Online Experimental Design for Network Tomography

Xuchuang Wang
University of Massachusetts Amherst

Yu-Zhen Janice Chen
University of Massachusetts Amherst

Matheus Guedes de Andrade
University of Massachusetts Amherst

Mohammad Hajiesmaili
University of Massachusetts Amherst

John C.S. Lui
The Chinese University of Hong Kong

Ting He
Pennsylvania State University

Don Towsley
University of Massachusetts Amherst

1. INTRODUCTION

Network tomography [5] is an essential approach 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 are difficult to access. A stochastic network tomography problem consists of a network with a set $\mathcal{L}$ of $L \coloneqq |\mathcal{L}|$ 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 $M \coloneqq |\mathcal{M}|$ 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] and loss network tomography [2]. That is, after obtaining these stochastic observations from probes, prior works aim to devise good *static* estimators to infer the network parameters. While the 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) [1] 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, 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. This complex feedback mechanism makes the design of the online algorithm and the analysis of its regret challenging.

2. MODELING AND FORMULATION

We formulate a general network tomography task as an online experimental design problem. Denote $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ as a network topology with a set of $N\in \mathbb{N}^+$ nodes $\mathcal{N} \coloneqq \{1, 2, \dots, N\}$ and a set of $L \in \mathbb{N}^+$ links $\mathcal{L} \coloneqq \{1, 2, \dots, L\}$. Each link $\ell \in \mathcal{L}$ is associated with a stochastic model X_ℓ representing the characteristic of the link, e.g., a Bernoulli random variable for signal loss [2]. 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 $\boldsymbol{\mu} \coloneqq (\mu_\ell)_{\ell\in\mathcal{L}}$ (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 $\mathcal{M} \coloneqq \{1, 2, \dots, M\}$ (called *probes* later), examples of which are unicast [4] and multicast [2]. Each probe $m \in \mathcal{M}$ is associated with one path or a set of paths consisting of one source node (sender) and a set of destination nodes (receivers) denoted as $\mathcal{N}_m \subseteq \mathcal{N}$, and 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 $f_m(\boldsymbol{Y}_m; \boldsymbol{\mu})$. 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 [4] that the network parameters $\boldsymbol{\mu}$ can be uniquely determined from the feedback of sufficiently many probing experiments in $\mathcal{M}$. Given a set of $T \in \mathbb{N}^+$ probe-feedback pairs $\mathcal{H} \coloneqq \{(m_t, \boldsymbol{Y}_{m_t,t})\}_{t=1}^T$, we define the log-likelihood function as $L(\boldsymbol{\mu}; \mathcal{H}) \coloneqq \sum_{t=1}^T \log f_{m_t}(\boldsymbol{Y}_{m_t,t}; \boldsymbol{\mu})$. The maximum likelihood estimation (MLE) estimator for the network parameters $\boldsymbol{\mu}$ is defined as $\texttt{MLE}(\mathcal{H}) \coloneqq \arg\max_{\boldsymbol{\mu}} L(\boldsymbol{\mu}; \mathcal{H})$.[1]

[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.

Algorithm 1 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), $\texttt{MLE}$ (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 $t = 1, 2, \dots, T$ **do**

2:   **if** $\exists m \in \mathcal{M} : S_{m,t} < S_{m,0}$ **then** ▷ `Initial phase`

3:     $m_t \leftarrow$ randomly pick from $\{m\in\mathcal{M}: S_{m,t} < S_{m,0}\}$

4:   **else** ▷ `Chasing phase`

5:     $\hat{\boldsymbol{\mu}}_t \leftarrow \texttt{MLE}(\mathcal{H}_{t-1})$
▷ `Use MLE to update the estimates` $\hat{\boldsymbol{\mu}}_t$

6:     $\hat{\boldsymbol{\phi}}_t^* \leftarrow \arg\min_{\boldsymbol{\phi}} F(\hat{\boldsymbol{\mu}}_t; \boldsymbol{\phi})$
▷ `Calculate estimated optimal allocation`

7:     $m_t \leftarrow \arg\max_{m\in\mathcal{M}} \hat{\phi}_{m,t}^* - (S_{m,t}/t)$
▷ `Chase the estimated allocation`

8:   **end if**

9:   Perform probe m_t and observe signals $\boldsymbol{Y}_{m_t,t} = (Y_{n,t})_{n\in\mathcal{N}_{m_t}}$ for its destination nodes

10:   $\mathcal{H}_t \leftarrow \mathcal{H}_{t-1} \cup \{(m_t, \boldsymbol{Y}_{m_t,t})\}$, $S_{m_t,t} \leftarrow S_{m_t,t-1}+1$ and $S_{m,t} \leftarrow S_{m,t-1}$ for all $m \neq m_t$

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)

Experimental Design We denote $\boldsymbol{\phi} \in \Delta^{M-1}$ as 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 performing probes $\sum_{m\in\mathcal{M}} S_m$, where Δ^{M-1} is the probability simplex in $\mathbb{R}^M$. Denote $\boldsymbol{I}(\boldsymbol{\mu};\boldsymbol{\phi})$ as the Fisher information matrix (FIM) of probing the network via the mixture of multiple probes according to an allocation $\boldsymbol{\phi}$ with respect to the network parameters $\boldsymbol{\mu}$. We denote $F(\boldsymbol{\mu};\boldsymbol{\phi})$ as 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[2] For example, in A-optimal scenario, we have $F(\boldsymbol{\mu};\boldsymbol{\phi}) = \operatorname{tr}\big(\boldsymbol{I}^{-1}(\boldsymbol{\mu};\boldsymbol{\phi})\big)$, and in D-optimal scenario, we have $F(\boldsymbol{\mu};\boldsymbol{\phi}) = (\det(\boldsymbol{I}(\boldsymbol{\mu};\boldsymbol{\phi})))^{-1}$.

3. OPTIMAL PROBE ALLOCATION

OPAL (presented in Algorithm 1) consists of two phases: (a) an initial sampling phase (Line 3) and (b) a chasing phase (Lines 5-7). The initial phase takes the 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 at Line 5 are not too far away from the true parameters $\boldsymbol{\mu}$.

The chasing phase proceeds for each of the remaining steps $t = T_0+1, T_0+2, \dots, T$. Denote $S_{m,t}$ as the number of times that the 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 inadequacy of the allocated fractions to each probe. Last, the algorithm performs the probe m_t with the worst (highest) allocation inadequacy 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).

The above estimated optimal allocation chasing step is the core of OPAL, which dynamically adjusts the probe allocation based on the estimated optimal allocation $\hat{\boldsymbol{\phi}}_t^*$ at each time step t. After both phases, the algorithm outputs the final MLE estimates $\hat{\boldsymbol{\mu}}_{T+1}$ and final estimated optimal allocation $\hat{\boldsymbol{\phi}}_{T+1}$.

4. ANALYSIS

We present two general conditions for the network tomography problem and then, the main theorem that characterizes the regret of OPAL.

Condition 1 (Lipschitz Continuity)**.** *The optimal experimental design criterion $F(\boldsymbol{\mu};\boldsymbol{\phi})$ is Lipschitz continuous with respect to the allocation $\boldsymbol{\phi}$. That is, for some 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 the parameter identifiability,*[3] *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 with respect 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.*

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

$$\begin{aligned} R_T &= F(\boldsymbol{\mu};\boldsymbol{\phi}_T) - F(\boldsymbol{\mu};\boldsymbol{\phi}^*) \\ &\leqslant \frac{\alpha}{(1-\xi)T} + 4\alpha\beta c_{\max}\left(\frac{\log(LT)}{\xi T}\right)^{\gamma_{\min}} \\ &+ \alpha\xi\mathbb{1}\left\{\exists m \in \mathcal{M} : \phi_m^* < \xi + 4\alpha\beta c_{\max}\left(\frac{\log(LT)}{\xi T}\right)^{\gamma_{\min}}\right\}, \end{aligned}$$

where $C = O(\alpha\beta c_{\max})$ is a constant that depends on the probing experiments and network, independent of time horizon T, α and β are from Condition 1, parameters $c_{\max} :=$

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

[3]The set $\mathcal{F}$ contains all interior points and perhaps a part of the boundary points of the simplex set Δ^{M-1}. For example, in the classical loss star network tomography via unicasts in Section 5, all unicast probes are necessary for identifying the network parameters. Hence, any allocation on the boundary (i.e., allocations contain zero entries) does not correspond to a feasible allocation, and $\mathcal{F} = \Delta^{M-1} \setminus \partial\Delta^{M-1}$.

$\max_{\ell \in \mathcal{L}} \sum_{m \in \mathcal{M}} c_{\ell,m}$ *and* $\gamma_{\min} := \min_{(\ell,m) \in \mathcal{L} \times \mathcal{M}: \gamma_{\ell,m} > 0} \gamma_{\ell,m}$ *are from Condition 2, and* $\mathbb{1}\{\cdot\}$ *is the indicator function.*

5. OPAL FOR UNICAST STAR NETWORK

In this case study, we examine loss tomography in a star network with L links via unicast.

Lipschitz Continuity Based on the derivation in [4, Theorem 6], we compute the trace of the inverse of the Fisher information matrix $\boldsymbol{I}(\boldsymbol{\mu}; \boldsymbol{\phi})$ for the L-link star network as follows,

$$F(\boldsymbol{\mu}; \boldsymbol{\phi}) = \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu}; \boldsymbol{\phi}) = \sum_{m=1}^{M} \frac{1}{\phi_m} A_m(\boldsymbol{\mu}), \tag{1}$$

where $A_m(\boldsymbol{\mu}) := \frac{1-\nu_m}{\nu_m} \sum_{\ell=1}^{L} \mu_\ell^2 \kappa_{\ell,m}$ is a function of the link parameters $\boldsymbol{\mu}$, and $\kappa_{\ell,m} := (\boldsymbol{Q}^{-1})_{\ell,m}$ denotes the (ℓ, m)-entry of the inverse of the measurement matrix $\boldsymbol{Q}$.

By the method of Lagrange multipliers, the A-optimal solution for minimizing $F(\boldsymbol{\mu}; \boldsymbol{\phi})$ is

$$\phi_m^* = \frac{\sqrt{A_m(\boldsymbol{\mu})}}{\sum_{m' \in \mathcal{M}} \sqrt{A_{m'}(\boldsymbol{\mu})}}, \quad \forall m \in \mathcal{M}. \tag{2}$$

As the analytical solutions of the A-optimal design in (1) and (2) only contain basic calculations and are thus differentiable, we verify the Lipschitz continuity in Condition 1.

Estimator Concentration We also prove that the concentration rate of the MLE estimator for the link parameters $\boldsymbol{\mu}$ in the star network as follows,

Theorem 2. *For unicast star network we have, with a probability of at least* $1 - \delta$*,*

$$\mu_\ell \in \left(\hat{\mu}_\ell - \sum_{m \in \mathcal{M}: \kappa_{\ell,m} \neq 0} c_{m,\ell} \sqrt{\frac{\log(\delta^{-1}/M)}{S_m}}, \right. \\ \left. \hat{\mu}_\ell + \sum_{m \in \mathcal{M}: \kappa_{\ell,m} \neq 0} c_{m,\ell} \sqrt{\frac{\log(\delta^{-1}/M)}{S_m}} \right), \tag{3}$$

where the $c_{m,\ell}$ *for each probe* m *and link* ℓ *pair is a constant that depends on the network topology and link parameters.*

This verifies Condition 2 regarding the concentration rate of the MLE estimator for the link parameters in the unicast star network, where the exponent parameters are $\gamma_{m,\ell} \in \{0, \frac{1}{2}\}$, resulting in a minimum nonzero exponent $\gamma_{\min} = \frac{1}{2}$.

Regret of OPAL *on the Unicast Star Network*

With the confidence interval in (3) and $\delta = 1/LT^2$, we know that $O(\log T)$ samples for each unicast probe suffice to estimate the link parameters $\boldsymbol{\mu}$ with relatively good accuracy. Furthermore, utilizing the Lipschitz continuity of the A-optimal design allocation in (2), one can derive a lower bound ξ for the optimal allocation ratio ϕ_m^* for all probes $m \in \mathcal{M}$. Consequently, the regret of OPAL for the A-optimal design under the loss star network is expressed as $R_T = \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu}; \boldsymbol{\phi}_T) - \operatorname{tr} \boldsymbol{I}^{-1}(\boldsymbol{\mu}; \boldsymbol{\phi}^*) = O\left(\left(\frac{\log T}{T}\right)^{\frac{1}{2}} \right)$.

Empirical Evaluation of OPAL

We present numerical simulations to evaluate the performance of OPAL) in the above tomography task. We consider three baselines: an iterative algorithm proposed in [4, Algorithm 2], the A-optimal allocation, and a uniform allocation. The iterative algorithm is batch-based. It performs probes in each batch (consisting of 100 time steps) according to its iteratively updated allocation. The initial sampling phase of OPAL is set to $T_0 = 0.05 \cdot T$ uniformly over all probes. The A-optimal allocation policy follows the $\boldsymbol{\phi}^*$ that minimize the A-optimal experimental design $F(\boldsymbol{\mu}, \boldsymbol{\phi})$ assuming the prior knowledge of link parameters $\boldsymbol{\mu}$, and the uniform allocation policy follows the $\boldsymbol{\phi}_{\text{unif.}} = \boldsymbol{1}/L$. Figure 1 reports the mean square error (MSE), $\mathbb{E}[\|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_2^2]$, of these four algorithms in the a 3-link star network, which OPAL performs as well as the A-optimal and outperforms other baselines.

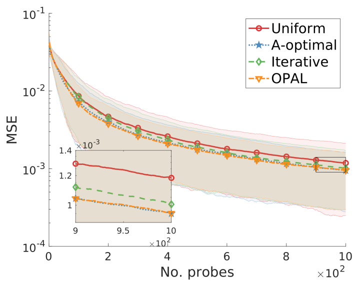

Figure 1: MSE Comparison in 3-Link Star Network

6. 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 classical loss unicast networks, representing key use cases for network tomography. Empirically, we illustrate the superior performance of OPAL compared to existing methods.

7. REFERENCES

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

[2] R. Cáceres, N. G. Duffield, J. Horowitz, and D. Towsley. Multicast-based inference of network-internal loss characteristics. *IEEE Transactions on Information theory*, 45(7):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.*, volume 3, pages 1351–1360. IEEE, 2000.

[4] 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 Performance Evaluation Review*, 43(1):389–402, 2015.

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