2026 International Conference on Quantum Communications, Networking, and Computing (QCNC)

Online Optimal Probe Allocation for Quantum Network Tomography

Xuchuang Wang
CICS, UMass Amherst
Amherst, MA, USA
xuchuangwang@cs.umass.edu

Yu-Zhen Janice Chen
CICS, UMass Amherst
Amherst, MA, USA
yuzhenchen@cs.umass.edu

Matheus Guedes de Andrade
CICS, UMass Amherst
Amherst, MA, USA
mguedesdeandrade@cs.umass.edu

Mohammad Hajiesmaili
CICS, UMass Amherst
Amherst, MA, USA
hajiesmaili@cs.umass.edu

John C.S. Lui
CSE, CUHK
Shatin, N.T., Hong Kong
csliu@cse.cuhk.edu.hk

Ting He
EECS, PSU
University Park, PA, USA
tinghe@psu.edu

Don Towsley
CICS, UMass Amherst
Amherst, MA, USA
towsley@cs.umass.edu

***Abstract*—How to efficiently perform network tomography is a fundamental problem in network management and monitoring. A network tomography task usually consists of applying multiple probing experiments, e.g., across different paths or via different casts (e.g., unicast and multicast). We study how to optimize the network tomography process through online sequential decision-making. From the methodology perspective, we introduce an online probe allocation algorithm that sequentially performs network tomography based on the principles of optimal experimental design and the maximum likelihood estimation. We rigorously analyze the regret of the algorithm under the conditions that *i)* the optimal allocation is Lipschitz continuous in the parameters being estimated and *ii)* the parameter estimators satisfy a concentration property. From the application perspective, we demonstrate that the quantum bit-flip network fulfills the two theoretical conditions and provide their corresponding regrets when deploying our proposed online probe allocation algorithm. Besides case studies with theoretical guarantees, we also conduct simulations to compare our proposed algorithm with existing methods and demonstrate our algorithm's effectiveness.**

***Index Terms*—Network tomography, quantum network, online learning, experimental design.**

I. Introduction

NETWORK tomography [1] is essential for inferring the internal network (e.g., link) parameters, such as the loss rate, delay, and bandwidth, via end-to-end (external) measurements. These external measurements are practically more accessible than the network's internal components (e.g., routers), which may be owned by different internet service providers (ISPs) [2] and therefore difficult to access. A stochastic network tomography problem consists of a network with a set $\mathcal{L}$ of $L := |\mathcal{L}|$ links, each link $\ell \in \mathcal{L}$ with an unknown parameter μ_ℓ that characterizes its stochastic property (e.g., μ_ℓ may represent loss rate or average delay of the link), and a set $\mathcal{M}$ of $M := |\mathcal{M}|$ probes. Performing a probe on the network refers to generating one or multiple stochastic measurements (i.e., *observations*) that depend on the network parameters. Network tomography aims to estimate the network parameters μ_ℓ from the observations.

Most prior works on network tomography focus on devising estimators for the network parameters from the stochastic observations, e.g., in packet delay tomography [3, 4] and loss network tomography [2, 5] (detailed in Section I-B). That is, after obtaining these stochastic observations from probes, prior works aim to devise good *static* estimators to infer the network parameters. While static estimator design is crucial for network tomography, another essential yet less explored task is, without knowing the network parameters $(\mu_\ell)_{\ell\in\mathcal{L}}$, how to *dynamically* collect online observations to *efficiently* estimate the network parameters using the least number of probes. In short, this is a sequential decision-making problem (a.k.a., online learning) [6] for network tomography.

However, to apply online learning techniques to network tomography, one needs to tackle unique challenges not present in common online learning scenarios. A key technical challenge comes from the complex feedback mechanism in network tomography. In standard online learning settings, e.g., basic multi-armed bandits [7] or linear bandits [8], the feedback is usually a scalar stochastic reward or loss, which is only related to the action (probe) taken by the algorithm. In contrast, in network tomography, feedback consists of the stochastic observations (e.g., a random vector) generated by the probes via the compound effect of the links across a set of paths traversed by the probe. The observations often have a non-linear relation with the network link parameters, which may correlate with each other, making the combinatorial bandits framework [9] difficult to apply. This complex feedback mechanism makes the design of the online algorithm and the analysis of its regret challenging.

A. Contributions

This paper considers the network tomography task as an online sequential decision making problem and focuses on two new perspectives: (1) A general framework that allows us to handle both classical and quantum network tomography tasks [10]; (2) fine-grained analysis, where we develop new algorithms to perform online network tomography with regret guarantees. Specifically, we make the following contributions:

979-8-3315-6110-9/26/$31.00 ©2026 IEEE
DOI 10.1109/QCNC69040.2026.00012

2026 International Conference on Quantum Communications, Networking, and Computing (QCNC) | 979-8-3315-6110-9/26/$31.00 ©2026 IEEE | DOI: 10.1109/QCNC69040.2026.00012

► First, we formulate a general network tomography task as an online experimental design problem (Section II). One attractive feature of the model is that it covers both classical and quantum tomography scenarios. The problem of adaptively collecting probe observations and efficiently inferring the network parameters is formulated as a *regret* minimization problem, where the objective is to minimize the difference between the performance of an online algorithm and the optimal policy in terms of optimal experimental design (OED) criteria [11]. To our knowledge, this paper is the first to design network tomography algorithms for minimizing regret.

► Second, we devise the Online Probe Allocation (OPAL) algorithm for the general network tomography task (Section III). The design of this algorithm is inspired by the idea of "chasing optimal bounds," which is the backbone of a class of online learning algorithms with optimal performance [12]. We refer to the actual probe allocation ratio (the allocation fraction of each probe) of the algorithm as *actual allocation* and the optimal allocation ratio based on the estimated network parameters as *estimated optimal allocation*. OPAL uses the difference between the actual allocation and the estimated optimal allocation to direct the algorithm to sample the most inadequately performed probes (compared with the estimated allocation ratio) sequentially.

► Third, we propose a new analysis framework for studying the regret of OPAL (Section IV). Analyzing the regret of OPAL is challenging: The chased estimated optimal allocation—determined by the estimated network parameters whose accuracy depends on the past probing decisions—varies over time. Furthermore, the online algorithm's actual allocation depends on all its past probe decisions, which lags behind the time-varying estimated optimal allocation it chases. We propose two critical conditions (Section IV-A): *Lipschitz continuity* (Condition 1) and *confidence interval concentration* (Condition 2), that abstract the property of any network tomography task. With these two theoretical conditions, we then rigorously prove that OPAL achieves a $\tilde{O}((\xi T)^{-\gamma_{\min}} + \xi\mathbb{1}\{\exists m \in \mathcal{M} : \phi_m^* < \xi\})$ regret (convergence rate), where T is the total number of probes, $\gamma_{\min} \in (0, 1/2)$ is the smallest concentration rate of Condition 2, $\xi \in (0, 1)$ is an input constant of OPAL, $\mathbb{1}\{\}$ is the indicator function, ϕ_m^* is the optimal allocation ratio of the m-th probe, and the $\tilde{O}(\cdot)$ notation hides poly-logarithmic $\log T$ terms.

► Last but not least, we demonstrate the effectiveness of our proposed framework and algorithms from both theoretical and empirical aspects in quantum network tomography (QNT) (Sections V). We verify the two theoretical conditions for bit-flip channel tomography in a quantum star network—the state-of-the-art quantum network studied in QNT literature [10, 13]—with multicast probes and prove a $\tilde{O}(T^{-1/3})$ regret for this quantum tomography task (Section V-A). Finally, we provide the empirical results of OPAL for this quantum network tomography task (Section V-B) and show that OPAL outperforms alternative baseline policies.

B. Related Works

Network tomography has been studied for decades [2, 3, 4, 14, 15, 16, 5, 17, 18, 19, 10, 13], and most prior works focus on network parameter estimation. For example, for delay network tomography [16], Lo Presti et al. [4] study estimator design for unicast probes, and Duffield and Lo Presti [3] study the multicast probes. For loss network tomography, Cáceres et al. [2] devise the maximum likelihood estimators, and Xi et al. [5] propose EM-based estimations, both of which are for multicast probes. Recently, network tomography in quantum networks is studied by De Andrade et al. [10, 13], where, due to the complexity of the quantum network and information, only multicast probes on a quantum star network have been considered. The estimators and probes studied in these previous works form a solid foundation for our study from the online experimental design perspective. But, as our work focuses on the sequential probe selection and theoretical guarantees, the existing works are not directly applicable to our study.

Optimal experimental design (OED) aims to design experiments to achieve specific statistical criteria [11]. It has been applied in a wide range of applications, including clinical trial [20], synthetic biology [21], and chemistry [22], etc. OED for network tomography has also been studied, He et al. [19], Kveton et al. [23], Xi et al. [5], to name a few. Among them, the most related work to this paper is He et al. [19], where they propose a heuristic iterative algorithm in terms of OED for the unicast network tomography task, and only an asymptotic convergence guarantee is provided. In this paper, we make a concrete step forward by proposing a rigorous *online* experimental design algorithm, OPAL, for network tomography, and we are the *first* to provide a fine-grained regret analysis for the network tomography task. Outside the network tomography literature, online experimental design for linear regression is studied by Fontaine et al. [24], Allen-Zhu et al. [25]. However, the linear regression model is too simple to apply to network tomography, whose feedback can be non-linear and correlated.

II. MODELING AND FORMULATION

In this section, we first formulate the general network tomography problem in Section II-A and then present the quantum bit-flip network tomography problems in Section II-B to further elaborate on the tomography tasks. Preliminaries of the optimal experimental design are introduced in Section II-C. Lastly, we cast the network tomography problem as an online experimental design problem in Section II-D.

A. General Network Tomography

Let $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ denote the network topology with a set of $N \in \mathbb{N}^+$ nodes $\mathcal{N} := \{1, 2, \ldots, N\}$ and a set of $L \in \mathbb{N}^+$ links $\mathcal{L} := \{1, 2, \ldots, L\}$. Associated with each link $\ell \in \mathcal{L}$ is a stochastic model X_ℓ representing the characteristic of the link, e.g., a Bernoulli random variable for signal loss [2], or the noise level or fidelity for quantum channel [26]. We refer to a connection between two nodes (source and destination) across several consecutive links as a *path* in the network. The tomography task assumes that the link model is known, but that its parameters, denoted as a vector $\boldsymbol{\mu} := (\mu_\ell)_{\ell \in \mathcal{L}}$ (e.g., the link success rates), are unknown.

The network tomography task aims to estimate these unknown link parameters via a given set of $M \in \mathbb{N}^+$ probing experiments $\mathcal{M} := \{1, 2, \ldots, M\}$ (called *probes* later), examples of which are unicast [19] and multicast [2]. Each probe $m \in \mathcal{M}$ is associated with one path or a tree rooted at a source node (sender) whose leaves are the set of destination nodes (receivers) denoted as $\mathcal{N}_m \subseteq \mathcal{N}$. The feedback $\boldsymbol{Y}_m \in \mathbb{R}^{|\mathcal{N}_m|}$ from the probing experiment m consists of received signals at the destination nodes in $\mathcal{N}_m$. Denote the probability of observing feedback $\boldsymbol{Y}_m$ given network parameters $\boldsymbol{\mu}$ and probing experiment m as $f_m(\boldsymbol{Y}_m; \boldsymbol{\mu})$, where $f_m(\cdot; \boldsymbol{\mu})$ is the probability density function. The goal is to estimate the network parameters based on the feedback from the probing experiments. For ease of presentation, we follow the common identifiability assumption in network tomography literature [15, 19] that the network parameters $\boldsymbol{\mu}$ can be uniquely determined from feedback from an appropriate set of probing experiments in $\mathcal{M}$. Given a set of $T \in \mathbb{N}^+$ probe-feedback pairs $\mathcal{H} := \{(m_t, \boldsymbol{Y}_{m_t,t})\}_{t=1}^T$, we define the log-likelihood function as $L(\boldsymbol{\mu}; \mathcal{H}) := \sum_{t=1}^T \log f_{m_t}(\boldsymbol{Y}_{m_t,t}; \boldsymbol{\mu})$. The maximum likelihood estimation (MLE) estimator for the network parameters $\boldsymbol{\mu}$ is defined as $\texttt{MLE}(\mathcal{H}) := \arg\max_{\boldsymbol{\mu}} L(\boldsymbol{\mu}; \mathcal{H})$.[1]

Next, we introduce quantum bit-flip network tomography as examples to illustrate network tomography.

B. Quantum Bit-flip Network Tomography

Quantum networks are communication systems used to interconnect quantum devices such as quantum processors. Despite differences between the physical properties of classical and quantum information systems, network tomography extends naturally to the quantum setting. Quantum network tomography (QNT) addresses the characterization of quantum channels representing links in a quantum network through end-to-end measurements [10]. QNT is defined under the assumption that the quantum channel representing link $\ell \in \mathcal{L}$, has the parametric form $\mathcal{E}(\rho) = \sum_k M_{\ell k}^\dagger(\mu_\ell) \rho M_{\ell k}(\mu_\ell)$, where $\mu_\ell \in \mathbb{R}^{d_\ell}$ for $d_\ell \in \mathbb{N}^+$ are the parameters, ρ is an input state, $\{M_{\ell k}\}$ denote the Kraus operators for $\mathcal{E}_\ell$. The goal of QNT is to estimate μ_ℓ for each link $\ell \in \mathcal{L}$.

In this work, we study the quantum network tomography of bit-flip channels in a quantum network with a star topology [10, 13]. Quantum network links are assumed to be *single-qubit, bit-flip* channels of the form $\mathcal{E}(\rho) = \mu_\ell \rho + (1 - \mu_\ell)\sigma_X \rho \sigma_X$, where $\mu_\ell \in (0, 1)$ denotes the probability of no bit-flip for link $\ell \in \mathcal{L}$, and σ_X is the Pauli X operator [26]. We utilize probes defined in [13] that rely on multipartite quantum state distribution, focusing on the *Root Independent* distribution strategy depicted in Fig 1, which we now describe. The number of nodes N in the star network equals the number of links L plus one, i.e., $N = L+1$. The process starts with the selection of a root node (link) ℓ_{root}, which sends a qubit in state

[1]Throughout this paper, we assume the outputs of the $\arg\min$ and $\arg\max$ functions are unique; otherwise, one can break the tie arbitrarily.

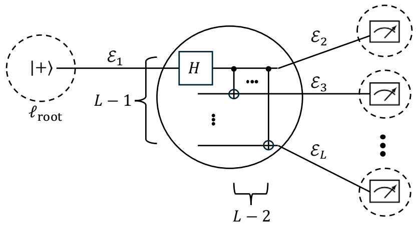

Fig. 1: Root independent multicast for quantum star network tomography. Root $\ell_{\text{root}} = 1$ prepares state $|+\rangle$, transmits it to the central node, and an $(L-1)$ qubit state is prepared from the application of CNOT gates. Each output is transmitted to an end node via the links representing the quantum channels $\mathcal{E}_\ell$, for $\ell = 2, \ldots, L$. σ_Z basis measurements performed in the end nodes generate boolean random variables $\boldsymbol{Y}_{\ell_{\text{root}}} \in \{0, 1\}^{L-1}$.

$|+\rangle$ to the intermediate node.[2] Upon receiving the qubit from ℓ_{root}, the intermediate node applies a Hadamard gate [26] in the qubit and performs an $(L-1)$-qubit CNOT gate on $(L-2)$ ancilla qubits initialized in state $|0\rangle$ controlled by the received qubit. Each output of the CNOT gate is sent to the end nodes of the star, excluding ℓ_{root}, such that an $(L-1)$-qubit quantum state is prepared in the rest $(L-1)$ end nodes. The process terminates with σ_Z-basis measurements performed on the qubits received by the end nodes. The probability of obtaining the $|b_2\, b_3 \ldots b_L\rangle$ state (for $b_\ell \in \{0, 1\}$) as the classical outcome at the $L-1$ end nodes after sending $|+\rangle$ at node $\ell_{\text{root}} = 1$ is $f_{m=1}(Y_m = |b_2\, b_3 \ldots b_L\rangle\, ; \boldsymbol{\mu}) = \prod_{\ell=2}^L \mu_\ell^{b_\ell}(1 - \mu_\ell)^{1-b_\ell}$.

Later, we provide a detailed case study of applying our new algorithm to this quantum bit-flip network tomography model in Section V. The reason for choosing the simple quantum bit-flip star network is mainly because the current research on QNT is still quite limited, and the QNT protocols for the bit-flip star are already state-of-the-art [10, 13]. We emphasize that this paper focuses on devising an online experimental design algorithm to optimize network tomography procedure. Once new QNT protocols are proposed for more practical quantum network infrastructures, one can apply our algorithm to optimize the probe allocation over these protocols as well.

C. Preliminaries of Optimal Experimental Design

Fisher Information Matrix and Cramér-Rao Bound. The Fisher information matrix (FIM) of probing the network via probe $m \in \mathcal{M}$ regarding link parameters $\boldsymbol{\mu}$ is as $\boldsymbol{I}_m(\boldsymbol{\mu}) := \left\{ \mathbb{E}_{\boldsymbol{Y}_m} \left[\frac{\partial \log f_m(\boldsymbol{Y}_m; \boldsymbol{\mu})}{\partial \mu_\ell} \frac{\partial \log f_m(\boldsymbol{Y}_m; \boldsymbol{\mu})}{\partial \mu_{\ell'}} \middle| \boldsymbol{\mu} \right] \right\}_{(\ell, \ell') \in \mathcal{L} \times \mathcal{L}}$. Let $\boldsymbol{\phi} \in \Delta^{M-1}$ denote the allocation ratio (probability distribution) of the number of times of each of the M probes S_m over the total number of times of probes, where Δ^{M-1} is the probability simplex in $\mathbb{R}^M$. Let $\boldsymbol{I}(\boldsymbol{\mu}; \boldsymbol{\phi})$ denote the expected FIM of a single probe chosen according to distribution $\boldsymbol{\phi}$ with respect to the network parameters $\boldsymbol{\mu}$. Given the linearity of

[2]We use Dirac notation for quantum states sparingly in this paper for rigor. However, understanding these notations is not necessary for following the content. We refer interested readers to Nielsen and Chuang [26].

Procedure 1 Online Network Tomography

```
Input: 𝓜 (set of probes), 𝓝 (set of nodes), 𝓛 (set of links),
    T (number of decision rounds), F (optimal experimental
    design criterion)
for each time slot t = 1, 2, ..., T do
    Select a probe m_t from set 𝓜 to perform the network
    tomography once
    Observe Y_{m_t,t} = (Y_{n,t})_{n∈𝓝_{m_t}} from all destination
    nodes n ∈ 𝓝_{m_t,t} of the probe m_t
    Update the parameters and estimates μ̂_t
end for
Output: φ̂* (estimated optimal allocation)
```

FIM, we have $\boldsymbol{I}(\boldsymbol{\mu};\boldsymbol{\phi}) = \sum_{m=1}^{M} \phi_m \boldsymbol{I}_m(\boldsymbol{\mu})$. The Cramér-Rao bound (CRB) states that the covariance matrix of any unbiased estimator $\hat{\boldsymbol{\mu}} := (\hat{\mu}_\ell)_{\ell\in\mathcal{L}}$ of the network parameters $\boldsymbol{\mu}$ (under allocation $\boldsymbol{\phi}$) is lower bounded by the inverse of the FIM, i.e., $\mathrm{Cov}(\hat{\boldsymbol{\mu}}) \succcurlyeq \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi})$ [27], where $\boldsymbol{A} \succcurlyeq \boldsymbol{B}$ denotes that $\boldsymbol{A}-\boldsymbol{B}$ is a positive semi-definite matrix. In particular, when the MLE estimator $\hat{\mu}_{\text{MLE}}$ is unbiased, it achieves the CRB, i.e., $\mathrm{Cov}(\hat{\boldsymbol{\mu}}_{\text{MLE}}) = \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi})$.

Let $F(\boldsymbol{\mu};\boldsymbol{\phi})$ denote a general optimal experimental design criterion and denote $\boldsymbol{\phi}^*(\boldsymbol{\mu}) := \arg\min_{\boldsymbol{\phi}} F(\boldsymbol{\mu};\boldsymbol{\phi})$ as the optimal allocation function given parameter $\boldsymbol{\mu}$ and OED criterion F.[3] For example, in the A-optimal scenario [11], we have $F(\boldsymbol{\mu};\boldsymbol{\phi}) = \mathrm{tr}\big(\boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi})\big)$. This quantity is a lower bound for the average of the variances of the estimated link parameters $\sum_{\ell\in\mathcal{L}} \mathrm{Var}(\hat{\mu}_\ell)$, when the estimator is unbiased.

D. Online Experimental Design and Objective Formulation

As tomography aims to reveal the values of the unknown network parameters $\boldsymbol{\mu}$, we formulate the task as an online experimental design problem with a total of $T \in \mathbb{N}^+$ decision rounds. In each decision round $t \leqslant T$, the learner selects one tomography probe m_t and then observes the received signals $\boldsymbol{Y}_{m_t,t}$ in the destination nodes $n \in \mathcal{N}_{m_t}$ generated by this probe. The learner then updates the estimates of the network parameters based on the received signals until the end of the T decision rounds. Let $S_{m,t}$ denote the number of times that a specific algorithm performs probe m up to time slot t, and let $\boldsymbol{\phi}_t = (S_{m,t}/t)_{m\in\mathcal{M}}$ denote the actual allocation ratio of an algorithm at time step t. The online experimental design problem is summarized in Procedure 1.

As the goal of optimal experimental design (OED) is to minimize function $F(\boldsymbol{\mu};\boldsymbol{\phi})$ in terms of the allocation $\boldsymbol{\phi}$, we take the difference between $F(\boldsymbol{\mu};\boldsymbol{\phi}_T)$ of the concerned algorithm and the minimum $F(\boldsymbol{\mu};\boldsymbol{\phi}^*)$ with optimal allocation $\boldsymbol{\phi}^*$ as our objective, that is, $R_T := F(\boldsymbol{\mu};\boldsymbol{\phi}_T) - F(\boldsymbol{\mu};\boldsymbol{\phi}^*)$. Following the online learning literature, we refer to this metric as *regret* since it describes the gap between the performed algorithm and the hindsight optimal.[4] It is important to note that it differs from cumulative regret, which is the sum over all T rounds. The cumulative regret is unsuitable in our setting since $\boldsymbol{\phi}_T$ already considers all probes up to the end time T. Our regret measures an algorithm's convergence rate: The smaller the regret, the better the algorithm performs.

[3] We slightly abuse the $\boldsymbol{\phi}^*$ notation without the input ($\boldsymbol{\mu}$) to denote the optimal allocation based on the actual parameters.

[4] Some literature may refer to this regret definition as *simple regret*, the performance gap at the end of the online learning process, to differentiate from the cumulative regret.

III. OPAL: An Online Experimental Design Algorithm for Network Tomography

An online experimental design algorithm aims is to estimate link parameters $\boldsymbol{\mu}$ and approach the optimal allocation $\boldsymbol{\phi}^*$ as close as possible. However, without knowledge of link parameters $\boldsymbol{\mu}$ a priori, it is challenging for an online algorithm to achieve the exactly optimal allocation $\boldsymbol{\phi}^*$ during a finite number of rounds, T. This difficulty is due to that the final allocation ratio $\boldsymbol{\phi}_T$ is the accumulation of its whole probe sequence over all T decisions, and many of the early probe decisions, accounted for in the final ratio $\boldsymbol{\phi}_T$, are made without accurate estimates of link parameters $\boldsymbol{\mu}$. Therefore, the actual allocation $\boldsymbol{\phi}_T$ will inevitably deviate from $\boldsymbol{\phi}^*$.

We propose an online tomography algorithm that dynamically chases the estimated optimal probe allocation, called Online Probe Allocation (OPAL). The key idea of OPAL is inspired by the "chasing the (theoretical) optimal bound" idea used in multi-armed bandits literature as exemplified by the Track-and-Stop algorithm achieving the optimal sample complexity for best arm identification [12].

While we use the high-level idea of "chasing" from prior literature, the design of OPAL for network tomography poses *several unique challenges* that make it impossible to directly apply existing algorithmic designs, thus requiring a substantially different analysis than previous bandit algorithms [12]. The main difference is that the feedback mechanism in network tomography is more complex: one scalar observation (e.g., obtained by a unicast probe) can be non-linearly related to multiple link parameters along a path, and a vector observation (e.g., obtained by a multicast probe) may depend on even more link parameters through a set of non-linear relations. Indeed, this non-linearity is far more involved than the simple feedback mechanism present in bandit learning problems, where one scalar feedback usually corresponds precisely to one unknown parameter or a linear combination of unknown parameters, and independence across observations is often assumed. The non-linearity and correlation in the feedback of network tomography make the problem more challenging, requiring a different algorithmic design and analysis.

A. OPAL *Algorithm Design*

OPAL (Algorithm 2) consists of two phases: (a). an initial sampling phase (Line 3) and (b). a chasing phase (Lines 5-7). The initial phase takes initial sample sizes $\mathcal{S}_0 = (S_{m,0})_{m\in\mathcal{M}}$ as input and performs each probe $m \in \mathcal{M}$ accordingly (Lines 2–3). Collecting these $T_0 := \sum_{m\in\mathcal{M}} S_{m,0}$ initial samples ensures that the first estimates $\hat{\boldsymbol{\mu}}_{T_0}$ for link parameters, Line 5 are not too far away from the true parameters $\boldsymbol{\mu}$. The input initial samples $\boldsymbol{S}_0$ can either be set universally as fixed values for guaranteed good performance (see Section IV-B) or

Algorithm 2 Optimal Probe Allocation (OPAL)

Input: $\mathcal{M}$ (set of probes), $\mathcal{N}$ (set of nodes), $\mathcal{L}$ (set of links), T (number of decision rounds), F (optimal experimental design criterion), MLE (estimator), $\boldsymbol{S}_0 := (S_{m,0})_{m\in\mathcal{M}}$ (initial sample sizes)

Initialize: $\mathcal{H}_0 \leftarrow \emptyset$ (history), $\hat{\boldsymbol{\mu}}_0 \leftarrow \mathbf{0}$ (parameter estimates), $T_0 \leftarrow \sum_{m\in\mathcal{M}} S_{m,0}$ (initial length)

```
1:  for each time step t = 1, 2, ..., T do
2:      if ∃m ∈ M : S_{m,t} < S_{m,0} then        ▷ Initial phase
3:          m_t ← randomly pick from {m∈M : S_{m,t} < S_{m,0}}
4:      else                                      ▷ Chasing phase
5:          μ̂_t ← MLE(H_{t-1})           ▷ Update estimates μ̂_t
6:          φ̂*_t ← arg min_φ F(μ̂_t; φ)
                                  ▷ Estimated optimal allocation
7:          m_t ← arg max_{m∈M} φ̂*_{m,t} − (S_{m,t}/t)
                             ▷ Chase the estimated allocation
8:      end if
9:      Perform probe m_t and observe signals Y_{m_t,t} =
        (Y_{n,t})_{n∈N_{m_t}} for its destination nodes
10:     H_t ← H_{t-1} ∪ {(m_t, Y_{m_t,t})}, S_{m_t,t} ← S_{m_t,t-1} + 1
        and S_{m,t} ← S_{m,t-1} for all m ≠ m_t
11: end for
```

Output: $\hat{\boldsymbol{\mu}}_{T+1} \leftarrow \texttt{MLE}(\mathcal{H}_T)$ (final estimates) and $\hat{\boldsymbol{\phi}}_{T+1} \leftarrow \arg\max_{\boldsymbol{\phi}} F(\hat{\boldsymbol{\mu}}_{T+1}; \boldsymbol{\phi})$ (final allocation)

determined by the specific network tomography task and prior knowledge of the network (see Section V-A).

The chasing phase proceeds step by step for each of the remaining time steps $t = T_0 + 1, T_0 + 2, \dots, T$. Denote by $S_{m,t}$ the number of times that probe m is performed up to and including time t, and use vector $\boldsymbol{S}_t := (S_{m,t})_{m\in\mathcal{M}}$ to represent all the sample sizes. For each time step $t > T_0$, the algorithm first updates the MLE estimates $\hat{\boldsymbol{\mu}}_t$ of the network parameters using the latest history $\mathcal{H}_{t-1} = \{(m_s, \boldsymbol{Y}_{m_s,s})\}_{s=1}^{t-1}$ (Line 5). With the updated estimates $\hat{\boldsymbol{\mu}}_t$, the algorithm generates an *estimated optimal allocation* $\hat{\boldsymbol{\phi}}_t^*$ based on the latest link parameter estimates $\hat{\boldsymbol{\mu}}_t$ (Line 6). Then, the algorithm subtracts the actual allocation $\boldsymbol{S}_t/t$ from the estimated optimal allocation $\hat{\boldsymbol{\phi}}_t^*$ element-wise (both are M-entry vectors), where the (possibly negative) entries of the output vector represent the inadequacies of the allocated fractions to each probe. Last, the algorithm performs the probe m_t with the worst (highest) allocation inadequacy $\hat{\phi}_{m,t}^* - (S_{m,t}/t)$ once (Line 7)—*chasing the estimated optimal allocation* $\hat{\boldsymbol{\phi}}_t^*$—to collect a new observation, and updates the sample sizes $\boldsymbol{S}_t$ (Lines 9–10). After both phases, the algorithm outputs the final MLE estimates $\hat{\boldsymbol{\mu}}_{T+1}$ and final estimated optimal allocation $\hat{\boldsymbol{\phi}}_{T+1}$.

IV. ANALYSIS

We start with two critical conditions for the network tomography tasks (Section IV-A) and then characterizes the regret (Section IV-B) based on these conditions.

A. Conditions on Network Tomography

We present two general conditions for the network tomography tasks, which are crucial for the regret analysis of OPAL. Both conditions are natural and typically fulfilled in practice. Later in Section V-A, we verify both conditions in the quantum network tomography case studies.

Condition 1 (Lipschitz Continuity)**.** *The optimal experimental design criterion $F(\boldsymbol{\mu}; \boldsymbol{\phi})$ is Lipschitz continuous regarding the allocation $\boldsymbol{\phi}$. That is, for constant $\alpha > 0$ and any two allocations $\boldsymbol{\phi}, \boldsymbol{\phi}' \in \mathcal{F}$ where $\mathcal{F} \subseteq \Delta^{M-1}$ is the set of all feasible allocations that support identifiability,*[5] *we have*

$$F(\boldsymbol{\mu}; \boldsymbol{\phi}) - F(\boldsymbol{\mu}; \boldsymbol{\phi}') \leqslant \alpha \|\boldsymbol{\phi} - \boldsymbol{\phi}'\|_\infty.$$

Further, the optimal allocation function $\boldsymbol{\phi}^*(\boldsymbol{\mu})$ is Lipschitz continuous to $\boldsymbol{\mu}$. That is, for some constant $\beta > 0$ and any two sets of feasible network parameters $\boldsymbol{\mu}, \boldsymbol{\mu}'$, we have

$$\|\boldsymbol{\phi}^*(\boldsymbol{\mu}) - \boldsymbol{\phi}^*(\boldsymbol{\mu}')\|_\infty \leqslant \beta \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty.$$

Condition 2 (Estimator Concentration / Confidence Interval)**.** *At any decision round t, the confidence interval (with confidence $1 - \delta$ for parameter $\delta \in (0, 1)$) for any parameter μ_ℓ of a link $\ell \in \mathcal{L}$ is an interval centered at its MLE estimate $\hat{\mu}_{\ell,t}$ with radius $\sum_{m\in\mathcal{M}} c_{\ell,m} \left(\log \delta^{-1}/S_{m,t}\right)^{\gamma_{\ell,m}}$, where $c_{\ell,m} \geqslant 0, \gamma_{\ell,m} > 0$ are parameters depending on the network and tomography probes, and $S_{m,t}$ is the number of times that probe m is performed up to round t.*

Lipschitz continuity (Condition 1) is a common condition in practical scenarios, as it implies that the network tomography problem is well-behaved, and the optimal allocation $\boldsymbol{\phi}^*$ is stable with respect to the network parameters $\boldsymbol{\mu}$. Condition 2 describes the concentration rate of the MLE (or any other) estimator in terms of the sample times $S_{m,t}$ of all probes m. With the central limit theorem, one can derive an asymptotic version of the confidence interval with radius $c_0\sqrt{1/S_{m,t}}$ [2, Theorem 4], for some positive constant c_0. This implies that the concentration rate $\gamma_{\ell,m}$ of probe m for the parameter of link ℓ approaches $1/2$ for large sample sizes.

B. Regret (Convergence Rate) of OPAL

In this section, we present the main theorem that characterizes the regret of OPAL. We recall the probe allocation ratio of the algorithm at time t as $\boldsymbol{\phi}_t := (S_{m,t}/t)_{m\in\mathcal{M}}$.

Theorem 1. *Given Conditions 1 and 2 with $\delta = 1/(LT^2)$, time horizon $T > 0$, parameter $\xi \in (0, 1)$, and initial probe times $S_{m,0} = \xi T, \forall m \in \mathcal{M}$, OPAL satisfies,*

$$\begin{aligned} R_T = F(\boldsymbol{\mu}; \boldsymbol{\phi}_T) - F(\boldsymbol{\mu}; \boldsymbol{\phi}^*) \leqslant C \left(\frac{\log(LT)}{\xi T}\right)^{\gamma_{\min}} \\ + \alpha\xi \mathbb{1}\left\{\exists m \in \mathcal{M} : \phi_m^* < \xi + C\left(\frac{\log(LT)}{\xi T}\right)^{\gamma_{\min}}\right\}, \end{aligned} \tag{1}$$

with probability of at least $1 - 1/T$, where $C = c_0 \cdot \alpha\beta c_{\max}$ is for some constant $c_0 > 0$ that depends on the probing experiments and network, independent of time horizon T. Parameters α and β come from Condition 1,

[5]The set $\mathcal{F}$ contains all interior points and perhaps a part of the boundary points of the simplex set Δ^{M-1}.

parameters $c_{\max} := \max_{\ell\in\mathcal{L}} \sum_{m\in\mathcal{M}} c_{\ell,m}$ *and* $\gamma_{\min} := \min_{(\ell,m)\in\mathcal{L}\times\mathcal{M}:\gamma_{\ell,m}>0} \gamma_{\ell,m}$ *come from Condition 2.*

Implications of Theorem 1. Theorem 1 characterizes the regret of OPAL. We defer to full version of this paper for the complete proof of Theorem 1. To better understand the regret, we discuss several special cases.

First, for many network tomography tasks, even without knowing the actual network channel parameters, one can select a compact set of probing experiments $\mathcal{M}$ such that each probe in the set is needed at least a minimum fraction of the time—that is, $\phi^*_m > \phi_{\text{threshold}}$ for all probes $m \in \mathcal{M}$—to identify the network parameters. One can set ξ in Theorem 1 to this threshold, and then the second term in the RHS of (1) vanishes, leading to a $\tilde{O}\left(T^{-\gamma_{\min}}\right)$ regret.

Second, one can always set $\xi = T^{-\frac{\gamma_{\min}}{1+\gamma_{\min}}}$ in Theorem 1, guaranteeing a decent $\tilde{O}(T^{-\frac{\gamma_{\min}}{1+\gamma_{\min}}})$ regret, independent of whether the second term of the regret in (1) can be canceled or not. For example, the second term cannot be canceled by any constant $\xi > 0$ in a scenario where the optimal allocation $\boldsymbol{\phi}^*$ contains some zero entries, i.e., $\phi^*_m = 0$ for some probe m (e.g., the case study in Section V-A). That is, when there are unavoidable inferior probes, letting $\xi = T^{-\frac{\gamma_{\min}}{1+\gamma_{\min}}}$ is a plausible choice, and OPAL yields a $\tilde{O}(T^{-\frac{\gamma_{\min}}{1+\gamma_{\min}}})$ regret, slightly worse than the $-\gamma_{\min}$ rate.

Third, from an asymptotic perspective (for large sample sizes) and by the central limit theorem, the concentration rate of Condition 2 is $\gamma_{\ell,m} = 1/2$ for all probes m and links ℓ. Assuming the second term in (1) is canceled, the regret of OPAL becomes $\tilde{O}(T^{-\frac{1}{2}})$. This rate is near-optimal in the sense that, even for the online linear regression (simpler than our network tomography task), minimizing the regret for A-optimal design criterion (with the number of probes greater than that of links, $M > L$) is known to suffer at least a $\Omega(T^{-1/2})$ regret [24, Theorem 4].

V. OPAL FOR QUANTUM NETWORK TOMOGRAPHY

This section presents a case study on the bit-flip quantum channel tomography in the quantum star network. Section V-A verifies the conditions outlined in Section IV-A and presents the specific regret of OPAL. Subsequently, in Section V-B, we provide empirical results for the quantum scenario, validating the superiority of OPAL over existing methods.

A. Analysis for Quantum Bit-Flip Star Network Multicast

We consider an L-link star network for A-optimal experimental design, where each link $\ell \in \mathcal{L}$ is a *quantum* bit-flip channel with an unknown bit-flip probability μ_ℓ. As no unicast protocol in quantum network tomography has been proposed for this setting, we consider the root-independent (RI) multicast protocol devised by De Andrade et al. [13]. The probing experiment set $\mathcal{M}$ consists of $M = L$ multicast probes, each selecting one link $\ell \in \mathcal{L}$ as the root, multicasting quantum states to the remaining $L-1$ links, and measuring the received qubits at these links. Denote $A_{m,\ell}$ as the total non-flip counts from the measurements of probe m for each link $\ell \neq m$. More details of the bit-flip channel and the RI multicast protocol are provided in Section II-B and Figure 1.

Algorithm 3 A-Optimal Allocation for Bit-Flip Tomograph in Quantum Star Networks

Input: link parameters $\boldsymbol{\mu}$ and candidate probe set $\mathcal{M}_{\text{opt}} \leftarrow \{1,2,3,\dots,M\}$ $(M = L)$

1: **while** $\sum_{\ell\in\mathcal{M}_{\text{opt}}} \sqrt{(1-\mu_\ell)\mu_\ell} - (|\mathcal{M}_{\text{opt}}| - 1)\max_{\ell'\in\mathcal{M}_{\text{opt}}} \sqrt{(1-\mu_{\ell'})\mu_{\ell'}} < 0$ **do**

2: $\quad \mathcal{M}_{\text{opt}} \leftarrow \mathcal{M}_{\text{opt}} \setminus \{\arg\max_{\ell'\in\mathcal{M}_{\text{opt}}} \sqrt{(1-\mu_{\ell'})\mu_{\ell'}}\}$

3: **end while**

Output: $\phi^*_\ell = \mathbb{1}\{\ell\in\mathcal{M}_{\text{opt}}\}\times$

$$\frac{\sum_{\ell'\in\mathcal{M}_{\text{opt}}} \sqrt{(1-\mu_{\ell'})\mu_{\ell'}} - (|\mathcal{M}_{\text{opt}}|-1)\sqrt{(1-\mu_\ell)\mu_\ell}}{\sum_{\ell'\in\mathcal{M}_{\text{opt}}} \sqrt{(1-\mu_{\ell'})\mu_{\ell'}}} \tag{2}$$

Lipschitz Continuity of the A-Optimal Design. We first derive the trace of the inverse of the Fisher information matrix for the quantum star network with bit-flip quantum channels as follows, $F(\boldsymbol{\mu};\boldsymbol{\phi}) = \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi}) = \sum_{m\in\mathcal{M}} \frac{(1-\mu_m)\mu_m}{1-\phi_m}$. This expression verifies the Lipschitz continuity of the A-optimal design criterion (first part of Condition 1). Next, by applying the method of Lagrange multipliers, we derive the A-optimal solution $\boldsymbol{\phi}^*$ for the quantum star network with bit-flip quantum channels. The optimal allocation depends on the link parameters, as presented in (2) of Algorithm 3. Specifically, one first sorts in descending order the links by the value of $\sqrt{(1-\mu_\ell)\mu_\ell}$, and then iteratively eliminates the link with the largest value of $\sqrt{(1-\mu_\ell)\mu_\ell}$ until the condition in Line 1 is not satisfied. The optimal allocation is then given by (2) for the remaining links in the candidate set $\mathcal{M}_{\text{opt}}$.

While the expression for the optimal allocation $\boldsymbol{\phi}^*$ in (2) is piecewise, these pieces are continuous at all critical points with respect to the link parameters—altering the inequality in Line 1 from "less than" to "less than or equal to" does not change the solution. Hence, the Lipschitz continuity of the A-optimal design (second part of Condition 1) is verified.

MLE Estimator for Link Parameters. The observations at each link ℓ from the RI multicast follow a Bernoulli distribution with success probability μ_ℓ [13, Table II and Eq. (17)]. Hence, the maximum likelihood estimator (MLE) for the link parameter μ_ℓ is calculated by the observations from all other RI probes $m \in \mathcal{M}\setminus\{\ell\}$ as follows:

$$\hat{\mu}_\ell = \frac{\sum_{m\in\mathcal{M}\setminus\{\ell\}} A_{m,\ell}}{\sum_{m\in\mathcal{M}\setminus\{\ell\}} S_m}, \quad \text{for all links } \ell \in \mathcal{L}, \tag{3}$$

where we recall that S_m is the total number of qubits sent by probe m, and $A_{m,\ell}$ is the total number of qubits not flipped at link ℓ from the measurements of probe m.

Confidence Interval of MLE Estimator. By Hoeffding's inequality, the confidence interval for the MLE estimator in (3)

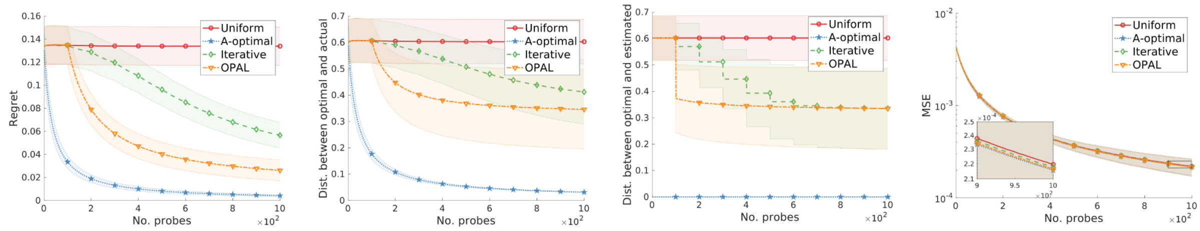

(a) Regret: $R_t = F(\boldsymbol{\mu},\boldsymbol{\phi}_t) - F(\boldsymbol{\mu},\boldsymbol{\phi}^*)$ (b) Dist. of actual: $\|\boldsymbol{\phi}^* - \boldsymbol{\phi}_t\|_2$ (c) Dist. of estimated: $\|\boldsymbol{\phi}^* - \hat{\boldsymbol{\phi}}_t\|_2$ (d) MSE: $\mathbb{E}[\|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_2^2]$

Fig. 2: Comparison of algorithms for 40-node quantum star networks

is given by, for some constant $C_2 > 0$,

$$\mu_\ell \in \left(\hat{\mu}_\ell - C_2\sqrt{\frac{\log \delta^{-1}}{\sum_{m\in\mathcal{M}\setminus\{\ell\}} S_m}}, \hat{\mu}_\ell + C_2\sqrt{\frac{\log \delta^{-1}}{\sum_{m\in\mathcal{M}\setminus\{\ell\}} S_m}} \right)$$

with confidence level $1-\delta$, for all links $\ell \in \mathcal{L}$. Therefore, Condition 2 is verified with all $\gamma_{m,\ell} = \frac{1}{2}$ and thus $\gamma_{\min} = \frac{1}{2}$.

Regret of OPAL on the Bit-Flip Quantum Star Network with RI Multicast. With the Lipschitz continuity of the A-optimal design criterion and the finite confidence interval with $\gamma_{\min} = \frac{1}{2}$, we apply Theorem 1 to derive the regret of OPAL on the quantum star network with bit-flip quantum channels as follows, $R_T = \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi}_T) - \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi}^*) \leqslant C\,(\log T/\xi T)^{\frac{1}{2}} + \alpha\xi$. Because the A-optimal allocation in (2) inherently includes probes with a zero allocation ratio, i.e., $\phi_m^* = 0$ for some m, the second term remains non-zero. This aligns with the intuition that it is difficult for an algorithm to avoid performing these ineffective probes without knowledge of the exact link parameters a priori, leading to a regret (convergence rate) slower than $\tilde{O}(T^{-\gamma_{\min}})$.

By selecting $\xi = T^{-\frac{1}{3}}$, we obtain the worst-case convergence (regret) guarantee of OPAL on the quantum bit-flip star network under the RI multicast, expressed as, $R_T = \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi}_T) - \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi}^*) \leqslant O((\log T/T)^{\frac{1}{3}})$.

Complexity of OPAL for Quantum Bit-Flip Star Network. The major computational cost of applying OPAL to the bit-flip star networks is from executing Algorithm 3 (Line 6 in OPAL), whose time complexity is $O(L)$.

B. Simulations for Quantum Network Tomography

This section reports the empirical performance of OPAL algorithm in a 40-node quantum bit-flip star network tomography task: one center node connects to 39 other root nodes via bit-flip quantum channels (links). The set of probing experiments consists of all 39 root-independent (RI) multicast probes, each selecting the end node of one link as the root node and the other 38 end nodes as the multicast receivers. We consider three baselines: an iterative algorithm proposed in [19, Algorithm 2], the A-Optimal allocation, and a uniform allocation. The iterative algorithm is a state-of-the-art experimental design solution for network tomography without prior knowledge of the link parameters. It performs probes in batch (each consisting of 100 time steps) according to its iteratively updated policy. The A-optimal allocation uses $\boldsymbol{\phi}^*$, minimizes the A-optimal experimental design, and represents an ideal solution with full knowledge of the link parameters, and the uniform allocation represents a simple baseline.

The initial sampling phase of OPAL is set to $T_0 = 0.1T$ rounds. The experiment is conducted with 20 scenarios, each with a set of randomly generated bit-flip probabilities $\boldsymbol{\mu}$ from a uniform distribution over $(0.1, 1)$ for the 40-node quantum star network. For each scenario, we run 100 Monte Carlo simulations and take the average as the performance of this scenario. We report the average performance of these 20 scenarios as solid lines, and the shaded areas represent their standard deviation across these scenarios.

The RI protocol considered in the experiment is specifically developed for quantum bit-flip star network, and it is already the most advanced in the QNT literature [13]. Therefore, other general QNT experiments under other more complex quantum channels and network topologies are invalid here. This quantum star network experiment serves as a proof-of-concept for the applicability of OPAL to QNT. Once more advanced QNT protocols are developed for more general network topologies, one can apply OPAL to optimizing the probe allocation over them in more general scenarios.

Figure 2 presents the four performance metrics of algorithms for bit-flip tomography in a 40-node quantum star network. The results demonstrate the superior performance of OPAL in the quantum network tomography scenario. Applying OPAL to QNT protocols for more complex quantum network topologies is expected to yield more significant improvements, which is infeasible to verify in this paper due to the lack of existing QNT protocols for general topologies.

VI. CONCLUSION

This paper introduces a novel and general online experimental design algorithm for network tomography, termed online probe allocation (OPAL). Theoretically, OPAL is the first algorithm to offer rigorous regret guarantees for network tomography. We establish these guarantees by identifying two critical conditions: Lipschitz continuity and confidence interval concentration. On the practical side, we validate these theoretical conditions in quantum bit-flip multicast networks, representing a key use case for network tomography. Empirically, we illustrate the superior performance of OPAL compared to existing methods in quantum network tomography scenarios.

ACKNOWLEDGMENT

This research is supported in part by the NSF grant CNS-2402861, NSF-ERC Center for Quantum Networks grant EEC-1941583. It is also sponsored in part by PiQSci project of Advanced Scientific Computing Research program, U.S. Department of Energy, and is performed at Oak Ridge National Laboratory managed by UT-Battelle, LLC for U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Ting He is supported by the NSF grant CNS-2106294. The work of John C.S. Lui was supported in part by the RGC SRFS2122-4S02 and GRF 14202925.

REFERENCES

[1] T. He, L. Ma, A. Swami, and D. Towsley, *Network tomography: identifiability, measurement design, and network state inference*. Cambridge University Press, 2021.

[2] R. Cáceres, N. G. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal loss characteristics," *IEEE Transactions on Information theory*, vol. 45, no. 7, pp. 2462–2480, 1999.

[3] N. G. Duffield and F. Lo Presti, "Multicast inference of packet delay variance at interior network links," in *Proceedings IEEE INFOCOM 2000*, vol. 3. IEEE, 2000, pp. 1351–1360.

[4] F. Lo Presti, N. G. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Transactions On Networking*, vol. 10, no. 6, pp. 761–775, 2002.

[5] B. Xi, G. Michailidis, and V. N. Nair, "Estimating network loss rates using active tomography," *Journal of the American Statistical Association*, vol. 101, no. 476, pp. 1430–1448, 2006.

[6] A. G. Barto, R. S. Sutton, and C. Watkins, *Learning and sequential decision making*. University of Massachusetts Amherst, MA, 1989, vol. 89.

[7] P. Auer, "Finite-time analysis of the multiarmed bandit problem," 2002.

[8] L. Li, W. Chu, J. Langford, and R. E. Schapire, "A contextual-bandit approach to personalized news article recommendation," in *Proceedings of the 19th international conference on World wide web*, 2010, pp. 661–670.

[9] W. Chen, Y. Wang, and Y. Yuan, "Combinatorial multi-armed bandit: General framework and applications," in *International conference on machine learning*. PMLR, 2013, pp. 151–159.

[10] M. G. De Andrade, J. Diaz, J. Navas, S. Guha, I. Montaño, B. Smith, M. Raymer, and D. Towsley, "Quantum network tomography with multi-party state distribution," in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2022, pp. 400–409.

[11] F. Pukelsheim, *Optimal design of experiments*. SIAM, 2006.

[12] A. Garivier and E. Kaufmann, "Optimal best arm identification with fixed confidence," in *Conference on Learning Theory*. PMLR, 2016, pp. 998–1027.

[13] M. G. De Andrade, J. Navas, I. Montaño, 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. 1. IEEE, 2023, pp. 1260–1270.

[14] M. Coates, A. O. Hero III, R. Nowak, and B. Yu, "Internet tomography," *IEEE Signal processing magazine*, vol. 19, no. 3, pp. 47–65, 2002.

[15] T. Bu, N. Duffield, F. Lo Presti, and D. Towsley, "Network tomography on general topologies," *ACM SIGMETRICS Performance Evaluation Review*, vol. 30, no. 1, pp. 21–30, 2002.

[16] Y. Tsang, M. Coates, and R. D. Nowak, "Network delay tomography," *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2125–2136, 2003.

[17] A. Gopalan and S. Ramasubramanian, "On identifying additive link metrics using linearly independent cycles and paths," *IEEE/ACM Transactions on Networking*, vol. 20, no. 3, pp. 906–916, 2011.

[18] L. Ma, T. He, A. Swami, D. Towsley, K. K. Leung, and J. Lowe, "Node failure localization via network tomography," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014, pp. 195–208.

[19] T. He, C. Liu, A. Swami, D. Towsley, T. Salonidis, A. I. Bejan, and P. Yu, "Fisher information-based experiment design for network tomography," *ACM SIGMETRICS*, vol. 43, no. 1, pp. 389–402, 2015.

[20] L. M. Haines, I. Perevozskaya, and W. F. Rosenberger, "Bayesian optimal designs for phase i clinical trials," *Biometrics*, vol. 59, no. 3, pp. 591–600, 2003.

[21] J. Gilman, L. Walls, L. Bandiera, and F. Menolascina, "Statistical design of experiments for synthetic biology," *ACS synthetic biology*, vol. 10, no. 1, pp. 1–18, 2021.

[22] R. Leardi, "Experimental design in chemistry: A tutorial," *Analytica chimica acta*, vol. 652, no. 1-2, pp. 161–172, 2009.

[23] B. Kveton, M. J. Amjad, C. Diot, D. Konomis, A. Soule, and X. Yang, "Optimal probing with statistical guarantees for network monitoring at scale," *Computer Communications*, vol. 192, pp. 119–131, 2022.

[24] X. Fontaine, P. Perrault, M. Valko, and V. Perchet, "Online A-optimal design and active linear regression," in *International Conference on Machine Learning*. PMLR, 2021, pp. 3374–3383.

[25] Z. Allen-Zhu, Y. Li, A. Singh, and Y. Wang, "Near-optimal discrete optimization for experiment design: A regret minimization approach," *Mathematical Programming*, vol. 186, pp. 439–478, 2021.

[26] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge university press, 2010.

[27] E. L. Lehmann and G. Casella, *Theory of point estimation*. Springer Science & Business Media, 2006.