

LINKSELFIE: Link Selection and Fidelity Estimation in Quantum Networks

Maoli Liu, Zhuohua Li*, Xuchuang Wang, John C.S. Lui
The Chinese University of Hong Kong
{mlliu, zhli, xcwang, cslui}@cse.cuhk.edu.hk

Abstract—Reliable transmission of fragile quantum information requires one to efficiently select and utilize *high-fidelity* links among multiple noisy quantum links. However, the fidelity, a quality metric of quantum links, is unknown a priori. Uniformly estimating the fidelity of all links can be expensive, especially in networks with numerous links. To address this challenge, we formulate the link selection and fidelity estimation problem as a best arm identification problem and propose an algorithm named LINKSELFIE. The algorithm efficiently identifies the optimal link from a set of quantum links and provides an accurate fidelity estimate of that link with low quantum resource consumption. LINKSELFIE estimates link fidelity based on the feedback of a vanilla *network benchmarking* subroutine, and adaptively eliminates inferior links throughout the whole fidelity estimation process. This elimination leverages a novel confidence interval derived in this paper for the estimates from the subroutine, which theoretically guarantees that LINKSELFIE outputs the optimal link correctly with high confidence. We also establish a provable upper bound of cost complexity for LINKSELFIE. Moreover, we perform extensive simulations under various scenarios to corroborate that LINKSELFIE outperforms other existing methods in terms of both identifying the optimal link and reducing quantum resource consumption.

Index Terms—Quantum Networks, Link Selection, Fidelity Estimation

I. INTRODUCTION

Quantum networks are capable of transmitting quantum information, represented by quantum bits or qubits, between multiple quantum systems, facilitating groundbreaking applications such as quantum cryptography [1], quantum key distribution (QKD) [2], clock synchronisation [3], and quantum internet-of-things (QIoT) [4]. The principles of quantum mechanics [5] enable quantum networks to achieve functionalities that remain unattainable with classical networks. However, harnessing the potential of quantum networks also introduces new challenges in network design and benchmarking. For instance, direct transmission of qubits via physical quantum links, such as optical fibers, in a large-scale quantum network is not feasible, as the probability of successful transmission diminishes exponentially with the length of quantum links due to quantum decoherence. Additionally, the no-cloning theorem [6] prevents the replication of an arbitrary qubit for re-transmission or amplification, adding further complexity to quantum network design. Besides, quantum information is inherently fragile and easily corrupted by noise, necessitating the benchmarking of quantum networks to ensure their reliability and performance.

*Zhuohua Li is the corresponding author. The work of John C.S. Lui was supported in part by the RGC GRF 14207721.

A quantum network usually consists of quantum nodes connected via quantum links [7]. Each quantum node can function as a source, a destination, or a repeater, with the ability to perform quantum operations and store qubits in its quantum memory. Quantum links, often formed by optical fibers or free-space optical links, can deliver qubits from one end to the other. To realize long-distance qubit transmission, people rely on *quantum entanglement*, a phenomenon where multiple qubits are correlated and the state of individual qubits cannot be described independently of the others. Quantum entanglement is regarded as an important resource for transmitting quantum information. Various experiments [8]–[10] have successfully demonstrated the distribution of quantum entanglement. Once two quantum nodes share entangled pairs, they can transmit quantum information to each other by a process known as quantum *teleportation* [11], regardless of their distance.

Reliable long-distance quantum information transmission is a fundamental requirement for many quantum applications. Many existing works [12]–[15] focus on long-distance quantum entanglement routing protocols. These protocols aim to establish end-to-end entanglement through quantum repeaters via an operation called *entanglement swapping* [11], [16] in a quantum network, with the goal of improving the network throughput, enhancing robustness and serving more users. However, due to the fragile nature of quantum information, qubits are susceptible to decoherence via interactions with the environment. For example, the generated entangled pairs may not be perfectly entangled, and attenuation in physical links and imperfect swapping operations may lead to corruption during the establishment of long-distance entanglement. As a result, the established end-to-end entanglement may not be at the desired states and cannot be used for reliable quantum information transmission. Usually, people use *fidelity* [5] to quantify the quality of an entanglement link. The value of fidelity is from 0 to 1, and it measures how well a quantum channel preserves quantum information. Despite some recent works [17]–[19] that take fidelity guarantee into consideration when designing entanglement routing protocols, it remains essential to explicitly verify the quality of entanglement links before transmitting important quantum information.

The main objective of this work is to efficiently estimate the fidelity of established entangled links. Our approach is based on a method called *network benchmarking* [20], which measures the average fidelity of quantum entanglement links. However, network benchmarking is designed to measure a single quantum

link. In cases where there are multiple links with unknown fidelities—a common scenario in quantum communications, one needs to apply network benchmarking to each link individually, leading to a rather high cost. In practice, one only needs to select a few high-fidelity quantum links to transmit quantum information. Precise fidelity estimation of links with low fidelity is unnecessary where the corresponding benchmarking cost is actually a waste and can be partially saved. Therefore, we consider identifying and eliminating low-fidelity links early on from a set of unknown links, so that we can efficiently obtain accurate fidelity estimates for the desired high-quality links and save quantum resources.

To tackle the challenging problem of link selection and fidelity estimation, we formulate it as a best arm identification problem, a classical sequential decision-making task in multi-armed bandits [21]. Specifically, each arm corresponds to an entangled link in a link set, and each arm is associated with an unknown stochastic reward representing the link fidelity. The network benchmarking method requires transmitting qubits through a link multiple times to obtain a fidelity estimate, which implies that one can only receive a reward or feedback after pulling an arm multiple times. This poses a significant challenge compared to the classical setting, where a single pull of an arm yields an immediate reward sample. Our objective is to identify the link with the highest fidelity from a link set and get its fidelity estimate while consuming as few quantum resources as possible. To achieve this, we design a phase-based elimination algorithm named LINKSELFIE (Link Selection and Fidelity Estimation). The elimination utilizes a novel confidence interval derived in this paper for the results of the vanilla network benchmarking subroutine. Based on the confidence interval, we prove that our algorithm identifies the optimal link with high confidence. We also introduce the notion of cost complexity, which corresponds to quantum resources used by the algorithm and provide a cost complexity upper bound of our algorithm, which theoretically shows that our algorithm is more efficient than the vanilla network benchmarking [20]. To evaluate our algorithm, we simulate a quantum network where a pair of quantum nodes are connected via multiple entanglement links associated with different levels of noise. Our goal is to determine the link with the highest fidelity. We conduct simulations with different noise models and link fidelity distributions. The results corroborate that our algorithm significantly reduces the total benchmarking cost compared to the vanilla network benchmarking method and the naive successive elimination method. Furthermore, our LINKSELFIE algorithm operates without requiring any knowledge of the network topology or the entanglement generation process. Therefore it can be easily integrated into the current multi-layer design of the quantum network protocol stack. This adaptability allows it to seamlessly function with different lower-level protocols, such as entanglement routing or network diagnostic protocols.

We summarize our contributions as follows.

- We formulate the link selection and fidelity estimation problem as a best arm identification problem, enabling




Fig. 1. End-to-end entanglement establishment.

efficient decision-making with limited quantum resources.

- We propose a new algorithm named LINKSELFIE, which utilizes the novel confidence interval for the network benchmarking subroutine. We prove that our algorithm can correctly output the optimal link and its accurate fidelity estimate with high probability. Additionally, we provide a provable cost complexity upper bound for the algorithm.
- We conduct extensive simulations, and the results show that LINKSELFIE can identify optimal links with significantly fewer quantum resources compared to other methods while providing comparable fidelity estimation accuracy.

The rest of the paper is organized as follows. We first present some background information on quantum networks and the network benchmarking method in Section II. In Section III, We formulate the link selection and fidelity estimation problem as a best arm identification problem. The detailed algorithm design and its theoretical analysis are presented in Section IV. Performance evaluations of our algorithm are conducted in Section V, and related work is discussed in Section VI. Finally, Section VII concludes the paper.

II. BACKGROUND

In this section, we review some prerequisites about quantum networks and the network benchmarking method [20].

A. Quantum Networks

A quantum network is composed of quantum nodes interconnected via quantum links. Each quantum node typically contains a quantum processor for performing quantum operations and measurements and has the ability to generate and store quantum states in its limited quantum memory. Quantum links, which can be optical fibers or free-space optical links, facilitate the physical transmission of qubits between quantum nodes, such as the transmission of photons through an optical fiber. However, since the successful transmission rate decreases exponentially with the length of quantum links, people have proposed entanglement-based networks and use quantum entanglement to enhance the transmission of quantum information.

1) *End-to-End Entanglement*: Quantum entanglement arises when the shared state between multiple qubits cannot be factored into a product of its individual qubit states. A classical example is the maximally entangled EPR (Einstein-Podolsky-Rosen) state $|\Psi^\pm\rangle = \frac{|01\rangle \pm |10\rangle}{\sqrt{2}}$. To overcome the distance limitation of end-to-end entanglement between two distant quantum nodes, quantum repeaters are positioned at

intermediate locations within the network. Given a source node S and a destination node D sharing no direct physical links, repeaters along a path of physical links connecting S and D are responsible for generating entanglement (Fig. 1 (a)) with adjacent nodes and executing entanglement swapping operations (Fig. 1 (b)) to establish an end-to-end entanglement link between S and D (Fig. 1 (c)). After the entanglement is successfully established, S can teleport an information qubit to D by consuming that entanglement.

2) *Quantum Noise and Average Fidelity*: When establishing end-to-end entanglement in quantum networks, various sources of noise can potentially arise. For instance, losses in optical fibers, imperfect hardware, and decoherence during qubit storage. Such noise can lead to imperfect entanglement being shared between the source and destination nodes, consequently causing errors in quantum state transmission. Quantum noise is often characterized by *quantum channels* [5]. Examples of noise channels include the bit-flip channel, the depolarizing channel, the dephasing channel, and so on. Specifically, the *depolarizing channel* $\mathcal{E}(\rho) := p\rho + (1-p)\frac{I}{2}$ leaves the input quantum state ρ unchanged with probability p and replaces it with the maximally mixed state $\frac{I}{2}$ (the quantum equivalent of a uniformly random classical bit) otherwise. The depolarizing channel model is a useful tool for characterizing quantum gates [22], [23] and quantum networks [20].

The *average fidelity* [24] of a quantum link associated with noise channel \mathcal{E} is defined by $F(\mathcal{E}) := \int d\psi \text{Tr}[\mathcal{E}(|\psi\rangle\langle\psi|)|\psi\rangle\langle\psi|]$, where the integral is taken uniformly over all pure quantum states $|\psi\rangle$ and Tr denotes the trace of a matrix. The average fidelity quantifies how well a quantum link preserves information, where a fidelity of 1 corresponds to a noiseless channel. For quantum links corresponding to the depolarizing channel with parameter p , the average fidelity is $(1+p)/2$ [5]. For simplicity, we use the term fidelity to refer to the average fidelity hereinafter.

In this work, we adopt the assumption of Markovian noise [20], which implies that the noise is memoryless, and entanglement links constructed via the same path always correspond to the same quantum channel, irrespective of their prior usage history. A pair of source and destination nodes may be interconnected by multiple physical paths, along which entanglement links can be established. Consequently, each entanglement link between two nodes is implicitly associated with a specific physical path within the network. In this paper, we do not consider the low-level details of path discovery or entanglement link generation but assume the set of distinct entanglement links is known.

B. Network Benchmarking

The idea behind network benchmarking is that by channel twirling process [22], which involves random applications of Clifford operations [25], one can transform arbitrary quantum channels into depolarizing channels with the same fidelity. By accessing these depolarizing channels repeatedly, one can estimate the average fidelity of them, which is equivalent to that of the original channel. The parameters used for network




Fig. 2. An example of a quantum network. We abstract away the irrelevant network topology between S and D and only assume there are L entanglement links.

benchmarking are \mathcal{M} and T . \mathcal{M} is called the bounce number set, which contains a series of integers, and T represents the repetition times for each bounce number $m \in \mathcal{M}$. A “bounce” refers to the process in which node S applies a random Clifford operation to the state and sends it to node D , which then does the same and returns it to S . To estimate the average fidelity of entanglement links between nodes S and D , the following procedure will be implemented repeatedly for T times for each $m \in \mathcal{M}$: (i) source node S generates an initial state; (ii) nodes S and D “bounce” the state m times; (iii) source node S applies a final operation and measures the state. The average value of the T measurement results is denoted by b_m , which is often called the survival probability. The survival probability is modeled by the exponential model $b_m = Ap^{2m}$, where A is a constant accounting for quantum state preparation and measurement errors, and p is the depolarizing parameter of the twirled channel. Thus, by fitting the exponential model $b_m = Ap^{2m}$ to the data $\{\mathcal{M}, \{b_m\}_{m \in \mathcal{M}}\}$, one can estimate the parameter \hat{p} and deduce the average fidelity $(1 + \hat{p})/2$.

In quantum networks, transmitting a qubit over a long distance is costly, and each bounce consumes two entangled links. Therefore, we use the number of bounces as a cost metric for network benchmarking. In practice, it is essential not only to precisely estimate the fidelity of the target links but also to conserve the consumption of bounces.

III. MODEL

In this section, we first define the link selection and fidelity estimation problem in quantum networks and provide a motivating example. Then we formulate our problem as a best arm identification problem in multi-armed bandits.

A. Problem Definition

Fig. 2 depicts a quantum network consisting of multiple nodes, among which we select a source node S and a destination node D . There are $L \in \mathbb{N}^+$ entanglement links between S and D , denoted by $\mathcal{L} = \{l_1, \dots, l_L\}$. Each of these links is established along a distinct path connecting S and D . Note that we do not consider the underlying network topology or any specific mechanism for establishing these entanglement links, as they are irrelevant to our purposes. Instead, we abstract away these details and assume the existence of a network

protocol capable of constructing these links. The fidelities of entanglement links can vary due to noise. We denote the fidelity of link $l_i \in \mathcal{L}$ by f_i . The optimal links are those with the highest fidelity, i.e., $l_* \in \arg \max_{l_i \in \mathcal{L}} f_i$.

To ensure high-quality quantum information transmission, it is essential to select entanglement links with the highest fidelity. Since the link fidelities are unknown beforehand, one approach is to use the network benchmarking method described in Section II-B to estimate the fidelity of each link respectively to a specific accuracy $\varepsilon > 0$, i.e., $|\hat{f}_i - f_i| < \varepsilon$ at a given confidence level, and then select the high-fidelity links to transmit quantum information. However, this method means that the benchmarking cost grows linearly with the number of available links L . This indicates that the benchmarking cost becomes prohibitive when the number of links is large.

Motivating Example. A qubit may be corrupted if transmitted through noisy links, causing negative impacts on quantum communications. Since links with low fidelity are unsuitable for transmitting quantum information, it is unnecessary to spend resources and time on accurately estimating these inferior links. For instance, consider two links l_1 and l_2 with unknown fidelities $f_1 = 0.9$ and $f_2 = 0.7$, respectively. When the estimation accuracy reaches 0.05, with a high probability the confidence intervals of the estimate fidelities no longer overlap, enabling us to identify that link l_2 is inferior to link l_1 . Then we can discard l_2 , thus saving the benchmarking cost that would have otherwise been spent on it, and focus on estimating the fidelity of l_1 more precisely. This motivates us to explore an online learning approach to solve the link selection and fidelity estimation problem.

B. Bandit Formulation

We now formulate the link selection and fidelity estimation problem as a best arm identification problem. Let us consider a stochastic multi-armed bandit whose arm set is denoted as $\mathcal{K} := \{1, \dots, L\}$. Each arm $i \in \mathcal{K}$ corresponds to a distinct link in $l_i \in \mathcal{L}$. The reward of each arm $i \in \mathcal{K}$ is associated with a stochastic random variable P_i , whose mean is p_i , satisfying that $f_i = (p_i + 1)/2$ and $p_i \in (0, 1)$, i.e., there is a simple linear relation between the link fidelities and reward means. Without loss of generality, we assume that reward means are sorted in descending order, i.e., $p_1 > p_2 > \dots > p_L$, and arm 1 is the unique optimal arm, since one can always relabel the link index. We denote the reward mean gaps as $\Delta_i := p_1 - p_i$ for suboptimal arm $i \neq 1$, and for $i = 1$, we set $\Delta_1 = \Delta_2$.

In the stochastic multi-armed bandit, the decision maker obtains samples drawn from the reward distribution of an arm by pulling that arm. The empirical reward mean can then be estimated by averaging multiple samples of the arm. However, in the fidelity estimation of quantum entanglement links, the empirical reward (fidelity estimate) is computed via regression by the vanilla network benchmarking subroutine. This implies that in our specific scenario, arms (links) can only be pulled in batches. For simplicity, we regard the act of applying this vanilla network benchmarking subroutine with the bounce number set \mathcal{M} and the repetition times T to link l_i once

equivalent to pulling arm i for T times. Besides, the quantum resources consumed by the subroutine increase linearly in terms of the repetition times T and the summation of all bounces in \mathcal{M} , i.e., $\sum_{m \in \mathcal{M}} m$.

Our goal is to design an algorithm that correctly outputs the optimal arm (link) 1 with the probability of at least $1 - \delta$ (confidence parameter $\delta \in (0, 1)$) with as small costs as possible. We denote the total cost incurred by an algorithm \mathcal{A} as *cost complexity*, which can be expressed as follows,

$$\text{Cost}(\mathcal{A}) := \sum_{i \in \mathcal{L}} N_{i,\tau} \sum_{m \in \mathcal{M}} m,$$

where $N_{i,s}$ is the total number of pulls to link $i \in \mathcal{L}$ up to time s , and τ is the stopping time at which the algorithm \mathcal{A} identifies the optimal arm.

IV. ALGORITHMS

In this section, we first elaborate on the design of our link selection and fidelity estimation algorithm, which we refer to as LINKSELFIE, and present the algorithmic details. We also provide a theoretical analysis of its cost complexity.

A. Algorithm Design

Here, we propose an online algorithm (LINKSELFIE) based on multi-armed bandits results. Unlike the classical bandit setting where each arm can be flexibly pulled, links in our problem can only be pulled in “batches” by the network benchmarking subroutine. This inflexibility requires one to assign multiple pulls for each link in a batched manner. Therefore, we consider phase-based algorithms, where in each phase we assign a fixed number of pulls to each link and run the network benchmarking subroutines to estimate their fidelities. One key challenge in the algorithm design is how to appropriately allocate the number of pulls to different phases, which shall (a) be based on the theoretical estimation performance of this subroutine (see Lemma 1), and also (b) be adaptive to the relative magnitudes of all links fidelities (e.g., reward gaps Δ_i), which are unknown a priori.

To address this challenge, we propose a phase-based elimination algorithm, presented in Algorithm 1. The algorithm runs in phases denoted as $s = 0, 1, 2, \dots$ and proceeds with a link elimination mechanism. In the beginning, the algorithm initializes a candidate set \mathcal{S} as the full link set \mathcal{L} . In each phase, the algorithm executes the vanilla network benchmarking subroutines for all remaining links in the candidate set \mathcal{S} with appropriate pull times T_s according to Lemma 1 (address (a)). At the end of each phase, the algorithm eliminates the inferior links identified in this phase, and, therefore, the algorithm does not need to spend costs on these inferior links (adaptively) in future phases (address (b)).

LINKSELFIE takes three input parameters: the link set \mathcal{L} , the pre-configured bounce length set \mathcal{M} , and the confidence parameter $\delta \in (0, 1)$. T_s is the cumulative number of pulls for remaining links in \mathcal{S} in phase s , i.e., the assigned number of pulls for one link in phase s is $T_s - T_{s-1}$. We set $T_s = C2^{2s} \log \frac{s(s+1)|\mathcal{S}|}{\delta}$ (Line 3), where C is a constant in

Lemma 1. Let $\hat{p}_i^{(s)}$ be the estimate fidelity of link i in phase s and $b_{i,m}^{(s)}$ be the empirical survival probability of link i at the bounce number m after T_s trials. We integrate the vanilla network benchmarking illustrated in Section II-B as a subroutine here, i.e., Benchmarking (Line 5). The subroutine Fitting (Line 7) fits the model $b_m = Ap^{2m}$ to the input data and returns \hat{p} .

Specifically, at the beginning of each phase s , LINKSELFIE uniformly pulls links in the candidate set that have not been eliminated yet for $T_s - T_{s-1}$ times (i.e., run the network benchmarking subroutine with the repetition times $T_s - T_{s-1}$ (Line 5)). Following this, it calculates $b_{i,m}^{(s)}$ by weighted averaging the newly obtained data $b_{i,m}$ and the data $b_{i,m}^{(s-1)}$ from last phase (Line 6), and gets fidelity estimates $\hat{p}_i^{(s)}$ of the remaining links (Line 7). Then it determines the maximum estimate value \hat{p}_{\max} (Line 8) and discards links i satisfying the condition $\hat{p}_i^{(s)} + 2^{-s} < \hat{p}_{\max} - 2^{-s}$ (Line 9). The condition means that the upper confidence bound of link i is less than the lower confidence bound of the link with the current highest empirical fidelity, which implies that link i is a suboptimal link with high confidence and we can eliminate link i . When there is only one link left in the candidate set \mathcal{S} , the algorithm terminates and returns corresponding information.

Algorithm 1: LinkSelfIE: Link Selection & Fidelity Estimation

Input: path set \mathcal{L} , confidence parameter δ , bounce length set \mathcal{M}

Initialization: candidate set $\mathcal{S} \leftarrow \mathcal{L}$, $s \leftarrow 0$,
 $\hat{p}_i^{(0)} \leftarrow 0, \forall i \in \mathcal{L}, T_0 = 0$,
 $b_{i,m}^{(0)} = 0, \forall i \in \mathcal{L}, m \in \mathcal{M}$

```

1 while  $|\mathcal{S}| > 1$  do
2    $s \leftarrow s + 1$  //  $s$  is the current phase.
3    $T_s \leftarrow C \cdot 2^{2s} \log \frac{s(s+1)|\mathcal{S}|}{\delta}$  //  $C$  is a
   constant.
4   for  $i \in \mathcal{S}$  do
5      $\{b_{i,m}\}_{m \in \mathcal{M}} \leftarrow$ 
       Benchmarking( $i, \mathcal{M}, T_s - T_{s-1}$ )
6      $b_{i,m}^{(s)} \leftarrow \frac{b_{i,m}(T_s - T_{s-1}) + b_{i,m}^{(s-1)}T_{s-1}}{T_s}, \forall m \in \mathcal{M}$ 
7      $\hat{p}_i^{(s)} \leftarrow$  Fitting( $\mathcal{M}, \{b_{i,m}^{(s)}\}_{m \in \mathcal{M}}$ )
8    $\hat{p}_{\max} \leftarrow \max_{i \in \mathcal{S}} \hat{p}_i^{(s)}$ 
9    $\mathcal{S} \leftarrow \mathcal{S} \setminus \{i \in \mathcal{S} : \hat{p}_i^{(s)} + 2^{-s} \leq \hat{p}_{\max} - 2^{-s}\}$ 

```

Output: the remaining link in \mathcal{S} and \hat{p}_{\max}

B. Theoretical Analysis

In this subsection, We analyze the cost complexity upper bound of LINKSELFIE.

We begin by presenting the confidence interval for the estimate of the vanilla network benchmarking method described in Section II-B. Since the link fidelity is estimated from an exponential regression using data points of each bounce length $m \in \mathcal{M}$, rather than a simple average of independent random samples, directly applying the conventional analysis tools used

in classical bandit literature, such as Hoeffding's inequality, fails to capture the correlations between data, and thus yields a looser confidence interval. To address this issue, we examine the properties of non-linear regression and derive a new confidence interval in Lemma 1, which exploits information from all data points of each bounce length, offering an advantage over Hoeffding's inequality.

Lemma 1 (Confidence Interval for Network Benchmarking). *Given the input parameters \mathcal{M} and $T \in \mathbb{N}^+$, for a confidence parameter $\delta \in (0, 1)$, to benchmark a link with the true depolarizing parameter p , the depolarizing parameter \hat{p} estimated by the vanilla network benchmarking method satisfies,*

$$\Pr \left[|\hat{p} - p| \leq \sqrt{\frac{C}{T} \log \frac{1}{\delta}} \right] \geq 1 - \delta,$$

where C is a constant related to the bounce length set \mathcal{M} and measurement noise, which is explicitly expressed in Appendix A.

We refer interested readers to Appendix A for its detailed proof. Given Lemma 1, we have the following corollary.

Corollary 1. *Given the confidence radius ε and the confidence level $1 - \delta$, the required repetition times for the vanilla network benchmarking method should satisfy:*

$$T = \mathcal{O} \left(\frac{C}{\varepsilon^2} \log \frac{1}{\delta} \right).$$

From Corollary 1, we know how to set the repetition times T to achieve certain accuracy when benchmarking one entanglement link, which also implies the benchmarking cost.

Based on the above lemma, we derive a cost complexity upper bound for our Algorithm 1. The proof details of Theorem 1 is given in Appendix B.

Theorem 1 (Cost Complexity of Algorithm 1). *Given a bounce length set \mathcal{M} and a confidence parameter $\delta \in (0, 1)$, the cost complexity of Algorithm 1 is upper bounded as follows,*

$$\text{Cost}(\mathcal{A}_{\text{LINKSELFIE}}) \leq \mathcal{O} \left(\sum_{i \in \mathcal{L}} \frac{C}{\Delta_i^2} \log \left(\frac{L}{\delta} \log \frac{4}{\Delta_i} \right) \right) \cdot \sum_{m \in \mathcal{M}} m,$$

where L is the number of links and C is a universal constant.

Remark 1 (Modify Algorithm 1 to find an ε -optimal link). If we modify the algorithm proceeding condition in Line 1 to be more than one link remains in the candidate set and the estimation accuracy does not reach ε , i.e., $|\mathcal{S}| > 1$ and $2^{-s} > \varepsilon$, then Algorithm 1 can output ε -optimal links (whose fidelity is greater than the optimal one minus ε , i.e., $p_i > p_1 - \varepsilon$). Identifying such ε -optimal links is often practical in quantum networks [17]. With this modification, the cost complexity is upper bounded as follows,

$$\text{Cost}(\mathcal{A}_{\text{mod}}) \leq \mathcal{O} \left(\sum_{i \in \mathcal{L}} \frac{C}{\tilde{\Delta}_i^2} \log \left(\frac{L}{\delta} \log \frac{4}{\tilde{\Delta}_i} \right) \right) \cdot \sum_{m \in \mathcal{M}} m, \quad (1)$$

where $\tilde{\Delta}_i := \max\{\Delta_i, \varepsilon\}$.

Remark 2 (Compare to vanilla network benchmarking (uniform exploration)). If one applies the vanilla network benchmarking method to uniformly estimate all links' fidelities and stops at the accuracy ε , then the cost complexity would be upper bounded as follows,

$$\text{Cost}(\mathcal{A}_{\text{vanilla}}) \leq \mathcal{O}\left(\frac{CL}{\varepsilon^2} \log \frac{L}{\delta}\right) \cdot \sum_{m \in \mathcal{M}} m.$$

The above bound is greater than (1) especially when there are some links with large fidelity gaps Δ_i . This corresponds to that when Δ_i 's are large, Algorithm 1 can eliminate these links with high fidelity gaps in early phases, thereby conserving a large number of quantum resources.

V. PERFORMANCE EVALUATION

In this section, we evaluate our algorithm LINKSELFIE. We first elaborate on our experiment setups, including quantum network structure, quantum noise model, baseline algorithms, and the performance metrics we use. Then, we show our evaluation results in terms of efficiency and precision.

A. Experiment Setup

We evaluate our algorithm by simulating a quantum network with two nodes connected via several quantum entanglement links. We employ various common noise models and assign different levels of noise to the quantum links, and the objective is to identify the optimal link using as few quantum resources as possible. The entanglement links are generated by placing a quantum source in the middle and bidirectionally distributing entangled photon pairs through noisy quantum channels. We apply four standard and widely used noise models to simulate quantum noise [5]: (1) depolarizing noise model, (2) dephasing noise model, (3) amplitude damping noise model, and (4) bit flip noise model. For the sake of fair comparison among these noise models, given a fidelity value, we convert it into the corresponding noise parameters used to initialize each noise model. All the quantum mechanisms are simulated by an off-the-shelf quantum network simulation framework called NetSquid [26]. We compare our algorithm with two baselines, (1) the vanilla network benchmarking algorithm (VanillaNB) and (2) the successive elimination algorithm [27] (SuccElimNB). VanillaNB uniformly benchmarks all the quantum links for each bounce number with a fixed number of repetitions T , which we set to $T = 200$ in our experiments. SuccElimNB invokes the network benchmarking subroutine with repetition times $T = 4$ and treats the estimated fidelities as independent random samples, then it makes optimistic decisions based on upper confidence bounds derived via Hoefdding's inequality. It iteratively eliminates identified bad links as learning proceeds until there is only one link left. All the experiments were done on a Linux machine (kernel 6.1.38) with a 3.70 GHz Intel Xeon E5-1630 v4 CPU and 16GB RAM.

B. Quantum Resources Consumption

First, we evaluate the quantum resource consumption of these link selection algorithms. We fix a fidelity gap $\Delta = 0.05$

and simulate a two-node quantum network with L links, which have fidelities $1, 1 - \Delta, 1 - 2\Delta, \dots, 1 - (L - 1)\Delta$, respectively. Then we apply each link selection algorithm to the network for different values of L and measure the quantum resource consumption when the algorithms terminate. Note that we use the total number of bounces as the metric for quantum resource consumption and average the results over 10 trials. Figure 3 shows the plots of quantum resource consumption versus the number of links under different noise models. Since VanillaNB spends the same amount of resources for each link, its resource consumption is proportional to the number of links in the network. This leads to excessive resource waste when the network has many low-quality links. On the other hand, both LINKSELFIE and SuccElimNB can adaptively decide the resources spent by each link, and the cost complexity grows slowly as L increases because it mainly depends on the fidelity gap Δ . However, SuccElimNB only relies on Hoefdding's inequality to estimate the confidence interval, while LINKSELFIE benefits from the much tighter confidence intervals for the non-linear regression (Lemma 1). Therefore, LINKSELFIE collects sufficient information and eliminates the inferior arms more efficiently. The numerical results demonstrate that LINKSELFIE costs significantly fewer resources than the other two algorithms.




Fig. 3. Bounces v.s. the number of links L using different noise models.

Next, we examine the effect of the fidelity gap on the performance of the link selection algorithms. The fidelity gap reflects the hardness of the best arm identification problem. We fix the number of links $L = 4$ and vary the fidelity gap Δ , and we plot the quantum resource consumption versus the fidelity gap under different noise models, averaged over 10 trials. As shown in Figure 4, the cost of VanillaNB is predetermined so it is not affected by the fidelity gap, while LINKSELFIE and SuccElimNB can adaptively adjust the cost according to the fidelity gap. Specifically, the smaller the

fidelity gap is, the harder to identify the optimal link, and therefore more resources are needed. When the fidelity gap is large, both LINKSELFIE and SuccElimNB can identify the optimal link quickly. However, when the fidelity gap is small, benefiting from Lemma 1, LINKSELFIE can obtain a much tighter confidence interval and eliminate inferior links faster, while SuccElimNB has to perform much more benchmarking subroutines to achieve sufficient confidence.




Fig. 4. Bounces v.s. gap using different noise models.

C. Fidelity Estimation Accuracy

Finally, we show the fidelity estimation accuracy of LINKSELFIE. We initialize a two-node network with L links, where the fidelity of each link is generated as follows. We set $\mu_1 = 0.95$ and $\mu_i = 0.85$ for $i = 2, \dots, L$, then for each link i , we sample its fidelity f_i from a Gaussian distribution with mean μ_i and variance $1/4$. We apply each algorithm to the network, get the estimated fidelity of the identified optimal link, and calculate the relative error. Figure 5 plots the relative error versus the number of links L , varying from 2 to 20, and the results are averaged over 10 trials. As expected, LINKSELFIE can not only identify the optimal link but also evaluate its fidelity accurately. The relative error of LINKSELFIE is less than 1%, which has no significant difference compared with other algorithms. When the number of links is large, the relative error tends to decrease because the algorithms need to spend more bounces to distinguish the optimal link, resulting in a more accurate estimation. In summary, LINKSELFIE identifies the optimal link with significantly less quantum resource consumption while providing comparable fidelity estimation accuracy.

VI. RELATED WORK

Quantum networks have gained a lot of attention since their introduction. Several real-world quantum networks have been successfully tested outside the laboratory, such as SECOQC




Fig. 5. Relative error v.s. the number of links L using different noise models.

Vienna QKD network [28], SwissQuantum QKP network [29], Tokyo QKD network [30], and the space-to-ground network [28], demonstrating the promising potential of the upcoming quantum revolution. However, the susceptibility of quantum information to noise necessitates the characterization of quantum link noise for reliable quantum information transmission. Ruan [31] proposes a protocol to estimate the fidelity of entanglement shared by remote nodes, but it does not take quantum measurement errors into account. Helsen *et al.* [20] propose a network benchmarking method, which is robust to state preparation and measurement errors and can efficiently and accurately estimate the fidelity of quantum links, regardless of how these quantum links are formed. Recently, Andrade *et al.* [32] devise novel network tomography protocols (NTP) to characterize the channel noise in quantum networks, but it only considers bit-flip probabilities of quantum star networks. Liu *et al.* [33] propose quantum Border Gateway Protocol (BGP) to support entanglement routing across multiple quantum Internet Service Providers (qISPs) and integrate network benchmarking with the top- K arm identification problem. However, their formulation has more stringent assumptions than ours. Our online link selection algorithm LINKSELFIE leverages the idea of online learning and inherits the favorable properties of network benchmarking. When selecting the optimal link from a set of links, LINKSELFIE is much more efficient than network benchmarking.

The multi-armed bandit (MAB) [34], [35] is a well-known framework with numerous applications in various fields, such as crowdsensing [36], opportunistic channel access [37], and social networks [38]. In this literature, our model is related to best arm identification with fixed confidence, aiming to identify the optimal arm with high probability using as few samples as possible. To solve this problem, algorithms like the successive elimination algorithm [27], the $\text{li}'\text{UCB}$ algorithm [39], and the track-and-stop strategy [40] are proposed and theoretically

analyzed. The problem has also been extended to the best- K arms identification setting [41], [42], with the objective to select top K arms with the highest means with high confidence. Another variant of the best arm identification problem is to identify the best arm in the fixed-budget scenario [21], [43]–[45], which aims to minimize the probability of returning wrong arms at the end of the time budget. We are the first to exploit the MAB framework to solve the link selection and fidelity estimate problem in quantum networks. Our algorithm deliberately considers the property of the network benchmarking, with the objective of identifying the optimal high-fidelity link from a link set with high confidence and consuming as few resources as possible. Unlike using the standard confidence interval based on Hoeffding’s inequality to analyze the sample complexity of algorithms in classical bandits, we deduce a confidence interval for the network benchmarking subroutine and hereby analyze the cost complexity upper bound of our algorithm from a different perspective.

VII. CONCLUSION

In this paper, we consider the problem of link selection and fidelity estimation in quantum networks. To address this challenge, we formulate it as a best arm identification problem and design an efficient algorithm named LINKSELFIE. We derive a novel confidence interval of estimates for the vanilla network benchmarking, a key subroutine within LINKSELFIE, and prove that given a confidence parameter $\delta \in (0, 1)$, with the probability of at least $1 - \delta$, LINKSELFIE outputs the optimal link and its accurate estimate. We also prove a cost complexity upper bound for LINKSELFIE. To validate the performance of LINKSELFIE, we simulate a quantum network with quantum nodes connected via multiple entanglement links. Simulation results show that LINKSELFIE outperforms other methods in efficiently selecting the optimal link and accurately estimating the fidelity while consuming fewer quantum resources across various scenarios.

The authors have provided public access to their code and/or data at <https://zenodo.org/doi/10.5281/zenodo.10444443>.

APPENDIX

A. Proof of Lemma 1

Since the estimate fidelity \hat{p} is obtained by non-linear regression, we use the linear approximation method [46], [47] to derive the confidence interval of \hat{p} for the vanilla network benchmarking.

In the model $b_m = Ap^{2m}$, $\theta = (A, p)^\top$ are the parameters we want to estimate. We denote their true values by $\theta^* = (A^*, p^*)^\top$, and their least-square estimates by $\hat{\theta} = (\hat{A}, \hat{p})^\top$.

In network benchmarking, the observed data (m_i, b_{m_i}) for $i = 1, \dots, |\mathcal{M}|$ are i.i.d. random samples drawn from the non-linear model as follows:

$$b_{m_i} = f(m_i; \theta) + \varepsilon_i = Ap^{2m_i} + \varepsilon_i,$$

where ε_i is the noise term, which follows the normal distribution $\mathcal{N}(0, \sigma^2/T)$. Parameter σ^2 is the single-shot variance for

measuring the survival probability at sequence length $m \in \mathcal{M}$. Since b_{m_i} is the average of T samples, the variance of its noise ε_i is σ^2/T .

The goal of non-linear least squares is to find the parameters $\hat{\theta}$ that minimize the sum of squares of residual errors $S(\theta)$, where

$$S(\theta) = \sum_{i=1}^{|\mathcal{M}|} (b_{m_i} - f(m_i; \theta))^2.$$

We compute the confidence interval using the linear approximation method.

Taking the Taylor expansion of function $f(m_i; \theta)$ at point $\theta^* = (A^*, p^*)^\top$, we have

$$\begin{aligned} f(m_i; \theta) &\approx f(m_i; \theta^*) + \left. \frac{\partial f(m_i; \theta)}{\partial A} \right|_{\theta^*} (A - A^*) \\ &\quad + \left. \frac{\partial f(m_i; \theta)}{\partial p} \right|_{\theta^*} (p - p^*) \\ &= A^* (p^*)^{2m_i} + (p^*)^{2m_i} (A - A^*) \\ &\quad + 2m_i A^* (p^*)^{2m_i-1} (p - p^*), \end{aligned} \quad (2)$$

where the \approx hides a small $o(1)$ quantity.

Hence,

$$\begin{aligned} S(\theta) &= \sum_{i=1}^{|\mathcal{M}|} (b_{m_i} - f(m_i; \theta))^2 \\ &\stackrel{(a)}{\approx} \sum_{i=1}^{|\mathcal{M}|} \left(b_{m_i} - A^* (p^*)^{2m_i} - (p^*)^{2m_i} (A - A^*) \right. \\ &\quad \left. - 2m_i A^* (p^*)^{2m_i-1} (p - p^*) \right)^2 \\ &= \sum_{i=1}^{|\mathcal{M}|} (\varepsilon_i - (\theta - \theta^*)^\top \mathbf{F}_i)^2, \end{aligned} \quad (3)$$

where (a) is by substituting $f(m_i; \theta)$ with its Taylor expansion (2) above, and $\mathbf{F}_i = \left((p^*)^{2m_i}, 2m_i A^* (p^*)^{2m_i-1} \right)^\top$ and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_{|\mathcal{M}|})^\top$.

After the linear approximation, we can minimize (3) by the linear least squares and get:

$$\theta - \theta^* = \left(\sum_{i=1}^{|\mathcal{M}|} \mathbf{F}_i \mathbf{F}_i^\top \right)^{-1} \sum_{i=1}^{|\mathcal{M}|} \varepsilon_i \mathbf{F}_i \quad (4)$$

Equation (4) shows that $\theta - \theta^*$ is a linear combination of $\varepsilon_i \sim \mathcal{N}(0, \sigma^2/T)$. By the definition of multivariate normal distribution, we have

$$\hat{\theta} \sim \mathcal{N}_2 \left(\theta^*, \frac{\sigma^2}{T} \mathbf{Q} \right),$$

where $\mathbf{Q} = \left(\sum_{i=1}^{|\mathcal{M}|} \mathbf{F}_i \mathbf{F}_i^\top \right)^{-1}$.

Since we only care about the estimation of p^* , by the property of multivariate normal distribution, we can extract the second component of $\hat{\theta} = (\hat{A}, \hat{p})^\top$, which has the distribution:

$$\hat{p} \sim \mathcal{N} \left(p^*, \frac{\sigma^2}{T} \mathbf{Q}_{2,2} \right).$$

where $Q_{2,2}$ is the entry of Q in the second row and the second column.

Although $Q_{2,2}$ depends on the true values A^* and p^* , which are unknown to us, the inherent range of A^* and p^* [22] shows that $Q_{2,2}$ is bounded. We can approximate F_i using $\widehat{F}_i = (\widehat{p}^{2m_i}, 2m_i \widehat{A} \widehat{p}^{2m_i-1})^\top$ and so $\widehat{Q} = (\sum_{i=1}^{|\mathcal{M}|} \widehat{F}_i \widehat{F}_i^\top)^{-1}$. We approximate the sample variance s^2 as follows,

$$s^2 = \frac{S(\widehat{\theta})}{|\mathcal{M}| - 2} \approx \frac{|\mathcal{M}| \sigma^2}{(|\mathcal{M}| - 2)T}.$$

Finally, we get the asymptotic confidence interval of p^* by the property of t -distribution as follows,

$$\Pr \left[|\widehat{p} - p^*| > \frac{1}{\sqrt{T}} \sigma \sqrt{\frac{|\mathcal{M}|}{|\mathcal{M}| - 2}} \cdot \widehat{Q}_{2,2} \cdot t_{|\mathcal{M}| - 2, 1 - \alpha/2} \right] \leq \alpha,$$

where $t_{k,p}$ is the p -th quantiles of the t -distribution with k degrees of freedom. That is, the confidence interval depends on $\mathcal{O}(\frac{1}{\sqrt{T}})$.

Let

$$C = \sigma^2 \frac{|\mathcal{M}|}{|\mathcal{M}| - 2} \cdot \widehat{Q}_{2,2} \cdot t_{|\mathcal{M}| - 2, 1 - \alpha/2}^2$$

for a fixed value of α . By applying the powering lemma [48], i.e., repeating the algorithm $\mathcal{O}(\log(\frac{1}{\delta}))$ times and taking the median, the vanilla network benchmarking method can return \widehat{p} such that

$$\Pr \left[|\widehat{p} - p^*| > \sqrt{\frac{C}{T} \log \frac{1}{\delta}} \right] \leq \delta,$$

Then we can conclude Lemma 1.

B. Proof of Theorem 1

Proof. We now prove Theorem 1 by showing that, with probability of at least $1 - \delta$, (a) the optimal link will always remain in the candidate set \mathcal{S} , and (b) the number of pulls required for a suboptimal link i to be removed from \mathcal{S} is bounded by $\mathcal{O}(C \Delta_i^{-2} \log(L/\delta \log(4\Delta_i^{-1})))$.

First, we prove (a): Note that link i will be eliminated from \mathcal{S} only if there exists some $j \neq i$ and some $s > 0$ such that

$$\widehat{p}_i^{(s)} + 2^{-s} < \widehat{p}_j^{(s)} - 2^{-s}. \quad (5)$$

We denote the event $\{\forall s > 0, \forall i \in \mathcal{S}, |\widehat{p}_i^{(s)} - p_i| \leq 2^{-s}\}$ by \mathcal{E} . On the event \mathcal{E} , $p_i \in [\widehat{p}_i^{(s)} - 2^{-s}, \widehat{p}_i^{(s)} + 2^{-s}]$ holds for any link $i \in \mathcal{S}$ in any stage $s > 0$. and thus, $\widehat{p}_i^{(s)} + 2^{-s} \geq p_i$ and $\widehat{p}_j^{(s)} - 2^{-s} \leq p_j$. Plugging them into (5), we have

$$p_i \leq \widehat{p}_i^{(s)} + 2^{-s} < \widehat{p}_j^{(s)} - 2^{-s} \leq p_j,$$

showing that $p_i < p_j$, i.e., i is not an optimal link.

Therefore, if the event \mathcal{E} holds, Algorithm 1 never makes mistakes and must output the correct optimal link in the end. It is sufficient to show that the probability that the event \mathcal{E} happens is no less than $1 - \delta$.

Since

$$\Pr(\bar{\mathcal{E}}) \leq \sum_{s=1}^{\infty} \sum_{i \in \mathcal{S}} \Pr(|\widehat{p}_i^{(s)} - p_i| > 2^{-s}) \quad (6)$$

$$\leq \sum_{s=1}^{\infty} \sum_{i \in \mathcal{S}} \frac{\delta}{|\mathcal{S}|s(s+1)} \quad (7)$$

$$< \delta, \quad (8)$$

where (6) holds by the union bound, (7) is because of Lemma 1, and (8) is from $\sum_{s=1}^{\infty} \frac{1}{s(s+1)} < 1$, i.e., $\Pr(\mathcal{E}) \geq 1 - \delta$, we can conclude (a).

Then we prove (b): We consider the cost complexity when \mathcal{E} holds. We show that a suboptimal link i must have been eliminated when $\Delta_i \leq 4 \cdot 2^{-s}$; otherwise, it results in a contradiction. Suppose $\Delta_i > 4 \cdot 2^{-s}$ and link i is not eliminated in phase s , implying that $\widehat{p}_i^{(s)} + 2^{-s} \geq \widehat{p}_{\max} - 2^{-s}$. Then,

$$p_i + 2 \cdot 2^{-s} \geq \widehat{p}_i^{(s)} + 2^{-s} \geq \widehat{p}_{\max} + 2^{-s} \geq \widehat{p}_1^{(s)} - 2^{-s} \geq p_1 - 2 \cdot 2^{-s},$$

holds, i.e., $\Delta_i \leq 4 \cdot 2^{-s}$, which contradicts with $\Delta_i > 4 \cdot 2^{-s}$.

Denote the phase in which link i is eliminated by phase s_i . We have $\Delta_i \leq 4 \cdot 2^{-s_i}$. Therefore, the total number of pulls of link i is upper bounded as follows,

$$N_{i,s_i} = C 2^{2s_i} \log \frac{s_i(s_i+1)|\mathcal{S}|}{\delta} \quad (9)$$

$$\leq C 2^{2s_i} \log \frac{s_i(s_i+1)L}{\delta} \quad (10)$$

$$\leq C 2^{2s_i} \log \left(\frac{L}{\delta} \log \frac{4}{\Delta_i} (\log \frac{4}{\Delta_i} + 1) \right) \quad (11)$$

$$= \mathcal{O} \left(\frac{C}{\Delta_i^2} \log \left(\frac{L}{\delta} \log \frac{4}{\Delta_i} \right) \right), \quad (12)$$

where (9) is from the definition, (10) is because of $|\mathcal{S}| \leq L$, (11) and (12) are because of $\Delta_i \leq 4 \cdot 2^{-s_i}$. Then we have (b).

Summing over the number of pulls of all links and multiplying it with the cost per pull, we have:

$$\text{Cost}(\mathcal{A}_{\text{LINKSEL}}) \leq \mathcal{O} \left(\sum_{i \in \mathcal{L}} \frac{C}{\Delta_i^2} \log \left(\frac{L}{\delta} \log \frac{4}{\Delta_i} \right) \right) \cdot \sum_{m \in \mathcal{M}} m,$$

which concludes the proof. \square

REFERENCES

- [1] A. K. Ekert, "Quantum cryptography based on Bell's theorem," *Phys. Rev. Lett.*, vol. 67, pp. 661–663, Aug 1991.
- [2] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," in *Proceedings of the International Conference on Computers, Systems & Signal Processing*, 1984, pp. 175–179.
- [3] P. Komar, E. M. Kessler, M. Bishof, 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.
- [4] L. Chen, Q. Chen, M. Zhao, J. Chen, S. Liu, and Y. Zhao, "DDKA-QKDN: Dynamic On-Demand Key Allocation Scheme for Quantum Internet of Things Secured by QKD Network," *Entropy*, vol. 24, no. 2, p. 149, 2022.
- [5] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge university press, 2010.
- [6] W. K. Wootters and W. H. Zurek, "A Single Quantum Cannot be Cloned," *Nature*, vol. 299, no. 5886, pp. 802–803, 1982.

- [7] W. Kozłowski, A. Dahlberg, and S. Wehner, "Designing a Quantum Network Protocol," in *Proceedings of the 16th International Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '20, 2020, p. 1–16.
- [8] T. Inagaki, N. Matsuda, O. Tadanaga, M. Asobe, and H. Takesue, "Entanglement distribution over 300 km of fiber," *Optics express*, vol. 21, no. 20, pp. 23 241–23 249, 2013.
- [9] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. P. Lanyon, "Light-matter entanglement over 50 km of optical fibre," *npj Quantum Information*, vol. 5, no. 1, p. 72, 2019.
- [10] Y. Yu, F. Ma, X.-Y. Luo, B. Jing, P.-F. Sun, R.-Z. Fang, C.-W. Yang, H. Liu, M.-Y. Zheng, X.-P. Xie *et al.*, "Entanglement of two quantum memories via fibres over dozens of kilometres," *Nature*, vol. 578, no. 7794, pp. 240–245, 2020.
- [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, pp. 1895–1899, Mar 1993.
- [12] S. Shi and C. Qian, "Concurrent Entanglement Routing for Quantum Networks: Model and Designs," in *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '20, 2020, p. 62–75.
- [13] Y. Zeng, J. Zhang, J. Liu, Z. Liu, and Y. Yang, "Multi-entanglement routing design over quantum networks," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 510–519.
- [14] C. Li, T. Li, Y.-X. Liu, and P. Cappellaro, "Effective routing design for remote entanglement generation on quantum networks," *npj Quantum Information*, vol. 7, no. 1, pp. 1–12, Jan. 2021.
- [15] A. Farahbakhsh and C. Feng, "Opportunistic Routing in Quantum Networks," in *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, 2022, pp. 490–499.
- [16] M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, "Event-ready-detectors" Bell experiment via entanglement swapping," *Phys. Rev. Lett.*, vol. 71, pp. 4287–4290, Dec 1993.
- [17] J. Li, M. Wang, K. Xue, R. Li, N. Yu, Q. Sun, and J. Lu, "Fidelity-guaranteed entanglement routing in quantum networks," *IEEE Transactions on Communications*, vol. 70, no. 10, pp. 6748–6763, 2022.
- [18] M. Liu, J. Allcock, K. Cai, S. Zhang, and J. C. Lui, "Quantum Networks with Multiple Service Providers: Transport Layer Protocols and Research Opportunities," *IEEE Network*, vol. 36, no. 5, pp. 56–62, 2022.
- [19] Y. Zhao, G. Zhao, and C. Qiao, "E2E fidelity aware routing and purification for throughput maximization in quantum networks," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 480–489.
- [20] J. Helsen and S. Wehner, "A Benchmarking Procedure for Quantum Networks," *npj Quantum Information*, vol. 9, no. 1, p. 17, 2023.
- [21] J.-Y. Audibert and S. Bubeck, "Best Arm Identification in Multi-Armed Bandits," in *COLT - 23th Conference on Learning Theory - 2010*, Jun. 2010, p. 13 p.
- [22] E. Magesan, J. M. Gambetta, and J. Emerson, "Characterizing quantum gates via randomized benchmarking," *Physical Review A*, vol. 85, no. 4, p. 042311, 2012.
- [23] J. Helsen, J. J. Wallman, S. T. Flammia, and S. Wehner, "Multiqubit randomized benchmarking using few samples," *Physical Review A*, vol. 100, no. 3, p. 032304, 2019.
- [24] M. A. Nielsen, "A simple formula for the average gate fidelity of a quantum dynamical operation," *Physics Letters A*, vol. 303, no. 4, pp. 249–252, 2002.
- [25] D. Gottesman, "The Heisenberg representation of quantum computers," *arXiv preprint quant-ph/9807006*, 1998.
- [26] T. Coopmans, R. Knegijens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpedek, M. Skrzypczyk *et al.*, "Netsquid, a network simulator for quantum information using discrete events," *Communications Physics*, vol. 4, no. 1, p. 164, 2021.
- [27] E. Even-Dar, S. Mannor, and Y. Mansour, "Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems," *J. Mach. Learn. Res.*, vol. 7, p. 1079–1105, dec 2006.
- [28] Y.-A. Chen, Q. Zhang, T.-Y. Chen, W.-Q. Cai, S.-K. Liao, J. Zhang, K. Chen, J. Yin, J.-G. Ren, Z. Chen *et al.*, "An integrated space-to-ground quantum communication network over 4,600 kilometres," *Nature*, vol. 589, no. 7841, pp. 214–219, 2021.
- [29] D. Stucki, M. Legre, F. Buntschu, B. Clausen, N. Felber, N. Gisin, L. Henzen, P. Junod, G. Litzistorf, P. Monbaron *et al.*, "Long-term performance of the SwissQuantum quantum key distribution network in a field environment," *New Journal of Physics*, vol. 13, no. 12, p. 123001, 2011.
- [30] M. Sasaki, M. Fujiwara, H. Ishizuka, W. Klaus, K. Wakui, M. Takeoka, S. Miki, T. Yamashita, Z. Wang, A. Tanaka *et al.*, "Field test of quantum key distribution in the Tokyo QKD Network," *Optics express*, vol. 19, no. 11, pp. 10 387–10 409, 2011.
- [31] L. Ruan, "Minimization of the estimation error for entanglement distribution networks with arbitrary noise," *Phys. Rev. A*, vol. 108, p. 022418, Aug 2023.
- [32] M. G. De Andrade, J. Navas, I. Montañó, and D. Towsley, "On the Characterization of Quantum Flip Stars with Quantum Network Tomography," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, 2023, pp. 1260–1270.
- [33] M. Liu, Z. Li, K. Cai, J. Allcock, S. Zhang, and J. C. Lui, "Quantum BGP with Online Path Selection via Network Benchmarking," in *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications*, 2024.
- [34] T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 2020.
- [35] S. Bubeck, N. Cesa-Bianchi *et al.*, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Foundations and Trends® in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
- [36] G. Gao, J. Wu, M. Xiao, and G. Chen, "Combinatorial multi-armed bandit based unknown worker recruitment in heterogeneous crowdsensing," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2020, pp. 179–188.
- [37] K. Cai, X. Liu, Y.-Z. J. Chen, and J. C. Lui, "An online learning approach to network application optimization with guarantee," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 2006–2014.
- [38] X. Liu, J. Zuo, X. Chen, W. Chen, and J. C. S. Lui, "Multi-layered Network Exploration via Random Walks: From Offline Optimization to Online Learning," in *Proceedings of the 38th International Conference on Machine Learning*, vol. 139. PMLR, 18–24 Jul 2021, pp. 7057–7066.
- [39] K. Jamieson and R. Nowak, "Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting," in *2014 48th Annual Conference on Information Sciences and Systems (CISS)*. IEEE, 2014, pp. 1–6.
- [40] A. Garivier and E. Kaufmann, "Optimal Best Arm Identification with Fixed Confidence," in *29th Annual Conference on Learning Theory*, vol. 49. PMLR, 23–26 Jun 2016, pp. 998–1027.
- [41] M. Simchowitz, K. Jamieson, and B. Recht, "The Simulator: Understanding Adaptive Sampling in the Moderate-Confidence Regime," in *Proceedings of the 2017 Conference on Learning Theory*, vol. 65. PMLR, 07–10 Jul 2017, pp. 1794–1834.
- [42] H. Jiang, J. Li, and M. Qiao, "Practical algorithms for best-k identification in multi-armed bandits," *arXiv preprint arXiv:1705.06894*, 2017.
- [43] S. Bubeck, T. Wang, and N. Viswanathan, "Multiple Identifications in Multi-Armed Bandits," in *Proceedings of the 30th International Conference on Machine Learning*, vol. 28, no. 1. PMLR, 17–19 Jun 2013, pp. 258–265.
- [44] Y. Zhou, X. Chen, and J. Li, "Optimal PAC Multiple Arm Identification with Applications to Crowdsourcing," in *Proceedings of the 31st International Conference on Machine Learning*, vol. 32, no. 2. PMLR, 22–24 Jun 2014, pp. 217–225.
- [45] V. Gabillon, M. Ghavamzadeh, and A. Lazaric, "Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence," in *Advances in Neural Information Processing Systems*, vol. 25, 2012.
- [46] A. K. Bera, "The use of linear approximation to nonlinear regression analysis," *Sankhyā: The Indian Journal of Statistics, Series B*, pp. 285–290, 1984.
- [47] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta, "Investigating the limits of randomized benchmarking protocols," *Phys. Rev. A*, vol. 89, p. 062321, Jun 2014.
- [48] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani, "Random generation of combinatorial structures from a uniform distribution," *Theoretical Computer Science*, vol. 43, pp. 169–188, 1986.