Multicast-based Weight Inference in General Network Topologies

Yilei Lin*, Ting He*, Shiqiang Wang†, Kevin Chan‡, and Stephen Pasteris§
*Pennsylvania State University, University Park, PA 16802, USA. Email: {yjl5282,tzh58}@psu.edu
†IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA. Email: wangshiq@us.ibm.com
‡US Army Research Laboratory, Adelphi, MD 20783, USA. Email: kevin.s.chan.civ@mail.mil
§University College London, London WC1E 6EA, UK. Email: s.pasteris@cs.ucl.ac.uk

***Abstract*—Network topology plays an important role in many network operations. However, it is very difficult to obtain the topology of coalition networks due to the lack of internal cooperation. Network tomography provides a powerful solution that can infer the network routing topology from end-to-end measurements. Existing solutions all assume that routes from a single source form a tree. However, with the rapid deployment of Software Defined Coalition (SDC) and Network Function Virtualization (NFV), the routing paths in modern networks are becoming more complex. To address this problem, we propose a novel inference problem, called the weight inference problem, which infers the finest-granularity information from end-to-end measurements on general routing paths in general topologies. Our measurements are based on emulated multicast probes with a controllable "width". We show that the problem has a unique solution when the multicast width is unconstrained; otherwise, we show that the problem can be treated as a sparse approximation problem, which allows us to apply variations of the pursuit algorithms.**

I. Introduction

Topology information is at the foundation of many network operations such as path selection, service placement, overlay construction, and load balancing. Meanwhile, for networks owned by multiple coalition members, it is very hard to obtain the global topology information as such information is distributed across multiple network owners.

It is known that end-to-end performance measurements can reveal topology information. Techniques known as *network tomography* have been developed to *infer* the network routing topology from end-to-end measurements such as delays and losses, e.g., [2] and followups. However, all the existing solutions are designed for traditional communication networks, where probes from each source follow an (unknown) routing tree.

The tree assumption causes existing network tomography solutions to severely underestimate the complexity of modern communication networks, where technologies like Software Defined Coalition (SDC) and Network Function Virtualization (NFV) [3] can generate complex non-tree routing topologies. This triggers a research question: *Can we still infer useful topology information from end-to-end measurements in networks with general (possibly non-tree) routing topologies?*

In this work, we take a first step towards answering this question by inferring the network's internal performance at the "finest granularity" (see Section II-C), which provides valuable information about the routing topology.

An extended version of this work has been presented at ICC'19 [1].
This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

II. Problem Formulation

A. Network Model

We model the network (routing) topology as an edge-weighted directed graph $\mathcal{G} = (V, E)$. Each vertex $v \in V$ represents a probing source, a probing destination, or a branching/joining point between multiple measurement paths. Each edge $e \in E$ represents a connection between two vertices, which may map to a sequence of links. In this work, we focus on *loss-based weight*. Given an edge e, let α_e be the probability for a probe to successfully traverse edge e, then its weight is defined as $u_e := -\log \alpha_e$ [4] be its weight.

B. Observation Model

We measure the network from a given probing source s, which can send probes on n different paths, e.g., using different combinations of the header fields. Let $\{p_1, \ldots, p_n\}$ denote the entire set of measurement paths.

We use stripes of k unicast probes sent back to back on k paths to emulate multicast on these paths [4]. In this work, we use this trick to emulate multicast, where $k \in \{1, \ldots, n\} =: [n]$ is a design parameter. In the sequel, we call such an emulated multicast a "k-cast", and the parameter k the "width" of the multicast.

Let X_C ($C \subseteq [n]$, $0 < |C| \leq k$) be the indicator that all the unicast probes sent back to back on paths $\{p_i : i \in C\}$ successfully reach their destinations. Under the assumption that different edges exhibit independent losses, we have

$$\phi_C := -\log(\Pr\{X_C = 1\}) = -\log\Big(\prod_{e \in \bigcup_{i \in C} p_i} \alpha_e\Big) = \sum_{e \in \bigcup_{i \in C} p_i} u_e. \quad (1)$$

We define ϕ_C as the *cast weight* of a $|C|$-cast on paths $\{p_i : i \in C\}$. As we can estimate $\Pr\{X_C = 1\}$ by the fraction of joint successes on paths $\{p_i : i \in C\}$ among all the k-cast probes covering these paths, we can estimate ϕ_C consistently. Let $\mathcal{C} := \{C \subseteq [n] : 0 < |C| \leq k\}$ be the subsets of paths for which the cast weights can be measured.

C. Weight Inference Problem

We are interested in inferring the edge weights at the level of *(edge) categories* from the measured cast weights [5]. For $A \subseteq [n]$ and $A \neq \emptyset$, category Γ_A is defined as $\{e \in E : e \in p_i \text{ iff } i \in A\}$. Let $\mathcal{A} := 2^{[n]} \setminus \{\emptyset\}$ (where $2^{[n]}$ denotes the power set of $[n]$). Let w_A denote the sum of the weights of the edges in category Γ_A, referred to as *category weight*.

Definition 1. *The* weight inference problem *aims at inferring the category weights* $(w_A)_{A\in\mathcal{A}}$ *from the measured cast weights* $(\phi_C)_{C\in\mathcal{C}}$.

By definition, we have

$$\sum_{A\in\mathcal{A}: A\cap C\neq\emptyset} w_A = \phi_C, \quad \forall C \in \mathcal{C}. \tag{2}$$

III. Weight Inference Algorithms

Written in vector form, the weight inference problem aims at solving the linear system

$$\boldsymbol{D} \cdot \boldsymbol{w} = \boldsymbol{\phi} \tag{3}$$

under the constraint

$$\boldsymbol{w} \geq \boldsymbol{0}, \tag{4}$$

where $\boldsymbol{w} := (w_A)_{A\in\mathcal{A}}$, $\boldsymbol{\phi} := (\phi_C)_{C\in\mathcal{C}}$, and $\boldsymbol{D} := (\mathbb{1}_{A\cap C\neq\emptyset})_{C\in\mathcal{C},A\in\mathcal{A}}$ ($\mathbb{1}.$: indicator function).

Theorem III.1. *If $k = n$, then (3) has a unique solution.*

When $k < n$, the linear system (3) is underdetermined, and hence there is no unique solution. To resolve the ambiguity, we prefer the "simplest" solution over alternatives. Then the weight inference problem becomes

$$\min \ \|\boldsymbol{w}\|_0 \tag{5a}$$

$$\text{s.t. } (3), (4). \tag{5b}$$

Problem (5) is an instance of the *sparse approximation problem with noiseless observations*. Here we assume that sufficiently many probes have been sent to accurately measure the cast weights. Otherwise, we can easily incorporate measurement errors by relaxing (3) into an error bound.

Since such problems are generally NP-hard, approximate solutions have been proposed. A popular solution is *basis pursuit (BP)*, which aims to minimize the ℓ-1 norm. We can thus apply existing LP solvers. In particular, we find that the *simplex method* provides a guaranteed sparsity.

Theorem III.2. *BP based on the simplex method provides a solution to (5) with no more than $\sum_{i=1}^{k}\binom{n}{i}$ non-zero entries, i.e., the solution is $O(n^k)$-sparse if $k = O(1)$.*

Since the linear system (3) has $2^n - 1$ variables, the worst-case complexity of the simplex method is exponential in $2^n - 1$, which is *super-exponential* in the number of measurement paths n. To reduce the complexity, we borrow from the greedy heuristics for sparse approximation problems, known as *orthogonal matching pursuit (OMP)* (Algorithm 1), which iteratively finds non-zero entries one at a time [6].

Non-negative orthogonal matching pursuit (nOMP): OMP is designed by computing an orthogonal projection onto the selected atoms using least square programming. As shown in Algorithm 1, nOMP updates the solution using non-negative least square programming $\arg\min_{\boldsymbol{x}\geq\boldsymbol{0}} \|\boldsymbol{\phi} - \boldsymbol{D}_L\boldsymbol{x}\|_2$ (line 5 of Algorithm 1), where L denotes the set of indices of the selected atoms, $\boldsymbol{D}_L$ the sub-matrix of $\boldsymbol{D}$ formed by these atoms, and $\boldsymbol{w}_L$ the sub-vector of $\boldsymbol{w}$ with indices in L.

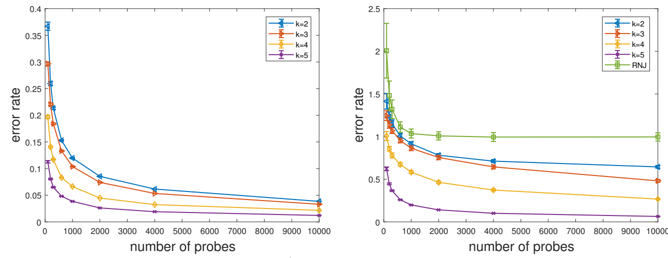

(a) measurement error for $\boldsymbol{\phi}$ (b) inference error for $\boldsymbol{w}$

Fig. 1. Relative error as k and the number of k-cast probes vary ($n = 5$).

Algorithm 1 Non-negative Orthogonal Matching Pursuit

1: **Initialization:** $L = \emptyset$, $\boldsymbol{w} = \boldsymbol{0}$ and $\boldsymbol{r} = \boldsymbol{\phi}$
2: **while** $\max \boldsymbol{D}^T\boldsymbol{r} > 0$ **do**
3: $\quad l \leftarrow \arg\max \boldsymbol{D}^T\boldsymbol{r}$
4: $\quad L \leftarrow L \cup \{l\}$
5: $\quad \boldsymbol{w}_L \leftarrow \arg\min_{\boldsymbol{x}\geq\boldsymbol{0}} \|\boldsymbol{\phi} - \boldsymbol{D}_L\boldsymbol{x}\|_2$
6: $\quad \boldsymbol{r} \leftarrow \boldsymbol{\phi} - \boldsymbol{D}_L\boldsymbol{w}_L$
7: **end while**

IV. Performance Evaluation

We evaluate relative error as k and the number of k-cast probes vary on Internet Service Provider (ISP) topologies from the Rocketfuel project.

Fig. 1 (a) shows the error in estimating the cast weights $\boldsymbol{\phi}$ (i.e., measurement error). Given the number of k-cast probes, the measurement error decreases as we increase k. Fig. 1 (b) shows the error of inferring category weights $\boldsymbol{w}$. These curves show the same trend as the curves in Fig. 1 (a), but with much larger gaps. While the proposed algorithms can find a feasible solution, it may not be the closest approximation to the ground truth. Nevertheless, increasing k still helps to reduce the error due to having more constraints. When $k = n$, we have a unique solution, which equals the ground truth if there is no measurement error. These plots show the importance of (emulating) broadcast in inferring general topologies.

References

[1] Y. Lin, T. He, S. Wang, K. Chan, and S. Pasteris, "Multicast-based weight inference in general network topologies," in *IEEE ICC*, May 2019.

[2] R. Caceres, N. G. Duffield, J. Horowitz, F. L. Presti, and D. Towsley, "Loss-based inference of multicast network topology," in *IEEE CDC*, 1999.

[3] "Network Functions Virtualisation — Introductory White Paper," White Paper, ETSI, 2012. [Online]. Available: https://portal.etsi.org/nfv/nfv_white_paper.pdf

[4] J. Ni, H. Xie, S. Tatikonda, and Y. R. Yang, "Efficient and dynamic routing topology inference from end-to-end measurements," *IEEE/ACM Transactions on Networking*, vol. 18, no. 1, pp. 123–135, February 2010.

[5] Y. Lin, T. He, S. Wang, K. Chan, and S. Pasteris, "Looking glass of NFV: Inferring the structure and state of NFV network from external observations," in *IEEE INFOCOM*, April 2019.

[6] M. Yaghoobi, D. Wu, and M. Davies, "Fast non-negative orthogonal matching pursuit," *IEEE Signal Processing Letters*, vol. 22, no. 9, pp. 1229–1233, September 2015.