

# UNDERSTANDING THE TOPOLOGY OF A TELEPHONE NETWORK VIA INTERNALLY-SENSED NETWORK TOMOGRAPHY

Michael G. Rabbat<sup>1</sup> John R. Treichler<sup>2</sup> Sally L. Wood<sup>3</sup> Michael G. Larimore<sup>2</sup>

<sup>1</sup>University of Wisconsin  
Electrical & Computer Engineering  
Madison, Wisconsin 53706

<sup>2</sup>Applied Signal Technology  
Sunnyvale, California 94086

<sup>3</sup>Santa Clara University  
Electrical Engineering  
Santa Clara, California 95053

## ABSTRACT

The ability to determine the topology of worldwide telephone networks offers the promise of substantially improving their operating efficiency. This paper explores the problem of identifying the topology of a telephone network using observations made within the network. Using tomographic methods inspired by medical imaging, we consider measurements made by transmitting probes (e.g., phone calls) between network endpoints. In general, these measurements alone do not suffice to reconstruct a unique network, and in fact, there are many network topologies from which the set of measurements could have been generated. We propose a topology reconstruction algorithm based on correlating measurements collected at different internal nodes, and identify conditions under which correctness of the inferred topology is guaranteed.

## 1. BACKGROUND

The operation and planning of voice and data communication networks requires accurate information about their topology, and their maintenance requires even more information, including understanding of the loading and performance of network elements such as switches, routers, and transmission links. This information has traditionally been determined in a number of ways, ranging from purely physical audit of the network, through active interrogation of network elements, to using `traceroute` to probe the paths by which IP packets flow.

Historically, the national telephone system was owned and operated by a single company and similar tools did not even exist for assessing telephone networks. With the recent explosion of the telecom business the international telephone system is a combination of multiple wired and wireless networks. Like the Internet, no single entity knows where all of the infrastructure lies or how calls are exactly routed through it.

This paper examines the problems and opportunities associated with an alternative approach to assessing internal network characteristics, that of using active probes through the network which are passively observed in transit. Although this paper focuses on determining the topology of a telephone network, the methods described here can be applied to packet-switched networks as well.

There is a growing body of literature on the topic of *network tomography*, the concept of using packet probes through a network to infer internal characteristics such as bandwidth, congestion level, or to determine the network topology [1, 2, 3]. In strong analogy to the medical imaging techniques from which they take their name [4], most of these methods can be described as “active,

endpoint-to-endpoint” techniques since they send probes – typically in the form of IP packets – from one side of the network to the other (see, e.g., [5, 6]).


The problem and technique studied in this paper differs in two ways – one simple and one more fundamental – than those studied in the existing network tomography literature. The simple differences: for probes we will use telephone call attempts from many locations throughout the world instead of using IP “pings”. The more fundamental difference here is that the sensors of the probes are located somewhere in the middle of the network rather than at the edge. Figure 1 illustrates this difference [7]. Additionally, unlike tomography for medical imaging where probes (e.g., X-rays) travel in roughly a straight line and are attenuated monotonically as they transit “the system”, telephone calls do not necessarily travel in a straight line through the phone network and there is no concept of measurable attenuation of a telephone call.

### 1.1. Problem Statement


To restate it concisely, the goal of this work is to infer a meaningful representation of the telephone network topology given information about which trunks are transited by calls made between known endpoints. For a given call, we know which trunks are transited but not the order of transit. We employ the standard graph structure for representing the telephone network, with nodes in the graph corresponding to switches (or public branch exchanges) and links in the graph corresponding to telephone trunks.

Suppose that a subset of trunks in the network can detect when they are used for a call between a known source and destination. In what follows we refer to these trunks as *measurement links* in the reconstruction graph, with the set of all measurement links denoted by  $\mathcal{M}$ . Let  $\mathcal{S}$  denote the set of sources, and  $\mathcal{D}$  denote the set of destinations. Measurements are collected by placing calls from a source  $s \in \mathcal{S}$  to a destination  $d \in \mathcal{D}$ , and recording at each measurement link  $m \in \mathcal{M}$  whether the call from  $s$  to  $d$  transited  $m$ . A total of  $N$  measurements are made,  $j = 1, \dots, N$ .

Measurement links are directed, meaning that all calls transiting a given link flow in the same direction. Additionally, for each of the  $N$  calls we know which links in  $\mathcal{M}$  are transited, but we do not know the order they are transited in. We assume that routing policies in the network are stationary during the measurement period. Thus, if calls were repeatedly placed from source  $s$  to destination  $d$  they would transit exactly the same internal links each time. We will also assume that the measurements are perfect, meaning that  $x_m(j) = 1$  if and only if call  $j$  transited link  $m$ . Finally, we assume that routes never contain cycles; a call will transit any link in the network at most once.



(a) Classical tomography – Both the probe transmitters and receivers are at the edge of the network.



(b) Internally-sensed tomography - Probe transmitters are at the edge of the network and probes are destined for receivers at the edge, but measurements are made within the network.

**Fig. 1.** Difference between classical and internally sensed network tomography. Nodes marked  $s$  are probe transmitters, and nodes marked  $m$  are measurement points.

The measurement data for each source, destination, and measurement link is summarized in a binary vector of length  $N$ . We associate with each element  $e \in S \cup \mathcal{D} \cup \mathcal{M}$  a vector  $x_e \in \{0, 1\}^N$ , where  $x_e(j) = 1$  indicates that  $e$  was involved with the  $j^{\text{th}}$  call. Thus, for each  $s \in \mathcal{S}$ , if  $x_s(j) = 1$  then  $s$  was the source of call  $j$ . Similarly, if  $x_d(j) = 1$  then  $d \in \mathcal{D}$  was the destination of call  $j$ , and for  $m \in \mathcal{M}$ ,  $x_m(j) = 1$  indicates that call  $j$  transited measurement link  $m$ . Again, we emphasize that the nature of these measurements is such that, e.g., we know measurement links  $m_1, m_2, \dots$  appear in the  $j^{\text{th}}$  call, but the measurements give no indication of the order of  $m_1, m_2, \dots$  in the call path.

## 2. METHODOLOGY


The main challenge of this problem stems from the fact that we do not know the order in which the observing links are transited for a given call. Without the order and without additional information which can be used to regularize the problem there will generally be many networks which are consistent with a given data set. While

we would like to come up with a reconstruction which exactly resembles the true physical network, this is generally not possible without additional information or prior knowledge.

Given the collection  $\{x_e\}$  for  $e \in S \cup \mathcal{D} \cup \mathcal{M}$ , we would like to reconstruct the network topology which meets the following criteria. First and foremost, the reconstructed topology should be consistent with the measurement data. If links  $m_1$  and  $m_2$  are transited in a call from  $s$  to  $d$ , then there should be a path in the reconstructed network from  $s$  to  $d$  which includes  $m_1$  and  $m_2$ , and does not include any other links in  $\mathcal{M}$ . Additionally, to the greatest extent possible we would like the reconstructed network to *only* reflect information gathered in the measurements. By this, we mean that if no call originating at source  $s$  transited link  $m$  then the topology inference algorithm should prefer networks which do not make it possible to get to  $m$  from  $s$ .

Now, if we knew the order of measurement links in each path we could reconstruct the network by connecting the observing links in the appropriate order. Based on this idea, our approach to network reconstruction is to try and estimate an ordering for the observing links, and then reconstruct the network accordingly. Intuitively, if measurement link  $m_1$  *always* observes calls being placed from source  $s$  then it seems likely that  $m_1$  is located near  $s$  in the network. Likewise, if calls made to destination  $d$  always transit link  $m_2$  then we are inclined to think that  $m_2$  is located near  $d$ . Measurement links which observe calls made to or from a variety of sources or destinations are thought of as being near the center of the network.

To make this more concrete, assume that the subgraph corresponding to paths going from source  $s_1$  to every destination forms a tree as depicted in Figure 2. From the figure it is clear that if only these paths are used then a call made from  $s_1$  will not transit  $m_2$  unless it also transits  $m_1$ . On the other hand, it is possible for  $m_1$  to be transited by calls which do not transit  $m_2$ . Thus, we expect more calls to transit  $m_1$  than  $m_2$ . Now, suppose calls from  $s_1$  are also observed on two other links,  $m_3$  and  $m_4$ , but we do not know their relative order. If more calls made from  $s_1$  transit  $m_3$  than they do  $m_4$ , we infer that  $m_3$  is closer to  $s_1$  than  $m_4$ . Similar reasoning can be applied to ranking links' relative distances from a destination if we assume that the paths from all sources to a specific destination form a tree in a complementary fashion.



**Fig. 2.** Example network where the paths from a source to all the destinations it calls form a tree. The bold links indicate the paths taken by calls made from source  $s_1$  to each destination  $d_1-d_4$ .

### 2.1. The Reconstruction Algorithm

To avoid reconstructing a network which allows paths between sources and destinations which were not observed, we begin the reconstruction process with all sources and destinations as nodes

connected to a “separator” node. We then compare the binary vectors  $x_m, x_{m'}$  for each pair of measurement links  $m, m' \in \mathcal{M}$ . If  $x_m = x_{m'}$  and  $m \neq m'$ , there is an ambiguity in the measurements. We will never be able to distinguish between  $m$  and  $m'$  given the current data set. For the purpose of understanding the structure of the network the precise order of  $m$  and  $m'$  does not matter, so we merge the two links and treat them as a single link in our reconstruction.

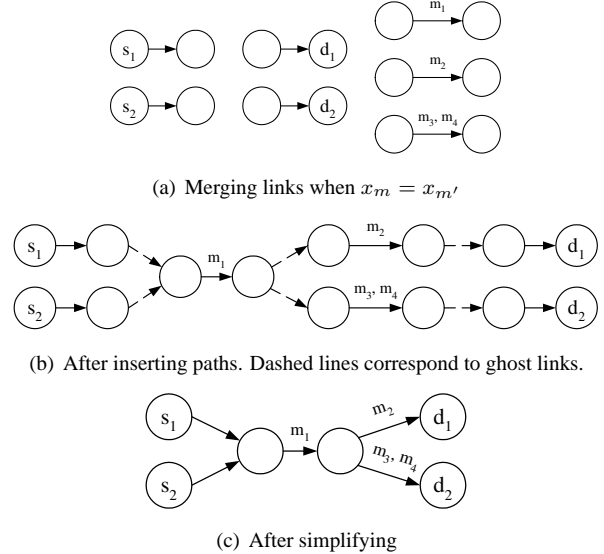
Next, for each call,  $j = 1, \dots, N$ , we associate a value  $\alpha_j(m)$  with each measurement link transited, as follows. Let  $s_j$  be the source and  $d_j$  be the destination of call  $j$ . For each measurement link  $m$  with  $x_m(j) = 1$ , we calculate  $\alpha_j(m) = x_{s_j}^T x_m - x_{d_j}^T x_m$ . Note that  $x_s^T x_m$  is equal to the number of calls from  $s$  which transited  $m$ , and  $x_d^T x_m$  is the number of calls which transited  $m$  on their way to  $d$ . Thus, an  $\alpha_j(m)$  which is large and positive suggests that  $m$  is closer to the source than it is to the destination. Likewise,  $\alpha_j(m)$  large and negative suggests that  $m$  is closer to the destination. To estimate the order of elements in the  $j^{\text{th}}$  call, we take measurement link  $m^* = \arg \max_{m \in \mathcal{M}} \{\alpha_j(m)\}$  and place it closest to  $s$ . The measurement link with the next largest  $\alpha_j$ -value follows  $m^*$ , and so forth, until all links are in order.

We then insert the measurement links, now in estimated order, into the network by first connecting the separator node of the source to the tail node of the measurement link with the greatest  $\alpha$ -value. If measurement link  $m_2$  follows link  $m_1$  in order (e.g.,  $s \rightarrow m_1 \rightarrow m_2 \rightarrow \dots \rightarrow d$ ), rather than merging the head node of  $m_1$  with the tail node of  $m_2$ , we insert a ghost link between the head and tail. This is to avoid creating paths in the reconstructed network which were not observed. Finally, after all paths have been inserted, we simplify the network by combining any links which are in series. This has the effect of removing any “artifacts” (e.g., separator nodes or ghost links which were unnecessarily introduced). Figure 3 depicts various stages of the reconstruction process for a simple example with two sources, two destinations, and four measurement links. The measurements for this example are presented in Table 1 below. Because we use separators and ghost links to avoid introducing unobserved paths and then simplify the network after inserting all calls, the order in which we insert call paths into the reconstruction network before simplifying does not effect the final inferred topology.

$j$	$s_1$	$s_2$	$d_1$	$d_2$	$m_1$	$m_2$	$m_3$	$m_4$
1	1	0	1	0	1	1	0	0
2	1	0	0	1	1	0	1	1
3	0	1	1	0	1	1	0	0
4	0	1	0	1	1	0	1	1

**Table 1.** An example set of measurements for the reconstruction depicted in Figure 3. Columns correspond to  $x_e$  for each element.

To further explore the performance of the reconstruction algorithm we have experimented with reconstructing larger networks. Figure 4(a) depicts a ring-topology network with four sources, four destinations, and eight measurement links. Calls are placed from each source to every destination. For this example the algorithm returns the correct topology (which is the one shown). To get a feel for how sensitive the procedure is to the amount of measurements available, we repeated the same experiment but without taking measurements on links  $m_1$  and  $m_2$ . The resulting reconstruction topology, shown in Figure 4(b), illustrates that the technique still works moderately well, with the additional presence of some



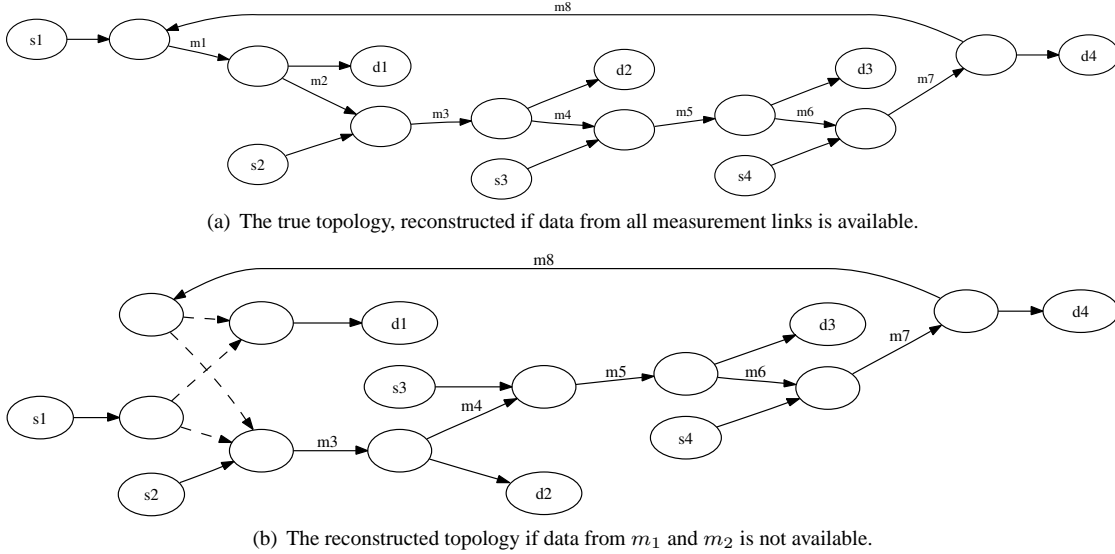
**Fig. 3.** Reconstructing a network topology for the example measurements given in Table 1.

ghost links as expected. Notice that two ghost links take the place of each missing measurement link in an attempt to assure that no paths exist in the network which were not measured.

### 3. DISCUSSION


The main principle guiding the design of our reconstruction algorithm is that when a measurement link is transited often by paths associated with a specific endpoint, the link is probably close to that endpoint. There is ambiguity in link order when two measurement links observe exactly the same set of calls. Beyond that case, when routing in the network is such that the paths going from a source to all destinations form a tree, we are guaranteed to correctly reconstruct the topology for the reasons discussed at the beginning of Section 2. A natural question one may ask is whether this situation ever arises in real life. One plausible situation where the routing-tree assumption holds is when shortest-path routing is used to determine call paths, with any ties resolved deterministically and identically for each path. Alternatively, if the network has a ring or star-shaped topology then there is only a single route from any source to any destination, and the assumptions hold.

However, in general, this condition is not necessarily true and the reconstruction algorithm may return the wrong network. Often telephone networks are designed so that there will be multiple routes between endpoints for robustness and efficient resource usage. We have proceeded with these somewhat unrealistic assumptions for the sake of taking an initial stab at the problem. In certain cases we will be able to detect that the assumptions do not hold from the measurements. For instance, if repeated calls are made from a source to a destination, and different paths are measured on different attempts we conclude that multiple routes are being used. Also, suppose calls are measured from source  $s$  to multiple destinations. If some of these calls transit link  $m_1$  but not  $m_2$ , some transit  $m_2$  but not  $m_1$ , and some transit both  $m_1$  and  $m_2$  then the paths from  $s$  to the destinations cannot form a tree and we know that our assumptions are violated.



**Fig. 4.** Example of a more complicated network which we correctly reconstruct.

An example of a situation where we cannot detect that the assumptions are violated is if the true network employs load balancing in the following manner. Consider the set of measurements give in Table 2. The true topology could be the one depicted in Figure 5(a) if calls to  $d_1$  get routed through  $m_2$  and  $m_4$  and calls to  $d_2$  are routed through  $m_3$  and  $m_5$ . Because all calls, regardless of destination, go through both  $m_1$  and  $m_6$ , it will appear to the algorithm that they are closer to the center of the network than the other links, and the algorithm will erroneously return the network shown in Figure 5(b).



**Fig. 5.** An example where our algorithm fails. Suppose (a) is the true network. If all calls to  $d_1$  are routed through  $m_2$  and  $m_4$ , and all calls to  $d_2$  are routed through  $m_3$  and  $m_5$  then the algorithm will return network (b).

We plan to address these issues in future work. Specifically, we intend to investigate how additional information such as the geographic location of measurement links or partial knowledge of link order can be used to condition the problem and possibly resolve ambiguities or detect the presence of load balancing-like structures. We also speculate that the presence of multiple routes between endpoints in fact provides more information about the in-

$s_1$	$s_2$	$d_1$	$d_2$	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$
1	0	1	0	1	1	0	1	0	1
1	0	0	1	1	0	1	0	1	1
0	1	1	0	1	1	0	1	0	1
0	1	0	1	1	0	1	0	1	1

**Table 2.** Measurements for the example depicted in Figure 5

ternal network and plan to investigate exactly how this information can be exploited in a network reconstruction algorithm.

#### 4. REFERENCES

- [1] M. Coates, A. Hero, R. Nowak, and B. Yu, "Internet tomography," *IEEE Signal Processing Magazine*, May 2002.
- [2] D. Arifler, G. de Veciana, and B. Evans, "Network tomography based on flow level measurements," in *Proc. IEEE ICASSP*, Montreal, Canada, May 2004.
- [3] M.-F. Shih and A.O. Hero, "Network topology discovery using finite mixture models," in *Proc. IEEE ICASSP*, Montreal, Canada, May 2004.
- [4] R.N. Bracewell, *Two-Dimensional Imaging*, Prentice-Hall, Upper Saddle River, NJ, 1995.
- [5] L. Ciavattone, A. Morton, and G. Ramachandran, "Standardized active measurements on a tier 1 IP backbone," *IEEE Comm. Mag.*, June 2003.
- [6] C. Fraleigh, et al., "Packet-level traffic measurements from the Sprint IP backbone," *IEEE Network*, vol. 17, no. 6, pp. 6–16, November 2003.
- [7] J.R. Treichler, M.G. Larimore, S.L. Wood, and M. Rabbat, "Determining the topology of a telephone network system using internally sensed network tomography," in *Proc. IEEE DSP Workshop*, Taos, NM, August 2004.