Link Identifiability in Communication Networks with Two Monitors

Liang Ma[†], Ting He[‡], Kin K. Leung[†], Ananthram Swami[§], and Don Towsley[*]
[†]Imperial College, London, UK. Email: {l.ma10, kin.leung}@imperial.ac.uk
[‡]IBM T. J. Watson Research Center, Yorktown, NY, USA. Email: the@us.ibm.com
[§]Army Research Laboratory, Adelphi, MD, USA. Email: ananthram.swami.civ@mail.mil
[*]University of Massachusetts, Amherst, MA, USA. Email: towsley@cs.umass.edu

***Abstract*—We investigate the problem of identifying individual link metrics in a communication network by measuring end-to-end metrics of selected paths between monitors, assuming that link metrics are additive and constant during the measurement, and measurement paths cannot contain cycles. Previous work showed that even the minimum number of monitors required for complete identification can be large in some practical networks, suggesting high monitor deployment cost. We therefore address this issue by proposing an efficient algorithm to place a fixed number of monitors for maximizing network identifiability, with concrete results for the two-monitor case. Evaluation on real ISP topologies shows that although a large number of monitors is required for complete identifiability, we can usually identify a substantial portion of links using two optimally placed monitors.**

I. INTRODUCTION

Accurate and timely knowledge of network internal states (e.g., delays on individual links) is essential for various network operations such as route selection, resource allocation, and fault diagnosis. Directly measuring the performance of individual network elements (e.g., nodes/links) is, however, not always feasible due to the lack of support at internal network elements for making such measurements. These limitations motivate the need for an *external* approach, *network tomography* [1], which provides a methodology for inferring network internal characteristics through externally-available end-to-end measurements. Besides eliminating the need for special cooperation from internal nodes, network tomography can also reduce traffic overhead.

In this paper, we consider the particular network tomography problem of inferring *additive* link metrics from end-to-end cycle-free (simple) path measurements, which enables to formulate this problem as a system of linear equations, with link metrics representing unknown variables and end-to-end path measurements representing the known constants. To completely identify a network, our previous work [2] shows that even the minimum required number of monitors can be large. In cases where the network operator does not have sufficient budget, it is still possible to identify a subset of the link metrics through a partial deployment. Therefore, we address this essential issue by solving two closely related subproblems: (a) determine the fraction of link metrics identifiable under a given monitor placement, and (b) the optimal placement of a given number of monitors so as to maximize the fraction.

II. PROBLEM FORMULATION AND FUNDAMENTALS OF NETWORK IDENTIFIABILITY

We assume that the network topology is known and model it as an undirected graph $\mathcal{G} = (V, L)$, where V and L are the sets of nodes and links, respectively. Without loss of generality, we assume $\mathcal{G}$ is connected, as different connected components have to be monitored separately. A subset of nodes in V are monitors, which can initiate/collect measurements. Let $\mathbf{w}$ be the column vector of all link metrics (which are unknown), and $\mathbf{c}$ the column vector of path measurements (which are observed, $|\mathbf{c}| = \gamma$). Then, we have a linear system

$$\mathbf{R}\mathbf{w} = \mathbf{c}, \tag{1}$$

where $\mathbf{R} = (R_{ij})$ is a $\gamma \times n$ *measurement matrix*, with each entry $R_{ij} \in \{0, 1\}$ indicating whether link j is present on path i. We assume that monitors can control the routing of measurement packets (i.e., source routing) as long as the path starts and ends at distinct monitors and does not contain repeated nodes (simple paths).

In this paper, we focus on the case where exactly two nodes are monitors, denoted by m_1 and m_2, which is the minimum number of monitors feasible under the measurement constraint of simple paths. For this case, we established following results.

Theorem II.1. [2] None of the *exterior links* (links incident to m_1 or m_2) is identifiable (m_1m_2 is identifiable if exists).

Theorem II.2. [2] When the interior graph $\mathcal{H}$ (subgraph obtained by removing the monitors and their incident links) of $\mathcal{G}$ is connected, and there is no link m_1m_2 in $\mathcal{G}$, the necessary and sufficient conditions for identifying *all* link metrics in $\mathcal{H}$ are

① The graph remaining after deleting any interior link in $\mathcal{G}$ is 2-edge-connected;
② The augmented graph after adding link m_1m_2 to $\mathcal{G}$ is 3-vertex-connected.

We developed algorithm MMP [2] to compute the minimum number of monitors (denoted by κ) required to identify all link metrics; however, even this minimum number can be large in some practical networks (see Table I).

III. PARTIAL IDENTIFIABILITY USING TWO MONITORS

1) Set of Identifiable Links for a Given Two-Monitor Placement: Our algorithm for determining the set of identifiable links originates from a simple observation rooted in Theorems II.1 and II.2: if the interior graph $\mathcal{H}$ is 3-vertex-connected and connects to each monitor via at least two exterior links, then $\mathcal{G}$

The extended version of this paper is to appear in *IEEE Globecom'13*.

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. 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.

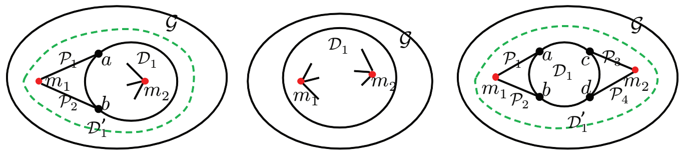

(a) (b) (c)

Fig. 1. Three categories in identifying links of a 3-vertex-connected subgraph $\mathcal{D}_1$.

satisfies the conditions in Theorem II.2, and hence the identifiability of all links can be easily determined (all interior links are identifiable, all exterior links are unidentifiable). Given an arbitrary network $\mathcal{G}$, if we decompose it into subgraphs such that each subgraph is 3-vertex-connected, then the above observation can be applied to determine link identifiability in each subgraph. Specifically, as illustrated in Fig. 1, let $\mathcal{D}_1$ denote a subgraph of $\mathcal{G}$ that is 3-vertex-connected. Our observation applies to three categories:

(1) Category 1: $m_1 \notin \mathcal{D}_1$ but $m_2 \in \mathcal{D}_1$ (or vice versa) and $\exists$ internally vertex disjoint paths $\mathcal{P}_1$ and $\mathcal{P}_2$ connecting m_1 to different nodes in $\mathcal{D}_1$ (Fig. 1-a): Our observation applies to the graph $\mathcal{D}_1'$ consisting of $\mathcal{D}_1$, m_1, $\mathcal{P}_1$ and $\mathcal{P}_2$ (modeled as single links), implying that all links in $\mathcal{D}_1$ except for those incident to m_2 are identifiable.

(2) Category 2: $\{m_1, m_2\} \in \mathcal{D}_1$ (Fig. 1-b): Our observation applies to $\mathcal{D}_1$, implying that all its links except for those incident to m_1 or m_2 are identifiable.

(3) Category 3: $\{m_1, m_2\} \notin \mathcal{D}_1$ and $\exists$ internally vertex disjoint paths $\mathcal{P}_1$ and $\mathcal{P}_2$ connecting m_1 to different nodes in $\mathcal{D}_1$, and paths $\mathcal{P}_3$ and $\mathcal{P}_4$ connecting m_2 to different nodes in $\mathcal{D}_1$ (Fig. 1-c): Abstracting each path as a single link, we can apply our observation to the graph $\mathcal{D}_1'$ formed by these four links together with $\mathcal{D}_1$ and the two monitors, which implies that all links in $\mathcal{D}_1$ are identifiable.

Based on these three categories, we develop Algorithm 1, Determining Identifiable Links under Two Monitors (DIL-2M). Our strategy is to first decompose the graph into subgraphs falling into one of the above categories, and then determine the set of identifiable links in each subgraph accordingly.

Note that for Category 2, the direct link (if any) connecting the 2 effective monitors cannot be immediately determined in the current triconnected component; nevertheless, it can be determined in a neighboring triconnected component. Triangle component of Category 3 is a special case requiring separate treatment, see [5] for details. The completeness of DIL-2M for determining all identifiable links is proved in [5]. Moreover, the overall complexity of DIL-2M is $O(|V| + |L|)$ (see [5]).

2) Optimal Two-Monitor Placement: Algorithm 1 enables us to maximize the number of identifiable links using a simple algorithm, referred to as *Optimal Monitor Placement (OMP)*: OMP enumerates all possible placements of m_1 and m_2 in $\mathcal{G}$ and computes the number of identifiable links for each placement so as to determine the optimal placement with the maximum number of identifiable links. The complexity of OMP is $O(|V|^3 + |V|^2|L|)$.

3) Evaluations on ISP Topologies: We evaluate OMP on the *Internet Service Provider (ISP)* topologies collected by the Rocketfuel project [6]. Given a network topology, OMP is applied to compute the maximum number of identifiable links ($n_{\max}$) that can be achieved by two monitors. Simulation results in Table I shows that to identify all link metrics, a large fraction of nodes must serve as monitors, ranging from 30% (AT&T) to more than 60% (Abovenet). In contrast, OMP can identify a significant portion of the links using only two monitors, achieving an identification ratio greater than 0.4 for all the networks and as large as 0.97 for Level3. This is because most ISP networks contain at least one densely-connected subnetwork, allowing a properly placed monitor pair to identify all links in this subnetwork.

In conclusion, we quantitatively characterized network identifiability by developing efficient algorithms to determine and then maximize the number of identifiable links in an arbitrary network using two monitors. Simulations driven by real networks show the high efficiency of two-monitor-based tomographic solutions.

Algorithm 1: Determining Identifiable Links under Two Monitors (DIL-2M)

```
input : Connected graph G
output: All identifiable links in G
partition G into triconnected components (according existing
  algorithms in [3] and [4]) and determine the category of each
  triconnected component by Algorithm A [5];
if ∃ direct link l_{m1 m2} connecting m1 and m2 then
  | l_{m1 m2} is identifiable;
end
foreach triconnected component T_i within a biconnected
  component with 2 effective monitors do
  if T_i is in Category 1 then
    all links in T_i except for the ones incident to the
    effective monitor (e.g., m2 in Fig. 1-a) are
    identifiable;
  else if T_i is in Category 2 then
    all links in T_i except for the ones incident to the two
    effective monitors (e.g., m1 and m2 in Fig. 1-b) are
    identifiable;
  else
    T_i must be of Category 3. all links in T_i are
    identifiable if T_i is not a triangle; otherwise, link
    identifiability in triangle T_i is determined by
    Algorithm B [5];
  end
end
```

TABLE I. ISP NETWORKS

ISP	$\|L\|$	$\|V\|$	κ	$n_{\max}$	$n_{\max}/\|L\|$
Abovenet (US)	294	182	117	126	0.43
Tiscali (Europe)	404	240	138	200	0.49
Exodus (US)	434	201	85	295	0.68
Telstra (Australia)	758	318	164	480	0.63
AT&T (US)	2078	631	208	1708	0.82
Verio (US)	2821	960	408	2210	0.78
Level3 (US)	5298	624	94	5146	0.97

REFERENCES

[1] Y. Vardi, "Estimating source-destination traffic intensities from link data," *Journal of the American Statistical Assoc.*, pp. 365–377, 1996.

[2] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley, "Identifiability of link metrics based on end-to-end path measurements," in *ACM IMC*, 2013.

[3] R. Tarjan, "Depth-first search and linear graph algorithms," *SIAM Journal on Computing*, vol. 1, pp. 146–160, 1972.

[4] J. E. Hopcroft and R. E. Tarjan, "Dividing a graph into triconnected components," *SIAM Journal on Computing*, vol. 2, pp. 135–158, 1973.

[5] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley, "Link identifiability with two monitors: Auxiliary algorithms and proof of selected theorems," Technical Report, Imperial College, London, UK., Aug. 2013. [Online]. Available: http://www.commsp.ee.ic.ac.uk/%7elm110/pdf/MaTRAug13.pdf

[6] "Rocketfuel: An ISP topology mapping engine," 2002. [Online]. Available: http://www.cs.washington.edu/research/networking/rocketfuel/