probability of 0.1 or more, while in our more realistic model, the link EP success probability is  $p_g^2 p_e^2 p_{ob} = 0.012$  for the default  $p_b$  value. We reiterate that our schemes require only 2 memory units per node, while the Q-Cast schemes requires  $2W$  units. The main reason for poor performance of Q-Cast (in spite of higher memory and link synchronization) is that, in the WaitLess model, the EP generation over a path is a very low probability event—essentially  $p^l$  where  $p$  is the link-EP success probability and  $l$  is the path length, for the case of  $W = 1$  (the analysis for higher  $W$ 's is involved [14]). Finally, our proposed techniques also outperform the SP algorithm, especially when the number of possible paths (trees) between  $(s, d)$  pair increases. In addition, we see that

performance increases with increase in  $p_b$ , number of nodes, or network link density, as expected due to availability of better trees/paths; it also increases with decrease in  $(s, d)$  distance as fewer hops are needed.

**QNR Problem Results.** We now present performance comparison of various schemes for the QNR problem. Here, we compare the following schemes: ITER-DPA, ITER-Bal, ITER-SP, Delft-LP, and Q-Cast with the optimal LP as the benchmark for comparison (LP wasn't feasible to run for more than 100 nodes). See Fig. 8. Our observations are similar to that for the QNR-SP problem results. We see that in all plots, LP being optimal performs the best, but is closely matched by ITER-DPA and the efficient heuristic ITER-Bal. We observe that the performance gap between our proposed techniques and ITER-SP is higher than in the QNR-SP case, as SP picks paths based on just number of links. Our schemes outperform both Delft-LP and Q-Cast by an order of magnitude, for the same reason as mentioned above.

**Fidelity and Long-Distance Entanglements.** We now investigate the fidelity of the EPs generated. First, we note that the Q-Cast and Delft-LP schemes will incur near-zero decoherence as they involve only transient storage. Decoherence for other schemes is also negligible as the EP generation latencies (10s of msecs) is much less than the coherence time. The operations-driven fidelity loss is expected to be similar for all schemes, as they all roughly use the same order of links. Overall, we observed fidelities of 94-97% across

all schemes (not shown), with our schemes also performing better sometimes due to smaller number of leaves.

**Long Path Graphs.** To test the limits of the schemes in terms of decoherence and fidelity, we consider a long path network and estimate fidelity of EPs generated by schemes for increasing distances and link-lengths (link success probability decreases with increasing link length). Fig. 10(a)-(b) shows EP generation rates and fidelity for path lengths of 500km and 1000km for varying link lengths, for the single-tree schemes DP-Approx and Balanced-Tree. Q-Cast and Delft-LP are not shown as their EP rate is near-zero ( $\leq 10^{-20}$ ) at these distances. We observe that our schemes yield EPs with qubit fidelities of 65-82% and 40-64% for 500km and 1000km paths respectively, with EP rates of 0.05 to 0.65/sec. These are viable results—since qubit copies with fidelities higher than 50% can be purified to smaller copies with arbitrarily higher fidelities [51], [52].

Now, in Fig. 10(c), we demonstrate the effect of decoherence time of quantum memories used in nodes. Here, we use 30-35 km links. We see that even with decoherence time of as low as 100 ms, DP-Approx is able to create EPs for up to 200 kms while Balanced-Tree can only create EP for paths up to 120 kms; they perform similarly for larger decoherence times. As all the links are almost of the same length, the optimal swapping will be largely-balanced trees wherein the EP generation rate depends only on the tree height. Due to this reason, the maximum achievable path-length graph is close to a step function. We add that our schemes produce 0.008 EPs/s for distance of more than 4000 kms.

Finally, in Fig. 10(d), we demonstrate the higher performance of non-balanced trees when the links on a path may have much different lengths. In particular, we pick link lengths randomly in the range of 10 to 50 kms. With this setting, we see that DP-Approx performs much better than Balanced-Tree, and in some cases, up to 100% better. Note that, Balanced-Tree and Caleffi have similar performance over linear graphs, as there is no path selection scheme needed.

**Validating the Analysis; Fairness.** Fig. 11(a) compares the EP generation rates as measured by the analytical formulae and the actual simulations for the QNR-SP algorithms DP-Approx and Balanced-Tree. We observe that they match closely, validating our assumption of 3/2 factor in Eqn. (1) and of exponential distributions at higher levels of the tree, and of the path metric  $M()$  for Balanced-Tree. Fig. 11(b)-(c) plots the EP generation rates for throttled and non-throttled trees. We see that the throttled tree underperforms the non-throttled tree by only a small margin for the single-tree case; however, for the multi-tree ITER-Bal algorithm, the throttled trees perform better as they are able to use the resources efficiently. Fig. 11(d) plots the average number of  $(s, d)$  pairs that get at least one tree/path for varying number of requests; we see that our schemes exhibit 90-99% fairness.

**Execution Times.** We ran our simulations on an Intel i7-8700 CPU machine, and observed that the WaitLess algorithms as well our Balanced-Tree and ITER-Bal heuristics run in fraction of a second even for a 500-node network; thus, they can be used in real-time. Note that since our problems depend on real-time network state (residual capacities), the algorithms must run very fast. The other algorithms (viz., DP-OPT, DP-Approx, and ITER-DPA) can take minutes to hours on large networks, and hence, may be impractical on large network without significant optimization and/or parallelization. See Appendix G for the plot.

## VIII. CONCLUSIONS

We have designed techniques for efficient generation of EP to facilitate quantum network communication, by selecting efficient swapping trees in a Waiting protocol. By extensive simulations, we demonstrated the effectiveness of our techniques, and their viability in generating high-fidelity EP over long distances (500-1000km). Our future work is focused on exploring more sophisticated generation structures, e.g., aggregated trees, taking advantage of pipelining across rounds, incorporating purification techniques, and to extend our techniques to multi-mode memories [53], [54].

## ACKNOWLEDGMENT

This work was supported by NSF awards FET-2106447 and CNS-2128187, and a Cisco industry grant.

## APPENDIX A LP FORMULATION FOR THE QNR PROBLEM

In this Appendix we provide an optimal LP-based solution to the QNR problem. Although polynomial-time, this solution has high complexity, so its main use is as a benchmark in evaluating the more efficient (but possibly sub-optimal) algorithms for the problem.

Our approach follows from the observation each swapping tree in a QN can be viewed as a special kind of path (called  $B$ -hyperpath [55]) over a hypergraph constructed from the network graph. We begin by describing the hypergraph construction for the single-pair case and ignoring fidelity constraints. We then extend traditional hypergraph-flow algorithm to incorporate losses (e.g., due to BSM failures), stochasticity, and the interaction between memory constraints and stochasticity. Finally, we extend the formulation to multiple  $(s, d)$  pairs and incorporate fidelity constraints.

Optimal generation of long-distance entanglement was posed as an LP problem in [56], but differs from our more general formulation work in three main ways. First of all, [56] assumes unbounded memory capacity at each swapping node to queue up incoming EPs. In contrast, our model has bounded memory capacity at each node, and consequently, our LP formulation deals with *expectations* over rates/latencies rather than scalar rate values. Secondly, our formulation accounts for node capacity constraints in addition to link constraints. Thirdly, our formulation poses




FIGURE 12. ST-hypergraph for a 4-node linear network. Not all *prod* nodes are shown.

the problem in terms of hypergraph flows, which permits us to easily incorporate fidelity and decoherence constraints.

#### A. HYPERGRAPH-BASED REPRESENTATION OF ENTANGLEMENT GENERATION

We begin by recalling standard hypergraph notions [55], [57], [58].

**Definition 2:** (Hypergraph) A directed hypergraph  $H = (V(H), E(H))$  has a set of vertices  $V(H)$  and a set of (directed) hyperarcs  $E(H)$ , where each hyperarc  $e$  is a pair  $(t(e), h(e))$  of non-empty disjoint subsets of  $V(H)$ . A weighted hypergraph is additionally equipped with a weight function  $\omega : E(H) \rightarrow R^+$ .  $\square$

Sets  $t(e)$  and  $h(e)$  are called the *tail* and *head*, resp., of hyperarc  $(t(e), h(e))$ . A hyperarc  $e$  is a *trivial* edge if both  $t(e)$  and  $h(e)$  are singleton; and *non-trivial* otherwise. A hyperarc  $e$  where  $|h(e)| = 1$ , i.e. whose head is singleton, is called a *B-arc*. A hypergraph consisting only of *B-arcs* is called a *B-hypergraph*.

**Definition 3:** (Connectivity and *B*-Hyperpaths) A vertex  $t$  is *B-connected* to vertex  $s$  in hypergraph  $H$  if  $s = t$  or there is a hyperarc  $e \in E(H)$  such that  $h(e) = \{t\}$  and every  $v \in t(e)$  is *B-connected* to  $s$  in  $H$ . A *B*-hyperpath from  $s$  to  $t$  is a minimal *B*-hypergraph  $P$  such that  $V(P) \subseteq V(H)$ ,  $E(P) \subseteq E(H)$ , and  $t$  is *B-connected* to  $s$  in  $P$ .  $\square$

**ST-hypergraph.** Given a QN and single  $(s, d)$  pair, we first construct a hypergraph that represents the set of all possible swapping trees rooted at  $(s, d)$ . Given a QN represented as an undirected graph  $G = (V, E)$  and a single  $(s, d)$  pair, its *ST-hypergraph* is a hypergraph  $H$  constructed as follows (see Fig. 12). All pairs of vertices below are unordered pairs.

- $V(H)$  consists of:
  - 1) Two distinguished vertices *start* and *term*
  - 2)  $prod(u, v)$  and  $avail(u, v)$  for all distinct  $u, v \in V$
- $E(H)$  consists of 5 types of hyperarcs:
  - 1) [Start]  $e = (\{start\}, \{avail(u, v)\}) \forall u, v \in V$ .

- 2) [Swap]  $e = (\{avail(u, w), avail(w, v)\}, \{prod(u, v)\})$ , for all distinct  $u, v, w \in V$ .
- 3) [Prod]  $e = (\{prod(u, v)\}, \{avail(u, v)\}) \forall u, v \in V$ .
- 4) [Term]  $e = (\{avail(s, d)\}, \{term\})$ .

In an ST-hypergraph, vertices *start* and *term* represent source and sink nodes of a desired hypergraph-flow (see below). Other vertices represent EPs over a pair of nodes in  $G$ . Hyperarcs represent how the tail EPs contribute to that at the head. For ease of accounting, we categorize generated EPs using different types of vertices: *start* represents link-level EPs generated over links in  $G$ , *prod* represent EPs produced by atomic entanglement-swapping, and *avail* represent EPs generated from either of the above. “Start” and “Prod” arcs turn the *start* and *prod* EPs respectively into *avail* EPs and thus make them available for further swapping. “Swap” arcs represent swapping over the triplets of nodes  $(u, w, v)$ . Note that an ST-hypergraph is a *B*-hypergraph, as “Swaps” are the only non-trivial hyperarcs, and their head is singleton.

**Swapping Trees as *B*-Hyperpaths.** Given a QNR problem with a single pair  $(s, d)$ , it is easy to see that any swapping tree generating  $(s, d)$  EPs can be represented by a unique *B*-hyperpath from *start* to *term* in the above ST-hypergraph. Thus, it easily follows that a QNR problem of selection of (multiple) swapping trees is equivalent to finding an optimal hypergraph flow from *start* to *term* in  $H$ . Note that  $H$  has  $O(|V|^2)$  vertices and  $O(|V^3|)$  hyperarcs.

#### B. ENTANGLEMENT FLOW AS $\text{LP}$

We now develop an LP formulation to represent the QNR problem over  $(s, d)$  in  $G$  as a hypergraph-flow problem in  $H$ . In contrast to the classic hypergraph-flow formulation [55], we need to consider *lossy* flow, with loss arising from two sources: (i) ES operations have a given success probability, and (ii) waiting for both qubits to arrive before performing ES leads to losses since the arrival of EPs follow independent probability distributions. For the latter, we make use of Observation 1. The proposed LP formulation is as follows.

- **Variables:**  $z_a$ , for each hyperarc  $a$  in  $H$ , represents the EP generation rate over each of the (one or two) node-pairs in  $a$ 's tail. This enforces the condition that EP rates over the two node-pairs in *prod* hyperarc's tail are equal. Thus, the  $\text{LP}$  solution will result in *throttled* swapping trees.
- **Capacity Constraints:**  $z_a \in R^+$  for all hyperarcs  $a$  in  $H$ . We use Eqn. (2) to add the following constraints due to nodes in  $G$ .

$$1/t_g \geq \sum_{x \in E(i)} z_{a(x)} / (p_g^2 p_e^2 p_{ob}) \quad \forall i \in V.$$

Above  $a(x)$  is the hyperarc in  $H$  of the form  $(start, avail(x))$  where  $x$  is an edge in  $G$ .

- **Flow Constraints** which vary with vertex types. Below, we use notations  $out(v)$  and  $in(v)$  to represent outgoing and incoming hyperarcs from  $v$ . Formally,  $out(v)$  is

$\{a \in E(H) : v \in t(a)\}$  and  $in(v)$  is  $\{a \in E(H) : v \in h(a)\}$ .

– For each vertex  $v$  s.t.  $v = avail(\cdot)$ :

$$\sum_{a \in in(v)} z_a = \sum_{a' \in out(v)} z_{a'}$$

That is, there is no loss in making already generated entanglements available for further swapping.

– For each vertex  $v$  s.t.  $v = prod(\cdot)$ :

$$\sum_{a \in in(v)} z_a p_b(2/3) = \sum_{a' \in out(v)} z_{a'}$$

The  $(2/3)p_b$  factor follows from Observation 1, and accounts for loss due to swapping failures as well as due to waiting for arrival of both EPs for swapping.

• **Objective:** Maximize  $\sum_{a \in in(term)} z_a$

**Multiple-Pairs Multi-Path:** The above LP formulation for the single-pair QNR problem can be readily extended to the multiple-pairs case. Let  $\{(s_1, d_1), (s_2, d_2), \dots, (s_n, d_n)\}$  be a set of source-destination pairs. The *only change* is that the hypergraph  $H$  now has  $n$  arcs ( $\{avail(s_i, d_i)\}, \{term\}$ ) for all  $i$ . The other arcs model the generation of EPRs independent of the pairs, and thus are unchanged. It is interesting to note that the multi-pairs problem, typically formulated as multi-commodity flow in classical networks, is posed here as single-commodity flow over hypergraphs.

### C. FIDELITY

Constraints on loss of fidelity due to noisy BSM operations and from decoherence due to the age of qubits can be added to the LP formulation, as follows. Recall that constraint on operation-based fidelity loss is modelled by limiting the number leaves of the swapping tree, and in §IV-B, we formulated the decoherence constraint by limiting the heights of the left-most and right-most descendants of the root’s children. These *structural* constraints on swapping trees can be lifted to the LP formulation by adding the leaf count and heights as parameters to *prod* and *avail* vertices; and (ii) swapping the EPs generated from only the compatible vertices.

In particular, we generalize the ST-hypergraph to a *fidelity-constrained* one called  $H^{(F)}$ , where the *prod* and *avail* vertices are parameterized by  $u, v \in V$ , and in addition by  $(n, h)$  where  $n$  is the number of leaves and  $h = (h_{ll}, h_{lr}, h_{rl}, h_{rr})$  represents the depths of left-most and right-most descendants of the root’s children, of the trees rooted at  $(u, v)$  with those parameter values. In terms of edges, the most interesting difference  $H^{(F)}$  and  $H$  is in “Swap” edges. In  $H^{(F)}$ , “Swap” edges are  $\{(avail(u, w, n', h') \text{ } avail(w, v, n'', h'')\}, \{prod(u, v, n, h)\}$  only if  $n = n' + n''$ , and  $h, h', h''$  are such that  $h_{ll} = h'_{ll} + 1, h_{lr} = h'_{rr} + 1, h_{rl} = h''_{ll} + 1$  and  $h_{rr} = h''_{rr} + 1$ . The above constraints ensure that only compatible subtrees are composed into bigger trees. The other changes are for bookkeeping: “Gen” are from  $gen(u, v)$  to  $avail(u, v, 1, (0, 0, 0, 0))$ ; “Prod” are from  $prod(u, v, n, h)$  to  $avail(u, v, n, h)$ ; and finally “Term” are

from  $avail(s, t, n, h)$  to *term* for  $n \leq \tau_l$  and  $h$  such that  $f(h) \leq \tau_d$ ; here,  $f(h)$  gives the tree’s age based on  $h$  values (following §IV-B) while using the link rates based on 50% node-capacity usage.

### APPENDIX B PROOF OF THEOREM 1

**Proof (sketch):** We provide a main intuition behind the claim in Theorem 1. The key claim is that at any instant the WaitLess protocol generates an EP, the Waiting protocol will also be able to generate an EP. Consider an instant  $t$  in time when the WaitLess protocol  $X$  generates an EP, as a result of all the underlying processes succeeding at time  $t$ . Right before time  $t$ , consider the state of the EPs in the swapping-tree  $T$  of the Waiting protocol  $Y$ : Essentially, some of the nodes in  $T$  have (generated) EPs that are waiting for their sibling EP to be generated; note that these generated EPs have not aged yet, else they would have been already discarded by  $Y$ . Now, at time  $t$ , during  $X$ ’s execution, all the underlying processes succeed instantly—it is easy to see that in the protocol  $Y$  too, all the un-generated EP would now be generated instantly<sup>11</sup>—yielding a full EP at the root (using qubits that have not aged beyond the threshold). Finally, since the number of operations in  $T$  is the same as the number of BSM operations incurred by  $X$  to generate an EP, the fidelity degradation due to operations is the same in both the protocols.

### APPENDIX C PROOF OF LEMMA 1

**Proof.** We first prove the claim that given any swapping tree  $\mathcal{T}_{xy}$  over a path  $P : x \rightsquigarrow w \rightsquigarrow y$ , there exists a swapping tree  $\mathcal{T}_{xw}$  over a path  $P' : x \rightsquigarrow w$  such that  $P'$  is a subset of  $P$  and generation latency of  $\mathcal{T}_{xw}$  is less than that of  $\mathcal{T}_{xy}$ . This claim can be easily proved by induction as follows. Consider two cases: (i)  $w$  is the root of  $\mathcal{T}_{xy}$ , in which case  $\mathcal{T}_{xw}$  is the left child of the root. (ii)  $i$  and  $w$  have a common ancestor  $a$  that is other than the root of  $\mathcal{T}_{xy}$ . In this case,  $a = w$ , and the subtree rooted at  $a = w$  is the required  $\mathcal{T}_{xw}$ . (iii) The only common ancestor of  $i$  and  $w$  is the root  $a$  of  $\mathcal{T}_{xw}$ , which is not  $w$ . In this case, we apply the inductive hypothesis on right subtree  $\mathcal{T}_{ay}$  of  $\mathcal{T}_{xy}$ , to extract a subtree  $\mathcal{T}_{aw}$  which along with the left right subtree  $\mathcal{T}_{ia}$  of  $\mathcal{T}_{xy}$  — gives the required subtree  $\mathcal{T}_{xw}$ . This proves the above claim.

Now, to prove the lemma, let us consider the swapping trees  $\mathcal{T}_{ik}$  and  $\mathcal{T}_{kj}$  given to us. By the above claim, there are swapping trees  $\mathcal{T}_{iv}$  and  $\mathcal{T}_{vj}$ , which will satisfy the requirements of the given lemma’s claim.

### APPENDIX D PROOF OF THEOREM 2

**Proof.** We show that  $T[i, j, h, u_i, u_j]$  is indeed the optimal latency over the nodes  $(i, j)$  using a throttled swapping tree of height at most  $h$  and with  $u_j$  and  $u_j$  as the usage percentages at nodes  $i$  and  $j$ . We use proof by induction over  $h$ . The base

<sup>11</sup>Here, we have implicitly assumed that if  $n$  BSM operations succeed in  $X$  protocol at some instant  $t$ , then at the same instant,  $n$  BSM operations anywhere in  $Y$  will also succeed.

TABLE 1: Execution times of QNR-SP algorithm over small networks

Algorithm	Number of nodes					
	10	13	15	16	18	20
Balanced-Tree	239 $\mu$ s	360 $\mu$ s	373 $\mu$ s	492 $\mu$ s	530 $\mu$ s	552 $\mu$ s
DP-Approx	4ms	10ms	14.7ms	17.6ms	28ms	34ms
DP-OPT	148ms	363ms	572ms	706ms	1s	1.7s
Caleffi [18]	92ms	4.6s	14s	26mins	3.2hrs	12.8hrs

case is obvious. The inductive hypothesis is that the above statement is true for all heights  $\leq (h - 1)$ . Now, let  $\mathcal{T}$  be an optimal-latency swapping tree of height at most  $h$  between a pair of nodes  $(i, j)$ , for some height greater than 1, and node usage percentages at  $i$  and  $j$  of  $u_i$  and  $u_j$  respectively. Let the expected latency of  $\mathcal{T}$  be  $L$ . Let the two children subtrees of the root of  $\mathcal{T}$  be  $\mathcal{T}_1$  and  $\mathcal{T}_2$ , each of latency  $L_c$ ; note that, as  $\mathcal{T}$  is throttled, the expected latencies of  $\mathcal{T}_1$  and  $\mathcal{T}_2$  are equal. Thus, we have  $L_c = (\frac{3}{2}L + t_c + t_b)/p_b$  by Eqn. (1). Note that  $\mathcal{T}_1$  and  $\mathcal{T}_2$  are of heights at most  $h - 1$ , and, without loss of generality, we can assume  $\mathcal{T}_1$  and  $\mathcal{T}_2$  to be disjoint (as per Lemma 1). Let  $\mathcal{T}_1$  and  $\mathcal{T}_2$  be between the pairs of nodes  $(i, k)$  and  $(k, j)$  with end-nodes usage percentages of  $(u_i, u_k)$  and  $(u'_k, u_j)$  respectively. Now, optimal throttled trees over  $(i, k)$  and  $(k, j)$  must have a latency of at most that of  $\mathcal{T}_1$  and  $\mathcal{T}_2$ , i.e.,  $L_c$ . Finally, by Eqn. 4 and the inductive hypothesis, we have that the  $T[i, j, h, u_i, u_j]$  (and throttled) will be at most  $L$ .

## APPENDIX E EXECUTION TIMES OF Caleffi [18] ALGORITHM

Here, we give execution times of different algorithms especially Caleffi's for small networks of 10-20 nodes. See Table 1. We see that Balanced-Tree and DP-Approx take fractions of a second, while DP-OPT takes upto 2 seconds. However, as expected Caleffi's execution time increases exponentially with increase in number of nodes – with 20-node network takes 10+ hours. Below, we further estimate Caleffi's execution time for larger graphs.

### Rough Estimate of Caleffi's Execution Time for Large Graphs.

Consider a  $n$ -node network with an average node-degree of  $d$ . Consider a node pair  $(s, d)$ . We try to estimate the number of paths from  $s$  to  $d$  – the goal here is merely to show that the number is astronomical for  $n = 100$ , and thus, our analysis is very approximate (more accurate analysis seems beyond the scope of this work). Let  $P(l)$  be the number of simple paths from  $s$  to a node  $x$  in the graph of length at most  $l$ . For large graphs and large  $l$ , we can assume  $P(l)$  to be roughly same for all  $x$ . We estimate that  $P(l+1) = P(l) + P(l)*6*(1-l/n)$ . The first term is to count paths of length at most  $l - 1$ ; in the second term, the factor 6 comes from the fact the destination  $x$  has 6 neighbors and the factor  $(1 - l/n)$  is the probability that a path counted in  $P(l)$  doesn't contain  $x$  (to constrain the paths to be simple, i.e., without cycles). Using  $P()$ , the execution time of Caleffi can be roughly estimated to be at least  $P(n-1)*500/(5*10^9)$  seconds where the factor 500 is a conservative estimate of the

number of instructions used in computing the latency for a path and  $5*10^9$  is the number of instructions a 5GHz machine can execute in a second. The above yields executions times of a few seconds for  $n = 15$ , about an hours for  $n = 20$ , about 350 hours for  $n = 25$ , and  $10^{16}$  hours for  $n = 50$ , and  $10^{44}$  hours for  $n = 100$ . The above estimates for  $n = 15$  to 20 are within an order of magnitude of our actual execution times, and thus, validate our estimation approach.

## APPENDIX F COMPARISON WITH Caleffi: MORE DETAILS

Fig. 9 shows that DP-OPT outperforms Caleffi by a margin of around 10% when averaging multiple experiments. However, when we look at one experiment at a time and compute the Caleffi's performance relative to DP-OPT for each experiment, we see a larger difference between DP-OPT and Caleffi. Fig. 13 plots the error bar of the relative performance of three algorithms comparing to DP-OPT at each experiment. The lower cap of Caleffi at 0.2 atomic BSM success rate is 0.35, which means that at an extreme sample, the DP-OPT is almost 300% better than Caleffi. In that extreme sample, the number of hops between the source and destination is large (thus the overall EP rate is small, which affects little when averaging with other experiments in Fig. 9). Moreover, we observe that the larger the number of hops between the source and destination, the larger the gap of relative performance is between DP-OPT and Caleffi. This observation aligns with what is shown in Fig. 7(b): our DP-OPT has an larger advantage in ratio when the source and destination are far away.





FIGURE 13. Compare the performance with Caleffi relative to DP-OPT (the closer to 1 the better).



**FIGURE 14.** The execution time comparison of various algorithms for QNR-SP and QNR algorithms.

## APPENDIX G EXECUTION TIMES PLOT

We give here the plot for execution times of various schemes. See Fig. 14.

## REFERENCES

- [1] F. Arute *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, pp. 505–510, 2019.
- [2] J. Gambetta, “IBM’s Roadmap For Scaling Quantum Technology,” <https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/>, 2020.
- [3] M. Caleffi, A. S. Cacciapuoti, and G. Bianchi, “Quantum internet: From communication to distributed computing,” in *5th ACM International Conference on Nanoscale Computing and Communication*, 2018.
- [4] C. Simon, “Towards a global quantum network,” *Nature Photonics*, vol. 11, no. 11, pp. 678–680, 2017.
- [5] Z. Eldredge, M. Foss-Feig, J. A. Gross, S. L. Rolston, and A. V. Gorshkov, “Optimal and secure measurement protocols for quantum sensor networks,” *Physical Review A*, vol. 97, no. 4, p. 042337, 2018.
- [6] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, “The security of practical quantum key distribution,” *Reviews of modern physics*, vol. 81, no. 3, p. 1301, 2009.
- [7] P. Komar, E. M. Kessler, M. Bishop, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, “A quantum network of clocks,” *Nature Physics*, vol. 10, no. 8, pp. 582–587, 2014.
- [8] T.-Y. Chen, H. Liang, Y. Liu, W.-Q. Cai, L. Ju, W.-Y. Liu, J. Wang, H. Yin, K. Chen, Z.-B. Chen *et al.*, “Field test of a practical secure communication network with decoy-state quantum cryptography,” *Optics express*, vol. 17, no. 8, pp. 6540–6549, 2009.
- [9] M. Marcozzi and L. Mostarda, “Quantum consensus: an overview,” *arXiv preprint arXiv:2101.04192*, 2021.
- [10] D. Dieks, “Communication by EPR devices,” *Physics Letters A*, vol. 92, no. 6, pp. 271–272, Nov. 1982.
- [11] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, “Teleporting an unknown quantum state via dual classical and Einstein–Podolsky–Rosen channels,” *Phys. Rev. Lett.*, vol. 70, no. 13, pp. 1895–1899, 1993.
- [12] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, “Long-distance quantum communication with atomic ensembles and linear optics,” *Nature*, 2001.
- [13] N. Sangouard, C. Simon, H. De Riedmatten, and N. Gisin, “Quantum repeaters based on atomic ensembles and linear optics,” *Reviews of Modern Physics*, 2011.
- [14] S. Shi and C. Qian, “Concurrent entanglement routing for quantum networks: Model and designs,” in *SIGCOMM*, 2020.
- [15] K. Chakraborty, D. Elkouss, B. Rijssman, and S. Wehner, “Entanglement distribution in a quantum network: A multicommodity flow-based approach,” *IEEE Transactions on Quantum Engineering*, 2020.
- [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 Information*, 2020.
- [17] K. Chakraborty, F. Rozpedek, A. Dahlberg, and S. Wehner, “Distributed routing in a quantum internet,” 2019, <https://arxiv.org/abs/1907.11630>.
- [18] M. Caleffi, “Optimal routing for quantum networks,” *IEEE Access*, 2017.
- [19] O. Gühne and G. Tóth, “Entanglement detection,” *Physics Reports*, vol. 474, no. 1, pp. 1–75, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0370157309000623>
- [20] W. K. Wootters and W. H. Zurek, “A single quantum cannot be cloned,” *Nature*, vol. 299, no. 5886, pp. 802–803, 1982.
- [21] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, “Optimal architectures for long distance quantum communication,” *Scientific reports*, vol. 6, no. 1, pp. 1–10, 2016.
- [22] S. J. Devitt, W. J. Munro, and K. Nemoto, “Quantum error correction for beginners,” *Reports on Progress in Physics*, vol. 76, no. 7, p. 076001, jun 2013. [Online]. Available: <https://doi.org/10.1088/0034-4885/76/7/076001>
- [23] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, “Quantum repeaters: The role of imperfect local operations in quantum communication,” *Phys. Rev. Lett.*, vol. 81, pp. 5932–5935, Dec 1998. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.81.5932>
- [24] J. Roffe, “Quantum error correction: an introductory guide,” *Contemporary Physics*, vol. 60, no. 3, p. 226–245, Jul 2019. [Online]. Available: <http://dx.doi.org/10.1080/00107514.2019.1667078>
- [25] D. Press, K. De Greve, P. L. McMahon, T. D. Ladd, B. Friess, C. Schneider, M. Kamp, S. Höfling, A. Forchel, and Y. Yamamoto, “Ultrafast optical spin echo in a single quantum dot,” *Nature Photonics*, vol. 4, no. 6, pp. 367–370, 2010.
- [26] H. Wang, Y.-M. He, T.-H. Chung, H. Hu, Y. Yu, S. Chen, X. Ding, M.-C. Chen, J. Qin, X. Yang *et al.*, “Towards optimal single-photon sources from polarized microcavities,” *Nature Photonics*, vol. 13, no. 11, pp. 770–775, 2019.
- [27] Y. Sagiv, I. Almog, and N. Davidson, “Process tomography of dynamical decoupling in a dense cold atomic ensemble,” *Phys. Rev. Lett.*, vol. 105, p. 053201, Jul 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.105.053201>
- [28] E. Vetsch, D. Reitz, G. Sagué, R. Schmidt, S. T. Dawkins, and A. Rauschenbeutel, “Optical interface created by laser-cooled atoms trapped in the evanescent field surrounding an optical nanofiber,” *Phys. Rev. Lett.*, vol. 104, p. 203603, May 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.104.203603>
- [29] C. Deutsch, F. Ramírez-Martínez, C. Lacroûte, F. Reinhard, T. Schneider, J. N. Fuchs, F. Piéchon, F. Lalòë, J. Reichel, and P. Rosenbusch, “Spin self-rephasing and very long coherence times in a trapped atomic ensemble,” *Phys. Rev. Lett.*, vol. 105, p. 020401, Jul 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.105.020401>
- [30] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland, “Long-lived qubit memory using atomic ions,” *Phys. Rev. Lett.*, vol. 95, p. 060502, Aug 2005. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.95.060502>
- [31] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, “Single ion qubit with estimated coherence time exceeding one hour,” *Nature communications*, vol. 12, no. 1, pp. 1–8, 2021.
- [32] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev *et al.*, “Experimental demonstration of memory-enhanced quantum communication,” *Nature*, vol. 580, no. 7801, pp. 60–64, 2020.
- [33] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, “Optimal architectures for long distance quantum communication,” *Scientific Reports*, vol. 6, 2016, <https://doi.org/10.1038/srep20463>.
- [34] W. Tittel, M. Afzelius, R. L. Cone, T. Chanelière, S. Kröll, S. A. Moiseev, and M. Sellars, “Photon-echo quantum memory,” 2008, <https://arxiv.org/abs/0810.0172>.
- [35] J. J. Longdell, E. Fraval, M. J. Sellars, and N. B. Manson, “Stopped light with storage times greater than one second using electromagnetically induced transparency in a solid,” *Phys. Rev. Lett.*, 2005.
- [36] E. Fraval, M. J. Sellars, and J. J. Longdell, “Dynamic decoherence control of a solid-state nuclear-quadrupole qubit,” *Phys. Rev. Lett.*, 2005.
- [37] M. Steger, K. Saeedi, M. Thewalt, J. Morton, H. Riemann, N. Abrosimov, P. Becker, and H.-J. Pohl, “Quantum information storage for over 180 s using donor spins in a 28si “semiconductor vacuum”,” *Science*, vol. 336, no. 6086, pp. 1280–1283, 2012.
- [38] K. Saeedi, S. Simmons, J. Z. Salvail, P. Dluhy, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, J. J. Morton, and M. L. Thewalt, “Room-temperature quantum bit storage exceeding 39 minutes using ionized donors in silicon-28,” *Science*, vol. 342, no. 6160, pp. 830–833, 2013.
- [39] M. Zhong, M. P. Hedges, R. L. Ahlefeldt, J. G. Bartholomew, S. E. Beavan, S. M. Wittig, J. J. Longdell, and M. J. Sellars, “Optically addressable nuclear spins in a solid with a six-hour coherence time,” *Nature*, vol. 517, no. 7533, pp. 177–180, 2015.

[40] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, "Single ion qubit with estimated coherence time exceeding one hour," *Nature Communications*, vol. 12, no. 1, Jan 2021. [Online]. Available: <http://dx.doi.org/10.1038/s41467-020-20330-w>

[41] K. Azuma, K. Tamaki, and H.-K. Lo, "All-photonic quantum repeaters," *Nature communications*, vol. 6, no. 1, pp. 1–7, 2015.

[42] T. Coopmans, S. Brand, and D. Elkouss, "Improved analytical bounds on delivery times of long-distance entanglement," *Phys. Rev. A*, vol. 105, p. 012608, Jan 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.105.012608>

[43] B. Li, T. Coopmans, and D. Elkouss, "Efficient optimization of cut-offs in quantum repeater chains," in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2020, pp. 158–168.

[44] L. Jiang, J. M. Taylor, N. Khaneja, and M. D. Lukin, "Optimal approach to quantum communication using dynamic programming," *Proceedings of the National Academy of Sciences*, vol. 104, no. 44, pp. 17291–17296, 2007. [Online]. Available: <https://www.pnas.org/content/104/44/17291>

[45] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben *et al.*, "A link layer protocol for quantum networks," in *SIGCOMM*, 2019.

[46] W. Kozlowski, A. Dahlberg, and S. Wehner, "Designing a quantum network protocol," in *CoNEXT*, 2020.

[47] L. Bugalho, B. C. Coutinho, and Y. Omar, "Distributing multipartite entanglement over noisy quantum networks," *arXiv preprint arXiv:2103.14759*, 2021.

[48] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. Oliveira, S. Wehner *et al.*, "Netsquid, a discrete-event simulation platform for quantum networks," *Communications Physics*, 2021.

[49] B. Waxman, "Routing of multipoint connections," *IEEE Journal on Selected Areas in Communications*, 1988.

[50] P. van Loock, W. Alt, C. Becher, O. Benson, H. Boche, C. Deppe, J. Eschner, S. Höfling, D. Meschede, P. Michler, F. Schmidt, and H. Weinfurter, "Extending quantum links: Modules for fiber- and memory-based quantum repeaters," *Advanced Quantum Technologies*, vol. 3, no. 11, p. 1900141, 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/qute.201900141>

[51] 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," *Physical Review Letters*, 1996.

[52] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, "Mixed-state entanglement and quantum error correction," *Physical Review A*, 1996.

[53] C. Simon, H. de Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, "Quantum repeaters with photon pair sources and multimode memories," *Phys. Rev. Lett.*, vol. 98, p. 190503, May 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.98.190503>

[54] O. A. Collins, S. D. Jenkins, A. Kuzmich, and T. A. B. Kennedy, "Multiplexed memory-insensitive quantum repeaters," *Phys. Rev. Lett.*, vol. 98, p. 060502, Feb 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.98.060502>

[55] I. Beckenbach, "Matchings and flows in hypergraphs," Ph.D. dissertation, Freie Universität at, Berlin, 2019.

[56] W. Dai, T. Peng, and M. Z. Win, "Optimal remote entanglement distribution," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 3, pp. 540–556, 2020.

[57] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen, "Directed hypergraphs and applications," *Discrete Applied Mathematics*, vol. 42, no. 2, pp. 177–201, 1993.

[58] M. Thakur and R. Tripathi, "Linear connectivity problems in directed hypergraphs," *Theoretical Computer Science*, vol. 410, no. 27, pp. 2592–2618, 2009.



MOHAMMAD GHADERIBANEH received the B.S. degree in electrical engineering from Azad University, Urmia, Iran in 2008 and the M.S. degree in telecommunication engineering from Shahid Beheshti University, Tehran, Iran in 2011. He is currently pursuing the Ph.D. degree in computer science at Stony Brook University, NY, USA.

From 2012 to 2018, he worked at industry as a Software and Telecommunication Engineer developing wireless communication networks like TETRA. In 2022, he interned at Google as a Software Engineer designing a machine learning model to improve YouTube's review process. His research interests includes quantum networks and machine learning application in wireless networks.



CAITAO ZHAN received his B.S. degree in computer science and technology from China University of Geosciences in Wuhan, China in 2017. He then joined the PhD program in Stony Brook University at the Department of Computer Science and is now a PhD candidate. He does research in the broad area of computer networks and his interest lies at the intersection of wireless networks and machine learning. He is also interested in quantum communication and quantum sensing.



HIMANSHU GUPTA is a professor of computer science at Stony Brook University, where he has been a faculty since 2002. His area of research has been in wireless networks, with recent focus on free-space optical communication networks and spectrum management. His current research focuses on quantum networks and communication, and distributed quantum algorithms. He graduated with an M.S. and Ph.D. in Computer Science from Stanford University in 1999, and a B.Tech. in Computer Science and Engineering from IIT Bombay in 1992.



C. R. RAMAKRISHNAN is a professor of computer science at Stony Brook University, where he has been a faculty since 1997. His area of research has been in logic programming and verification, with recent focus on analyzing properties of probabilistic programming. His current research focuses on quantum networks and communication, and distributed quantum algorithms. He graduated with an Ph.D. in Computer Science from Stony Brook University in 1995, and an M.Sc. (Tech.) in Computer Science from BITS Pilani in 1987.