OCCAM: An Optimization Based Approach to Network Inference

Anirudh Sabnis
UMass Amherst
asabnis@cs.umass.edu

Ramesh K. Sitaraman
UMass Amherst
ramesh@cs.umass.edu

Donald Towsley
UMass Amherst
towsley@cs.umass.edu

ABSTRACT

We study the problem of inferring the structure of a communication network based only on network measurements made from a set of hosts situated at the network periphery. Our novel approach called "OCCAM" is based on the principle of occam's razor and finds the "simplest" network that explains the observed network measurements. OCCAM infers the internal topology of a communication network, including the internal nodes and links of the network that are not amenable to direct measurement. In addition to network topology, OCCAM infers the routing paths that packets take between the hosts. OCCAM uses path metrics measurable from the hosts and expresses the observed measurements as constraints of a mixed-integer bilinear optimization problem that can then be feasibly solved to yield the network topology and the routing paths. We empirically validate OCCAM on a wide variety of real-world ISP networks and show that its inferences agree closely with the ground truth. Specifically, OCCAM infers the topology with an average network similarity score of 93% and infers routing paths with a path edit distance of 0.20. Further, OCCAM is robust to error in its measured path metric inputs, producing high quality inferences even when 20-30% of its inputs are erroneous. Our work is a significant advance in network tomography as it proposes and empirically evaluates the first method that infers the *complete* network topology, rather than just logical routing trees from sources.

1. INTRODUCTION

Enterprises rely heavily on communication networks, such as the Internet, for their operations. However, they lack explicit knowledge about the topological properties of their network, such as the routers, the links, and the routing table of the network. But, there are great benefits for enterprises to know the topological properties of their communication network. For instance, by deducing the graph structure of the network and the routing paths between their hosts, the enterprise can better understand the impact of node (i.e., router) and link failures on their mission-critical communication, leading to better disaster planning and recovery.

Consider an enterprise with hosts that are connected by the communication network shown in Figure 1. Inferring the underlying graph topology of the network will enable the enterprise to understand network bottlenecks, such as link (a, e) whose failure can impact the connectivity of host A to hosts C, D, E. Knowing the network topology can also help plan alternate routes for disaster recovery, such as using links (a, x), and (x, e) if (a, e) fails.

We require two types of topological metrics of the routing paths as inputs to our network inference: path sharing metrics (PSMs) and distance metrics (DMs).

Path Sharing Metrics (PSMs) PSMs measure whether and to what extent routes (i.e., paths) between hosts share links. Let $\pi(S, T)$ represent the set of links in the routing path from host S to host T. Given source S and destinations T_1 and T_2, $PSM(S, T_1, T_2)$ is the number of links that are shared between the path from S to T_1 and the path from S to T_2. More precisely, $PSM(S, T_1, T_2)$ equals $|\pi(S, T_1) \cap \pi(S, T_2)|$. Our work does not require knowing absolute values for the PSMs, but only relative ones. For instance, given a source S and three destinations T_1, T_2, and T_3, it suffices to know how $PSM(S, T_1.T_2)$ compares with $PSM(S, T_2, T_2)$.

Distance Metrics (DMs) A DM measures the distance, i.e., the number of links, in the path from a source host S to hosts T. More precisely, $DM(S, T)$ equals $|\pi(S, T)|$. Again, we do not require absolute values of the DMs, and relative ones will suffice. More precisely, given a source S and two destinations T_1 and T_2, it suffices to know how $DM(S, T_1)$ compares with $DM(S, T_2)$.

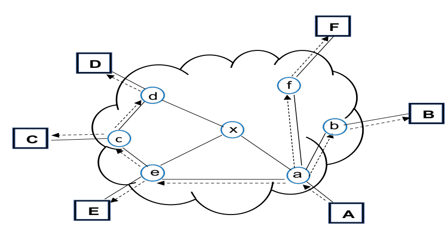

Figure 1: Network inference using network measurements from the hosts.

Problem Statement. The paper is focused on the following intuitive and important question: *can we infer the complete network topology and host-to-host routing paths of a communication network, knowing only the relative values of the path metrics (PSMs and/or DMs)?* Precisely, we are given *PSM constraints* of the form

$$PSM(S, T_1, T_2) < PSM(S, T_2, T_3), \quad (1)$$

for multiple combinations of sources S and destination hosts

T_1, T_2, and T_3. We are also given DM constraints of the form

$$DM(S,T_1) < DM(S,T_2), \quad (2)$$

for multiple combinations of sources S and destination hosts T_1 and T_2. Note that we also allow equality to be stated in both of the above constraints. Our goal is to find the "simplest" communication network $N = (G, H, P)$ that satisfies the given PSM and DM constraints, where $G = (V, E)$, is the underlying graph with a set of nodes V and a set of links E, $H \subset V$ is the set of hosts, P is the set of routing paths in G between host pairs in $H \times H$.

Our Contributions. To our knowledge, OCCAM is the first technique that can accurately infer the *complete* network topology and host-to-host routing paths of a communication network. Prior work on network tomography has been limited to inferring logical trees that are formed by the routes from a host as root to other hosts as leaves, using path sharing metrics derived from multicast probes [1]. Some of our contributions follow.

1) We propose a novel theoretical approach (OCCAM) that formulates as an optimization problem the Occam's razor principle of finding the "simplest" network that obeys observations. The solution to the optimization problem yields the inferred network and routing paths.

2) We evaluated OCCAM on several real-world ISP networks and showed that it produced high quality inferences.

3) Further, to more closely simulate real-world network inference, we implemented ISP topologies on the DETER [3] network emulator. We implemented unicast packet probes and derived the PSM and DM values from packet-level measurements. We then fed those values as inputs to OCCAM and showed that it produced high-quality inferences.

Full Paper. For a full version of the paper with the complete list of contributions, see [4].

2. THE OCCAM APPROACH

The key idea of our approach is to view network inference as an optimization problem where the "simplest" network that satisfies the observed path sharing metrics is produced as a solution. We view this approach as analogous to Occam's razor that is a heuristic element of the scientific method and advocates the construction of the simplest and most parsimonious model that obeys the empirical observations. We capture the existence of nodes and links, as well as the membership of links in routing paths, as indicator variables whose values are set by the optimization process.

Objective function. There are many notions of simplicity possible in a network setting. We use the notion that the inferred network should have (i) the smallest number of links, and (ii) the smallest total host-to-host shortest path distance. We can express the notion of simplicity as the objective function that need to minimized as follows.

$$\min: \ \alpha \sum_{s\in H}\sum_{t\in H} m_t^s + (1-\alpha) \sum_{i\in V}\sum_{j\in V} w_{ij}, \quad (3)$$

where m_t^s is an integer variable denoting the length of the path from source host $s \in H$ to destination host $t \in H$, w_{ij} is an indicator variable indicating if a link exists between node i and node j in the inferred network, and $0 \le \alpha \le 1$ weighs the relative importance of the two components.

Path sharing. Next, we encode the PSM inputs. For every path sharing metric $PSM(S, T_1, T_2) < PSM(S, T_2, T_3,)$ obtained from measurements, we add the following constraint

$$\sum_{i\in V} v_i^{S,T_1} v_i^{S,T_2} < \sum_{i\in V} v_i^{S,T_2} v_i^{S,T_3}, \quad (4)$$

where $v_i^{S,T}$ is an indicator variable indicating if node i is on the path from host S to host T. The *LHS* of the above inequality thus counts the number of nodes present in the intersection of routes $\pi(S,T_1)$ and $\pi(S,T_2)$. Similarly the *RHS* finds the number of nodes present in the intersection of paths $\pi(S,T_2)$ and $\pi(S,T_3)$.

Source-oblivous paths. Typically, a packet at a node $i \in V$ is forwarded to the next node $j \in V$ by consulting a routing table that provides the "next-hop" for each destination $T \in H$, independent of the packet's source. In particular, two packets at a node i from different sources is forwarded to the same next node j if they are going to the same destination T. We capture this constraint as follows:

$$\sum_{j\in V} d_{ij}^T \le 1 \qquad \forall i \in V \setminus H \qquad T \in H, \quad (5)$$

where d_{ij}^T is an indicator variable indicating if a link (i, j) is on any of the paths P^T, where $P^T \subset P$ is a set of routing paths in G between host pairs in $H \times \{T\}$. Above equation ensures that, if a node i is on any of paths to destination T, the number of possible forward hops is at most 1.

Other path properties. Since the host-to-host routing paths in a network are acyclic, we add constraints to enforce that. Further, let $P^S \subset P$ be the set of routing paths in G between host pairs $\{S\} \times H$. We add constraints to ensure that the links belonging to P^S forms a tree.

Dealing with inaccurate path sharing measurements. Path sharing measurements can be inaccurate in real-world networks when there is no multicast available and there is a significant amount of background traffic. Such was the case with some of our experiments on the DETER testbed. An optimization variant that we used to tackle measurement inaccuracies is to convert the hard constraints of equation 4 into soft constraints by moving them to the objective function. That is, we add a third component to the objective function in Equation 3 that represents the number of path sharing constraints that are violated. Thus, the new objective function finds the "simplest" network that obeys "most" of the observed measurements. With this approach, OCCAM made accurate network inferences even when 20-30% of the path sharing constraints were incorrect.

Solution approach. The optimization problem formed with the above objective function and constraints is a **M**ixed **I**nteger **B**ilinear **P**rogram (*MIBP*). The problem is then linearized using standard techniques to form a **M**ixed **I**nteger **P**rogram (*MIP*). We use the distributed parallel MIP feature of CPLEX to solve our problem on a server cluster.

3. EMPIRICAL RESULTS

We use several real-world ISP networks obtained from topology-zoo [2] to evaluate OCCAM.

Quality metrics for network inference. Given a communication network $N = (G, H, P)$ and an inferred network $N = (G', H, P')$, we introduce two metrics that quantitatively measures the quality of inference.

Network Similarity (NS) The *NS* score measures how close the inferred graph $G' = (V', E')$ is to the ground truth of $G = (V, E)$. Intuitively, we compute the "best" one-to-one function $\phi : V \to V'$ to match the vertices of one graph with the vertices of the other. We then compute the percentage of links that are matched under ϕ, i.e., percentage of links present in both graphs. Formally,

$$NS(G, G') = \max_{\phi : V \to V'} \left(\frac{\sum_{i,j \in V \times V} E_{i,j} \wedge E'_{\phi(i),\phi(j)}}{\sum_{i,j \in V \times V} E_{i,j} \vee E'_{\phi(i),\phi(j)}} \right) \times 100,$$

where $E_{i,j}$ (resp., $E'_{\phi(i),\phi(j)}$) are indicator variables that is set to 1 if the corresponding link is present in G (resp., G') and 0 otherwise, $\wedge$ is the boolean AND operator, and $\vee$ is boolean OR operator. Note that the numerator evaluates the number of links that are in common between the graphs and the denominator is the total number of links present in either graph.

Path Edit Distance (PED) The *PED* metric for path sets P and P' is the average path edit distance between the corresponding paths in P and P'. Note that given the one-to-one function ϕ, each path $\pi \in P$ has a corresponding path $\pi' \in P'$ such that the two corresponding paths connect the same host pairs under ϕ. Path edit distance between two paths π and π' is simply the number node insertions, deletions and substitutions required to convert one path to the other. The overall PED is simply the average PED of the individual path pairs.

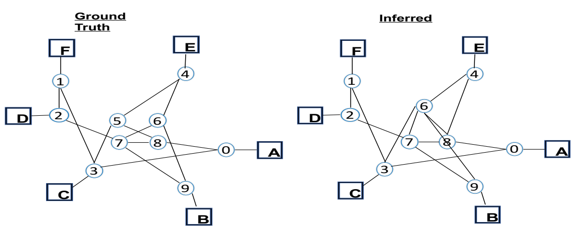

Figure 2: Columbus network inferred by OCCAM

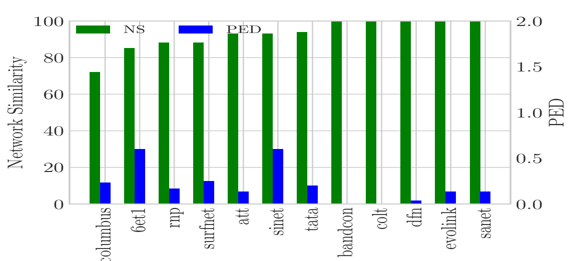

Figure 3: NS and PED for various real-world topologies from topology-zoo

3.1 Path metrics from ground truth

The first set of experiments the PSMs and DMs that are inputs to OCCAM are obtained from Ground Truth. We evaluate OCCAM on 12 networks obtained from topology zoo. Figure 2 shows the ground truth and the inferred topology of the COLUMBUS network. Notice that as we find the "simplest" network satisfying the constraints, we sometimes infer a more compact form of the original network. For instance, in Figure 2, nodes 5 and 6 in the ground truth has been merged into one node (node 6) in the inferred graph.

We further quantify our inference using the network similarity and path edit distance metrics. Over the tested networks, we obtain an average network similarity score of 93%. Further, we obtained an average path edit distance is 0.20.

3.2 Path Metrics from DETER

We used a packet-level network emulator called DETER [3] to emulate the TATA and 6et1 topologies. We obtain the PSMs and DMs from DETER experimentally, by sending unicast packet probes. A key difference with the previous experiments is that roughly 20-30% of the measured path sharing constraints were erroneous when compared to ground truth. Therefore, we use OCCAM with soft path sharing constraints as described in Section 2. Interestingly, OCCAM is able to suppress the erroneous constraints and obtain an accurate inference. For instance, Figure 4 shows the original and the inferred TATA network. The network similarity scores achieved for both TATA and 6et1 networks is about 89%. Further, we obtain a path edit distance of about 0.7 and 0.66 respectively.

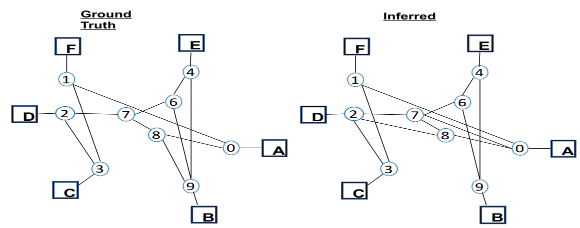

Figure 4: TATA network on DETER

4. CONCLUSION

Our work is the first to demonstrate the feasibility of inferring the complete network topology and routing paths using path sharing and path distance information measured at the hosts. However, many questions remain open. An interesting question is how much path sharing information is needed to infer the complete network. Our preliminary work suggests that highly accurate inference is possible even with partial and/or incorrect path sharing information. Another natural extension of our work is whether the topology and path inference can be extended to infer the link capacities in the network.

5. REFERENCES

[1] N. G. Duffield and F. L. Presti. Network tomography from measured end-to-end delay covariance. *IEEE/ACM Transactions on Networking (TON)*, 12(6):978–992, 2004.

[2] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The internet topology zoo. *IEEE Journal on Selected Areas in Communications*, 29(9):1765–1775, 2011.

[3] J. Mirkovic, T. V. Benzel, T. Faber, R. Braden, J. T. Wroclawski, and S. Schwab. The deter project: Advancing the science of cyber security experimentation and test. In *In Technologies for Homeland Security (HST), 2010 IEEE International Conference on*, page 7, 2010.

[4] A. Sabnis, R. K. Sitaraman, and D. Towsley. Occam: An optimization-based approach to network inference. *https://arxiv.org/abs/1806.03542*, 2018.