

Topological Constraints on Identifying Additive Link Metrics via End-to-end Paths Measurements

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, Hawthorne, 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 through measuring accumulated end-to-end metrics over selected paths, under the assumption that link metrics are additive (e.g., delay) and constant in the measurement duration. Based on linear algebra, we know that all the link metrics can be uniquely identified when the number of linearly independent paths is equal to the number of links in the network. There lacks, however, a fundamental theory to relate the number of linearly independent paths (and thus link identifiability) to externally observable parameters such as network topology, number of monitoring nodes, and routing restrictions. The aim of this paper, therefore, is to study constraints on the network topology for identifying additive link metrics, conditioned on the number of monitoring nodes being fixed, and cycles being prohibited in constructing measurement paths. Our first main result is that it is impossible to identify all the link metrics in any network with a nontrivial topology (having more than one link) using only two monitoring nodes; nevertheless, the interior links not incident with any monitoring node might be identifiable. Our second main result is a set of necessary and sufficient conditions for identifying all the interior links using two monitoring nodes. Furthermore, we show that these conditions have a natural extension to identifying the entire network using three or more monitoring nodes. To the best of our knowledge, this is the first work providing fundamental constraints on network topology for identifying additive link metrics using end-to-end measurements on cycle-free paths.

I. INTRODUCTION

Accurate and timely knowledge of the internal state of a network (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 traffic overhead of the measurement process and the lack of support at internal network elements for taking such measurements [1]. These limitations motivate the need for *external* approaches, where we infer the states

of internal network elements by measuring the performance along selected paths.

Depending on the granularity of observations, external approaches are based on either *hop-by-hop* measurements or *end-to-end* measurements. Hop-by-hop approaches rely on special diagnostic tools to reveal fine-grained performance metrics for each link on the probed paths. Standard tools for IP networks include *ping* and *traceroute*: ping reports round-trip loss and delay, and traceroute reports these metrics for each hop on the probed path by gradually increasing the time-to-live (TTL) field of probing packets. Traceroute has known accuracy issues due to asymmetry in routes and differences in the priorities of control and data packets. Its refinement, *pathchar* [2], returns hop-by-hop link capacities, delays, and loss rates. While providing fine-grained information, the hop-by-hop approach requires a monitoring functionality, Internet Control Message Protocol (ICMP), to be supported at each internal node. Moreover, the probing packets will cause extra traffic load and potentially congestion. In applications such as military coalition networks, security policies may block hop-by-hop measurements altogether.

Alternatively, end-to-end approaches provide a solution that does not rely on the cooperation of internal network elements or the equal treatment of control/data packets. They rely on end-to-end performance metrics (e.g., end-to-end delays) experienced by data packets to solve for the corresponding hop-by-hop metrics using *network tomography*. Network tomography, originated by Vardi [3], refers to the methodology of inferring internal network characteristics through controllable end-to-end accumulated measurements. Without requiring special cooperation from internal nodes, network tomography can utilize passive measurements from data packets to obtain aggregate path-level information [4], thus not affecting traffic load.

In many cases, link metrics are *additive*, i.e., the combined metric over multiple links is the sum of individual link metrics. For instance, delays are additive, while a multiplicative metric (e.g., packet delivery ratio) can be expressed in an additive form using $\log(\cdot)$ function. For additive metrics, we can cast the problem as a system of linear equations, where the unknown variables are the link metrics, and the known constants are the end-to-end path measurements, each equal to the sum of the corresponding link metrics. Thus, network tomography

Research was partially 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.

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE 20 SEP 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012
4. TITLE AND SUBTITLE Topological Constraints on Identifying Additive Link Metrics via End-to-end Paths Measurements		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER	
	5e. TASK NUMBER	
	5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) IBM T.J. Watson Research Center,19 Skyline Drive,Hawthorne,NY,10532		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES Presented at the Annual Conference of International Technology Alliance (ACITA) 2012 held 17-21 Sep in Southampton, UK. U.S. Government or Federal Rights License		
14. ABSTRACT We investigate the problem of identifying individual link metrics in a communication network through measuring accumulated end-to-end metrics over selected paths, under the assumption that link metrics are additive (e.g., delay) and constant in the measurement duration. Based on linear algebra we know that all the link metrics can be uniquely identified when the number of linearly independent paths is equal to the number of links in the network. There lacks, however, a fundamental theory to relate the number of linearly independent paths (and thus link identifiability) to externally observable parameters such as network topology, number of monitoring nodes, and routing restrictions. The aim of this paper, therefore, is to study constraints on the network topology for identifying additive link metrics, conditioned on the number of monitoring nodes being fixed, and cycles being prohibited in constructing measurement paths. Our first main result is that it is impossible to identify all the link metrics in any network with a nontrivial topology (having more than one link) using only two monitoring nodes nevertheless, the interior links not incident with any monitoring node might be identifiable. Our second main result is a set of necessary and sufficient conditions for identifying all the interior links using two monitoring nodes. Furthermore, we show that these conditions have a natural extension to identifying the entire network using three or more monitoring nodes. To the best of our knowledge, this is the first work providing fundamental constraints on network topology for identifying additive link metrics using end-to-end measurements on cycle-free paths.		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

essentially solves for the inverse of this linear system.

Based on the model of link metrics, existing work can be divided into algebraic and statistical approaches. Algebraic approaches model link metrics as unknown constants, and use linear algebraic techniques to compute link metrics from cumulative path metrics [5, 6]. Statistical approaches model link metrics as random variables with (partially) unknown probability distributions, and apply various parametric/nonparametric techniques to estimate the link metric distributions from realizations of path metrics [7, 1, 8].

Existing work on network tomography emphasizes extracting as much information about link metrics as possible from available measurements. However, past experiences show that it is frequently impossible to uniquely identify all link metrics from path measurements [5, 6, 9]. For example, if two links (not necessarily adjacent) are simultaneously included/excluded by all measurement paths, then we can at most identify their sum metric but not the individual metrics. Generally, many measurement paths are *linearly dependent* in that some paths are linear combinations of the rest, and hence their measurements do not provide new information. From the perspective of linear algebra, link metrics are uniquely identifiable if and only if the number of *linearly independent* measurement paths equals the number of links. There is, however, a lack of basic understanding of the topological conditions that ensure identifiability, even in the simplified scenario of constant link metrics.

In this paper, we consider a fundamental but unsolved problem: under what conditions can one uniquely identify link metrics from end-to-end measurements, given that all link metrics are additive and constant? Our “constant” link metric has two interpretations: (i) the link metric changes slowly relative to end-to-end measurements so that it can be considered constant during the measurement process, or (ii) the instantaneous link quality changes dynamically, but the metric of interest is a statistical measure (e.g., mean, variance) that is stable over time¹. Although the answer to the above question is straightforward in linear algebra (the number of linearly independent paths equals the number of links), a useful answer should be in terms of externally verifiable network properties such as the network topology, the number/placement of monitors, and the routing restrictions. To this end, we develop necessary and sufficient conditions on network topology for link metric identification, assuming that the number/placement of monitors is fixed, and cycles are prohibited in constructing measurement paths between the monitors.

A. Further Discussions on Related Work

With link metrics modeled as random variables, multicast, if supported, is exploited as a measurement method with broad coverage and low overhead [10, 11]. Sub-trees and unicast are employed in [4, 12] as alternatives, due to the inflexibility of multicasting to all receivers. Employing multicast, [4, 13]

derive the necessary and sufficient conditions on the multicast tree for identifying all link metric distributions. If most links do not exhibit severe losses or delays, [7] proposes algorithms to identify the worst performing links. A novel approach proposed in [9] employs the Fourier transform of the observable path metric distributions to estimate the unobservable link metric distributions. All the above methods implicitly assume the links to be identifiable, and the multicast-based methods require multiple monitors to participate in the measurement process. In contrast, we assume unicast measurements and focus on establishing topological conditions for identifying all link metrics using the minimum number of monitors.

When link metrics are constant, [6] shows that it is challenging to solve the inverse problem due to the presence of linearly dependent paths. If all but k link metrics are *zero*, compressive sensing techniques are used to identify the k non-zero link metrics [14, 15]. If all link metrics are *binary* (normal/failed), [16] proves that the network must be $(k + 2)$ -edge-connected to identify up to k failed links using one monitor measuring cycles. For link metrics of arbitrary values, few positive results are known. If the network is directed (links have different metrics in different directions), [8] proves that not all link metrics are identifiable unless every non-isolated node is a monitor. Even if every node is a monitor, unique link identification is still impossible if measurement paths are constrained to cycles [5]. If the network is undirected (links have equal metrics in both directions), [17] derives the first necessary and sufficient conditions on the network topology for identifying all link metrics, given that monitors can measure cycles or paths possibly containing cycles. Since routing along cycles is typically prohibited in real networks, it remains open what the conditions become if only cycle-free paths can be measured. In this regard, we investigate the fundamental relationships between link identifiability, network topology, and the number/placement of monitors.

B. Summary of Contributions

To the best of our knowledge, this is the first work studying fundamental topological conditions for identifying additive link metrics using end-to-end measurements on cycle-free paths. Our contribution is three-fold:

- 1) We prove that it is generally impossible to identify all link metrics using only two monitors, irrespective of the network topology and the placement of monitors;
- 2) We establish necessary and sufficient conditions for identifying the metrics of all *interior links* (links not incident to any monitor) using two monitors: the network graph is (i) 2-edge-connected after removing any interior link and (ii) 3-vertex-connected after adding a direct link between the monitors. These conditions are shown to be verifiable in $O(|L|(|V| + |L|))$ time ($|V|$ is the number of nodes and $|L|$ is the number of links);
- 3) We transform the above result into a necessary and sufficient condition for identifying all link metrics using κ ($\kappa \geq 3$) monitors by embedding the network graph in an extended

¹In this case, end-to-end measurements are also long-term statistics, e.g., path mean/variance. In the case of variance, we need the additional assumption of independent link qualities to make the metric additive.

TABLE I
NOTATIONS IN GRAPH THEORY

Symbol	Meaning
$V(\mathcal{G}), L(\mathcal{G})$	set of nodes/links in graph $\mathcal{G}$
$ \mathcal{G} $	degree of graph $\mathcal{G}$: $ \mathcal{G} = V(\mathcal{G}) $ (number of nodes)
$ \mathcal{G} $	order of graph $\mathcal{G}$: $ \mathcal{G} = L(\mathcal{G}) $ (number of links)
$\mathcal{H}$	interior graph (see Definition 1)
$L(v)$	set of links incident to node v
$N(v)$	set of neighbors of node v
$\mathcal{G} - l$	delete a link: $\mathcal{G} - l = (V(\mathcal{G}), L(\mathcal{G}) \setminus \{l\})$, where $l \in L(\mathcal{G})$ and “ $\setminus$ ” is setminus
$\mathcal{G} + l$	add a link: $\mathcal{G} + l = (V(\mathcal{G}), L(\mathcal{G}) \cup \{l\})$, where the endpoints of link l are in $V(\mathcal{G})$
$\mathcal{G} - v$	delete a node: $\mathcal{G} - v = (V(\mathcal{G}) \setminus \{v\}, L(\mathcal{G}) \setminus L(v))$, where $v \in V(\mathcal{G})$
$\mathcal{G}_s + v$	add a node: $\mathcal{G}_s + v = (V(\mathcal{G}_s) \cup \{v\}, L(\mathcal{G}_s) \cup L_v)$, where $\mathcal{G}_s$ is a subgraph of $\mathcal{G}$, $v \in V(\mathcal{G}) \setminus V(\mathcal{G}_s)$, and L_v is the set of all links between v and nodes in $V(\mathcal{G}_s)$
$\mathcal{G} \setminus \mathcal{G}'$	From $\mathcal{G}$, delete all common nodes with $\mathcal{G}'$ and their incident links
$\mathcal{G} \cap \mathcal{G}'$	intersection of graphs: $\mathcal{G} \cap \mathcal{G}' = (V(\mathcal{G}) \cap V(\mathcal{G}'), L(\mathcal{G}) \cap L(\mathcal{G}'))$
$\mathcal{G} \cup \mathcal{G}'$	union of graphs: $\mathcal{G} \cup \mathcal{G}' = (V(\mathcal{G}) \cup V(\mathcal{G}'), L(\mathcal{G}) \cup L(\mathcal{G}'))$
$\mathcal{P}(v_0, v_k)$	simple path connecting nodes v_0 and v_k , defined as a special graph with $V(\mathcal{P}) = \{v_0, \dots, v_k\}$ and $L(\mathcal{P}) = \{v_0v_1, v_1v_2, \dots, v_{k-1}v_k\}$
$\mathcal{C}$	cycle: if $(v_0, \dots, v_k)$ ($k \geq 2$) is a sequence of nodes on a simple path $\mathcal{P}$, then $\mathcal{C} = \mathcal{P} + v_kv_0$ is a <i>cycle</i>
$\mathcal{F}$	a non-separating cycle (see Definition 2)
m_i	$m_i \in V(\mathcal{G})$ is the i -th monitor in $\mathcal{G}$
$W_l, W_{\mathcal{P}}$	metric on link l and sum metric on path $\mathcal{P}$

graph, with two virtual monitors connecting to all the real monitors. This condition can be verified in $O(|V| + |L|)$ time;

We acknowledge that not all link characteristics can be modeled as additive metrics (e.g., bit error rates). Our goal is to characterize identifiable scenarios for additive link metrics; whether a metric is additive or not is beyond the scope of this paper. Once network identifiability is confirmed, linearly independent paths have to be constructed between monitors to take measurements; we will investigate efficient path construction in future work.

The rest of the paper is organized as follows. Section II formulates the problem. Section III summarizes our main results. Sections IV–V present identifiability conditions for the case of two monitors, and Section VI addresses the case of three or more monitors. Algorithms for testing network identifiability are presented in Section VII. Finally, Section VIII concludes the paper.

II. PROBLEM FORMULATION

A. Models and Assumptions

We assume that the network topology is known and is modeled as an undirected graph² $\mathcal{G} = (V, L)$, where V and L are the sets of nodes and links, respectively. Without loss of

²In this paper, the terms *network* and *graph* are used interchangeably.




Fig. 1. Sample network with three monitors: m_1 , m_2 , and m_3 .

generality, we assume $\mathcal{G}$ is connected, as different connected components have to be monitored separately. Denote the link incident to nodes i and j by ij ; links ij and ji are assumed to have the same metric. Certain nodes in V are monitors, where measurements can be initiated/collected. We assume that each link in $\mathcal{G}$ has two distinct end-points (i.e., no self-loop), and there is at most one link connecting a pair of nodes. Table I summarizes all graph theory notations used in this paper (following the convention in [18]).

Let $n := |L|$ be the number of links in $\mathcal{G}$, $\{l_i\}_{i=1}^n$ the link set in $\mathcal{G}$, $\mathbf{w} = (W_{l_1}, \dots, W_{l_n})^T$ the column vector of all link metrics, and $\mathbf{c}$ the column vector of path measurements. Each element in $\mathbf{c}$ is the sum of link metrics along the corresponding path. We assume that the 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. In the language of graph theory, we limit measurements to *simple paths* (in contrast, a *non-simple path* may contain repeated nodes). Given all the path measurements, we have a linear system:

$$\mathbf{R}\mathbf{w} = \mathbf{c}, \quad (1)$$

where $\mathbf{R} = (R_{ij})$ is a $\gamma \times n$ *measurement matrix* (γ is the number of all simple paths between monitors), with each entry $R_{ij} \in \{0, 1\}$ denoting whether link j is present on path i .

We say a link is *identifiable* if the associated link metric can be uniquely determined from path measurements; we say the network $\mathcal{G}$ is *identifiable* if all link metrics in $\mathcal{G}$ are identifiable. Otherwise, the link or the network is said to be *unidentifiable*. Given the above linear system, $\mathcal{G}$ is identifiable if and only if $\mathbf{R}$ in (1) has full column rank, i.e., $\text{rank}(\mathbf{R}) = n$. In other words, to uniquely determine $\mathbf{w}$, there must be n linearly independent simple paths between monitors.

B. Objective

Given κ ($\kappa \geq 2$) monitors, the objective of this paper is to derive necessary and sufficient conditions on the network topology and the placement of these monitors for identifying all link metrics in $\mathcal{G}$ (or certain sub-graphs of $\mathcal{G}$) by solving the linear system (1).

C. Illustrative Example

Fig. 1 displays a sample network with 3 monitors (m_1 – m_3) and 9 links (link 1–9). To identify all 9 link metrics, 9 end-to-end paths (four $m_1 \rightarrow m_2$ paths, three $m_1 \rightarrow m_3$ paths and three $m_2 \rightarrow m_3$ paths) are constructed to form the measurement matrix $\mathbf{R}$:




Fig. 2. Reorganizing graph $\mathcal{G}$ into monitors/exterior links and interior graph (a_i and b_j may be the same node).

$$\begin{aligned}
 & m_1 \rightarrow m_2 : 1 \ 2 \\
 & \quad \quad \quad 6 \ 5 \\
 & \quad \quad \quad 6 \ 9 \ 2 \\
 & \quad \quad \quad 1 \ 4 \ 7 \ 5 \\
 & m_1 \rightarrow m_3 : 1 \ 4 \ 8 \\
 & \quad \quad \quad 6 \ 7 \ 8 \\
 & \quad \quad \quad 6 \ 9 \ 4 \ 8 \\
 & m_2 \rightarrow m_3 : 3 \\
 & \quad \quad \quad 5 \ 7 \ 8 \\
 & \quad \quad \quad 2 \ 4 \ 8
 \end{aligned}
 \Rightarrow \mathbf{R} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix},$$

where $R_{ij} = 1$ if and only if link j is on path i . Then we have $\mathbf{R}\mathbf{w} = \mathbf{c}$, where $\mathbf{c}$ is the vector of measured metrics at the destinations (m_2 for the first four measurements and m_3 for the rest six). In this example, $\mathbf{R}$ is invertible, and thus $\mathbf{w}$ can be uniquely identified, i.e., $\mathbf{w} = \mathbf{R}^{-1}\mathbf{c}$. In Fig. 1, other simple paths can be measured as well, although they do not provide further information since the rank of the measurement matrix is at most 9. However, if we remove a monitor, say m_2 , then the remaining paths can no longer form an invertible measurement matrix. Note that a path such as 67495 cannot be used for path measurement because it contains a cycle.

III. MAIN RESULTS

Our main contributions are a set of necessary and sufficient conditions for network identification that are explicitly expressed in terms of network topology and the number/placement of monitors. First, we establish a negative result that no matter how we place the monitors, we cannot identify all link metrics using only two monitors.

Theorem III.1. For any given network topology $\mathcal{G}$ with $|\mathcal{G}| \geq 2$, $\mathcal{G}$ is unidentifiable using two monitors to measure simple paths, irrespective of the placement of the monitors.

Second, we examine the two-monitor case in more detail and discover that the unidentifiability issue only applies to a small subset of links, and that the majority of links can be identified under certain conditions. Specifically, given two monitors m_1 and m_2 , we can reorganize $\mathcal{G}$ into two parts³ as illustrated in Fig. 2.

Definition 1.

1) The *interior graph* $\mathcal{H}$ of $\mathcal{G}$ is the sub-graph obtained by removing the monitors and their incident links, i.e., $\mathcal{H} := (V \setminus M, L \setminus L_M)$ for $M = \{m_1, m_2\}$ and $L_M = L(m_1) \cup L(m_2)$.

³An area with a dashed border denotes a sub-graph (the nodes/links within the dashed border are also part of the sub-graph), and a solid line denotes a link/path/cycle.




Fig. 3. (a) $\mathcal{G}$ with κ ($\kappa \geq 3$) monitors; (b) $\mathcal{G}_{ex}$ with two virtual monitors.

2) We refer to links incident to monitors, i.e., $L(m_1) \cup L(m_2)$, as *exterior links*, and the rest of the links as *interior links*.

We show that the exterior links can never be identified using two monitors (see Corollary IV.1), but the interior links can be identified under the following conditions.

Theorem III.2. Assume that the interior graph $\mathcal{H}$ of $\mathcal{G}$ is connected, and that two monitors m_1 and m_2 are used to measure simple paths. The necessary and sufficient conditions for identifying all link metrics in $\mathcal{H}$ are

- ① $\mathcal{G} - l$ is 2-edge-connected for every interior link l in $\mathcal{H}$, and
- ② $\mathcal{G} + m_1 m_2$ is 3-vertex-connected.

Third, we show that the above conditions can be naturally extended to a necessary and sufficient condition for identifying *all* link metrics in $\mathcal{G}$ using three or more monitors. Our condition is based on an *extended graph* $\mathcal{G}_{ex}$ of $\mathcal{G}$ constructed as follows. As illustrated in Fig. 3, given a graph $\mathcal{G}$ with monitors $m_1, \dots, m_\kappa$, its extended graph $\mathcal{G}_{ex}$ is obtained by adding two *virtual monitors* m'_1 and m'_2 , and 2κ *virtual links* between each pair of virtual-actual monitors. The identifiability of $\mathcal{G}$ is characterized by a simple condition on $\mathcal{G}_{ex}$ as follows.

Theorem III.3. Assume that κ ($\kappa \geq 3$) monitors are used to measure simple paths. The necessary and sufficient condition on the network topology $\mathcal{G}$ for identifying all link metrics in $\mathcal{G}$ is that the associated extended graph $\mathcal{G}_{ex}$ is 3-vertex-connected.

Finally, we develop efficient algorithms that can test whether a given placement of monitors can identify all link metrics (see Section VII) in linear time w.r.t. network size (specifically, $O(|V(\mathcal{G})| + |L(\mathcal{G})|)$) and hence are suitable for large networks.

IV. UNIDENTIFIABILITY USING TWO MONITORS

At least *two* monitors are required to identify link metrics through monitoring simple paths. In this section, we investigate if two monitors suffice to identify all link metrics in the network. Suppose two distinct nodes are selected to serve as monitors. Each measurement starts at one monitor and terminates at the other via a controllable simple path. The termination node then reports the end-to-end metric, which becomes an entry in the measurement vector $\mathbf{c}$. From the perspective of graph theory, such a network can be represented as $\mathcal{G} = (\{m_1, m_2, v_0, \dots, v_k\}, L)$, where m_1 and m_2 are the monitors, $\{v_0, \dots, v_k\}$ the non-monitors, and $|L| = n$. Let $m_1 m_2$ be the direct link incident to m_1 and m_2 . Since $m_1 m_2$ can be easily identified through a one-hop measurement, we

assume without loss of generality that $m_1 m_2 \notin L(\mathcal{G})$ (i.e., there is no direct link) in the sequel unless otherwise specified.

A. Proof of Theorem III.1

Any $\mathcal{G}$ with $||\mathcal{G}|| \geq 2$ can be reorganized as in⁴ Fig. 2. We define $a := \{a_1, a_2, \dots, a_{k_1}\} = N(m_1)$ and $b := \{b_1, b_2, \dots, b_{k_2}\} = N(m_2)$ be the sets of neighbors of m_1 and m_2 , respectively, where $k_1 := |a|$, $k_2 := |b|$ and a, b can overlap ($m_1, m_2 \notin a \cup b$).

Assuming that $\mathcal{H}$ is connected⁵ and all link metrics in $\mathcal{H}$ are known, we can reduce any equation associated with a simple path $\mathcal{P}$ between m_1 and m_2 to the form:

$$W_{m_1 a_i} + W_{b_j m_2} = \phi_{ij} \quad (2)$$

for some $a_i \in a$ and $b_j \in b$. This is obtained by rewriting the original equation $W_{m_1 a_i} + W_{\mathcal{P}'_{ij}} + W_{b_j m_2} = W_{\mathcal{P}_{ij}}$ ($\mathcal{P}'_{ij}$ is the segment of $\mathcal{P}_{ij}$ in $\mathcal{H}$) to place the unknowns to the left-hand side, where $\phi_{ij} := W_{\mathcal{P}_{ij}} - W_{\mathcal{P}'_{ij}}$. Thus, we obtain $k_1 \times k_2$ equations from all simple paths between m_1 and m_2 , each corresponding to the sum of the metrics of one link incident to m_1 and one link incident to m_2 . The corresponding reduced measurement matrix is (each column corresponding to an unknown link metric):

$$\mathbf{R} = \begin{pmatrix} \begin{array}{c|c} W_{m_1 a_1} \cdots W_{m_1 a_{k_1}} & W_{b_1 m_2} \cdots W_{b_{k_2} m_2} \\ \hline 1 & 1 \\ 1 & \\ \vdots & \ddots \\ 1 & 1 \end{array} & \begin{array}{c|c} 1 & \\ \hline 1 & 1 \\ & \ddots \\ & 1 \end{array} \\ \hline \begin{array}{c|c} & \\ \hline & 1 \\ & \\ \vdots & \vdots \\ & 1 \end{array} & \begin{array}{c|c} & \\ \hline 1 & 1 \\ & \ddots \\ & 1 \end{array} \\ \hline \begin{array}{c|c} & \\ \hline & 1 \\ & \\ \vdots & \vdots \\ & 1 \end{array} & \begin{array}{c|c} & \\ \hline 1 & 1 \\ & \ddots \\ & 1 \end{array} \end{pmatrix},$$

where the blank entries are zero. We apply the following linear transformations to $\mathbf{R}$. For each $q = 1, \dots, k_1 - 1$ and $i = 2, \dots, k_2$, replace $\text{row}(qk_2 + i)$ by $\text{row}(qk_2 + i) - \text{row}(i) - \text{row}(qk_2 + 1) + \text{row}(1)$; it can be verified that the result is a row of zeros. Ignoring rows of zeros, $\mathbf{R}$ transforms into

$$\mathbf{R}' = \left(\begin{array}{c|c} W_{m_1 a_1} \cdots W_{m_1 a_{k_1}} & W_{b_1 m_2} \cdots W_{b_{k_2} m_2} \\ \hline 1 & 1 \\ 1 & \\ \vdots & \ddots \\ 1 & 1 \\ \hline & 1 \\ & \\ \vdots & \vdots \\ & 1 \end{array} \right) \begin{array}{l} \left. \begin{array}{l} k_2 \\ \text{rows} \end{array} \right\} \\ \left. \begin{array}{l} k_1 - 1 \\ \text{rows} \end{array} \right\} \end{array}, \quad (3)$$

where rows are linearly independent, and the number of rows equals $k_1 + k_2 - 1$. When $\mathcal{H}$ is not connected, some rows in $\mathbf{R}$

⁴If certain links in $\mathcal{G}$ cannot be included in any possible paths constructed from m_1 to m_2 in Fig. 2, then these links are unidentifiable, resulting in a disconnected or one-edge-connected interior graph $\mathcal{H}$.

⁵That is, any node in $\mathcal{H}$ is reachable from every other node within $\mathcal{H}$.

may not exist because there is no simple path connecting the corresponding nodes in a and b , and the rank of $\mathbf{R}$ may be even smaller. Since there are $k_1 + k_2$ unknown variables $(W_{m_1 a_i})_{i=1}^{k_1}$ and $(W_{b_j m_2})_{j=1}^{k_2}$, they cannot be uniquely determined even if all link metrics in $\mathcal{H}$ are already known. Therefore, $\mathcal{G}$ with $||\mathcal{G}|| \geq 2$ is unidentifiable using two monitors. ■

In fact, we can show a stronger result that none of the exterior link metrics is identifiable.

Corollary IV.1. None of the exterior link metrics can be identified when using two monitors to measure simple paths.

Proof: Assume all interior link metrics are known. From the proof of Theorem III.1, we see that the transformed measurement matrix $\mathbf{R}'$ in (3) gives a maximum set of linearly independent equations (one equation per row) regarding the exterior link metrics $(W_{m_1 a_i})_{i=1}^{k_1}$ and $(W_{b_j m_2})_{j=1}^{k_2}$. To identify a metric, say $W_{m_1 a_i}$, we must be able to find a subset of $k \geq 1$ equations with k unknown variables, including $W_{m_1 a_i}$. This is, however, impossible with the equations in $\mathbf{R}'$ as any subset of k equations must contain at least $k + 1$ unknown variables. Therefore, none of the exterior link metrics can be identified. ■

B. Discussions on Paths with Cycles

At the end of Section III-B in [17], the authors raise the question whether or not monitoring non-simple paths (i.e., paths that may contain cycles) between two monitors suffices to identify all link metrics in the network. According to Corollary IV.1, the exterior links cannot be identified even if all the interior link metrics are known; allowing cycles in the interior graph $\mathcal{H}$ provides no additional information regarding the exterior links. Consequently, the answer to that question in [17] is that monitoring (simple or non-simple) paths between two monitors is *not* sufficient to identify all link metrics.

V. IDENTIFIABILITY OF INTERIOR LINK METRICS USING TWO MONITORS

In real networks, network administrators are more interested in using end-to-end measurements to infer the qualities of links that are at least *one-hop away*, since monitors can use other techniques (e.g., RTS/CTS in 802.11 MAC) to estimate their adjacent link metrics. Therefore, in this section, we only focus on the interior graph $\mathcal{H}$ and derive necessary and sufficient conditions on the network topology $\mathcal{G}$ for identifying all the links in $\mathcal{H}$ using two monitors (m_1 and m_2).

Before going into details, we first point out that it is sufficient to solve the case in which $\mathcal{H}$ is a connected graph. This is because if $\mathcal{H}$ consists of K_H ($K_H \geq 2$) connected components $\mathcal{H}_i$ ($i = 1, \dots, K_H$), we can decompose the entire graph $\mathcal{G}$ into sub-graphs $\mathcal{G}_i := \mathcal{H}_i + m_1 + m_2$, with $\mathcal{G} = \cup_{i=1}^{K_H} \mathcal{G}_i$ (see the definition of graph union in Table I). Since none of the $m_1 \rightarrow m_2$ simple paths in $\mathcal{G}_i$ can traverse $\mathcal{G}_j$ ($i \neq j$), the identification of links within different $\mathcal{G}_i$'s is mutually independent. Therefore, in the rest of this section, we assume $\mathcal{H}$ to be connected with $||\mathcal{H}|| \geq 1$. Our result can be applied to each $\mathcal{G}_i$ separately when $\mathcal{H}$ is disconnected.




Fig. 4. Illustration of Condition ①, where (d) is an enlarged view of (c). Node sets $\{\dots, a_{1i}, \dots\}$ and $\{\dots, a_{2i}, \dots\}$, $\{\dots, b_{1j}, \dots\}$ and $\{\dots, b_{2j}, \dots\}$ in (d) may have overlap, respectively.

A. Proof of Theorem III.2: Necessary Part

Suppose all links in $\mathcal{H}$ are identifiable. We prove conditions ①–② by contradiction.

a) Let $l_1 \in L(\mathcal{H})$ be an arbitrary interior link. If $\mathcal{G} - l_1$ is disconnected, then l_1 is a bridge⁶ in $\mathcal{G}$ (shown in Fig. 4-a). If each of $\mathcal{G}_1$ and $\mathcal{G}_2$ contains a monitor, then l_1 is unidentifiable by Lemma II.1 in [19]. If m_1 and m_2 are both in $\mathcal{G}_1$ (or $\mathcal{G}_2$), then l_1 cannot be employed by any measurement path (otherwise, l_1 will be used more than once), and is thus unidentifiable. Both cases contradict the assumption.

b) Suppose there is a bridge l_2 in $\mathcal{G} - l_1$. If l_2 is an exterior link, as shown in Fig. 4-b, then by Lemma II.1 in [19], its adjacent interior links $ra_i \in L(\mathcal{H})$ are unidentifiable, contradicting the assumption that all interior links are identifiable. Thus, l_2 must be an interior link. Since by a), an interior link cannot be a bridge in $\mathcal{G}$, $\{l_1, l_2\}$ must be an edge cut as shown in Fig. 4-c/d. If both m_1 and m_2 are in $\mathcal{G}_1$ as in Fig. 4-c, then all $m_1 \rightarrow m_2$ paths traversing l_1 must traverse l_2 as well. Thus we can at most identify $W_{l_1} + W_{l_2}$, but not W_{l_1} and W_{l_2} individually. If m_1 is in $\mathcal{G}_1$ and m_2 is in $\mathcal{G}_2$ as in Fig. 4-d, then all $m_1 \rightarrow m_2$ paths must traverse either l_1 or l_2 . Consider the sub-graph of $\mathcal{G}$ formed by all simple $m_1 \rightarrow m_2$ paths traversing l_1 . Since l_1 must be a bridge in this sub-graph, Lemma II.1 in [19] shows that l_1 and its adjacent links $a_{1i}r_1$, s_1b_{1j} cannot be identified from paths traversing l_1 ; similar argument applies to l_2 . Moreover, similar analysis as in the proof of this lemma shows that none of these link metrics can be identified by jointly considering measurements involving l_1 and l_2 . This contradicts the assumption that all the interior links are identifiable.

Based on a) and b), we see that $\mathcal{G} - l_1$ must be 2-edge-connected for any $l_1 \in L(\mathcal{H})$ (i.e., condition ① holds).

c) We can also prove condition ② by contradiction; see Proposition II.2 in [19] for details.

B. Proof of Theorem III.2: Sufficient Part

Given conditions ①–②, we need to show that all links in $\mathcal{H}$ are identifiable. We first introduce two types of identifiable links. The idea is to show that every interior link belongs to one of these two types.




Fig. 5. Two types of identifiable links in $\mathcal{H}$.

1) *Cross-link*: A cross-link connects nodes on two simple paths between the monitors. Specifically, as shown in Fig. 5-a, link y is a *cross-link* if $\exists$ four $m_1 \rightarrow m_2$ paths $\mathcal{P}_A$, $\mathcal{P}_B$, $\mathcal{P}_C$, and $\mathcal{P}_D$ formed by simple paths⁷ $\mathcal{P}_1, \dots, \mathcal{P}_4$:

$$\begin{cases} \mathcal{P}_A = \mathcal{P}_1 \cup \mathcal{P}_2 \\ \mathcal{P}_B = \mathcal{P}_3 \cup \mathcal{P}_4 \end{cases}, \quad \begin{cases} \mathcal{P}_C = \mathcal{P}_1 \cup \mathcal{P}_y \cup \mathcal{P}_4 \\ \mathcal{P}_D = \mathcal{P}_3 \cup \mathcal{P}_y \cup \mathcal{P}_2 \end{cases}, \quad (4)$$

such that

$$\begin{cases} |\mathcal{P}_1 \cap \mathcal{P}_2| = 1 \\ |\mathcal{P}_3 \cap \mathcal{P}_4| = 1 \end{cases}, \quad \begin{cases} |\mathcal{P}_2 \cap \mathcal{P}_3| = 0 \\ |\mathcal{P}_1 \cap \mathcal{P}_4| = 0 \end{cases}. \quad (5)$$

See Table I for definitions of graph union/intersection and $|\cdot|$; note that paths are also graphs. The constraints in (5) are used to ensure that $\mathcal{P}_A - \mathcal{P}_D$ are *simple* paths, e.g., $|\mathcal{P}_1 \cap \mathcal{P}_2| = 1$ ($\mathcal{P}_1, \mathcal{P}_2$ have no common node other than v_1) prevents cycles in $\mathcal{P}_A$. This does not mean $\mathcal{P}_1 - \mathcal{P}_4$ have to be disjoint, e.g., $\mathcal{P}_1$ and $\mathcal{P}_3$ can have common nodes. The cross-link y can be identified by

$$W_y = \frac{1}{2}(W_{\mathcal{P}_C} + W_{\mathcal{P}_D} - W_{\mathcal{P}_A} - W_{\mathcal{P}_B}). \quad (6)$$

2) *Shortcut*: A shortcut connects the endpoints of a simple path whose metric is known. Specifically, as shown in Fig. 5-b, link y is a *shortcut* if $\exists$ a simple path $\mathcal{P}_3$ whose metric has been identified such that the following $m_1 \rightarrow m_2$ paths can be formed by simple paths $\mathcal{P}_1, \dots, \mathcal{P}_3$:

$$\mathcal{P}_A = \mathcal{P}_1 \cup \mathcal{P}_y \cup \mathcal{P}_2, \quad \mathcal{P}_B = \mathcal{P}_1 \cup \mathcal{P}_3 \cup \mathcal{P}_2, \quad (7)$$

satisfying $|\mathcal{P}_1 \cap \mathcal{P}_3| = 1$, $|\mathcal{P}_2 \cap \mathcal{P}_3| = 1$, and $|\mathcal{P}_1 \cap \mathcal{P}_2| = 0$. Again the constraints are used to guarantee that $\mathcal{P}_A - \mathcal{P}_B$ are simple paths. The shortcut y can be identified by

$$W_y = W_{\mathcal{P}_A} - W_{\mathcal{P}_B} + W_{\mathcal{P}_3}. \quad (8)$$

The key to the proof is to show that each interior link can be categorized as either a cross-link or a shortcut. To this end, we introduce a special kind of cycle as follows.

Definition 2. A *non-separating cycle* in $\mathcal{G}$, denoted by $\mathcal{F}$, is an induced sub-graph⁸ such that: (i) $\mathcal{F}$ is a cycle (see definition in Table I), and (ii) $\mathcal{F}$ does not separate any node from a monitor, i.e., each connected component in $\mathcal{G} \setminus \mathcal{F}$ contains at least one monitor.

For example, there are four non-separating cycles in Fig. 6: $v_1v_2v_3v_1$, $v_4v_3v_2v_5v_4$, $m_1v_1v_3v_4m_1$, and $v_5v_2m_2v_5$. Cycle

⁷Here $\mathcal{P}_y$ is the 1-hop path traversing a link y .

⁸An *induced sub-graph* $\mathcal{G}'$ of $\mathcal{G}$ is a sub-graph that contains all the original links as long as both end-points of the links are in $\mathcal{G}'$.

⁶A link separating its end-points is a bridge [18].




Fig. 6. Sample network with identifiable interior graph.




Fig. 7. Possible cases of interior link vw : (a) Case A, (b) Case B-1, (c) Case B-2.

$v_4v_3v_1v_2v_5v_4$ is not a non-separating cycle as it is not induced (due to link v_2v_3), neither is $v_4m_1v_1v_2v_5v_4$ as it separates v_3 from monitors.

Our *first observation* is that if we can find two non-separating cycles C_1 and C_2 as shown in Fig. 7-a, sharing only the link of interest (vw) and each connecting to a different monitor by a disjoint simple path, then link vw is a cross-link and can be identified as in (6) (where $\mathcal{P}_A$ – $\mathcal{P}_D$ are constructed by (4), using $\mathcal{P}_1$ – $\mathcal{P}_4$ marked in the figure). The challenge is that such cycles do not exist for every link. We have shown in Lemma II.4 in [19] that two other cases may occur, as shown in Fig. 7: (b) any path from $C_1 - v - w$ to a monitor intersects with $C_2 - v - w$, or (c) C_1 and C_2 have a common vertex other than v and w . The lemma guarantees that these three cases are complete, i.e., any interior link belongs to one of these cases.

Our *second observation* is that for any link not identifiable by (6) (Case B-1 or B-2), we can find a detour path connecting its endpoints, whose links are all cross-links and thus can be identified by (6). Consider the example in Fig. 5-b. We show in Lemma II.5 (a) in [19] that each Case-B (including B-1, B-2) link y resides on a non-separating cycle $\mathcal{P}_3 + y$ in the interior graph, where all the other links on this cycle are of Case A (cross-links). Furthermore, we show in Lemma II.5 (b) in [19] that there exist disjoint simple paths $\mathcal{P}_1$ and $\mathcal{P}_2$ connecting the end-points of this Case-B link to different monitors, each intersecting with the cycle only at the end-point. Thus, the Case-B link y is a shortcut and can be identified by (8) after identifying all the Case-A links. Therefore, any interior link, either Case-A or Case-B, is identifiable. See details in [19]. ■

As an example, Fig. 6 displays a network satisfying Conditions ① and ②.

VI. IDENTIFIABILITY USING THREE OR MORE MONITORS

Since two monitors are not sufficient to identify all link metrics in $\mathcal{G}$, in this section, we explore the case of three or more monitors. Lessons learned in Section V suggest that it is easier to identify links that are one-hop away from the monitors. This observation inspires us to construct an extended graph $\mathcal{G}_{ex}$ of $\mathcal{G}$ (see Fig. 3), so that all links of interest (actual links in $\mathcal{G}$) become one-hop away from the virtual monitors m'_1 and m'_2 . This construction immediately converts the problem

of identifying $\mathcal{G}$ using κ monitors to a problem of identifying the interior graph of $\mathcal{G}_{ex}$ using 2 monitors. Therefore, we can use Theorem III.2 to obtain the following result.

Lemma VI.1. Employing κ ($\kappa \geq 3$) monitors to measure simple paths, the necessary and sufficient condition on the network topology $\mathcal{G}$ for identifying all link metrics in $\mathcal{G}$ is that the associated extended graph $\mathcal{G}_{ex}$ has an identifiable interior graph, i.e., $\mathcal{G}_{ex}$ satisfies Conditions ① and ② in Theorem III.2.

Proof: Since $\mathcal{G}$ is the interior graph of $\mathcal{G}_{ex}$, it suffices to show that the information obtained by the real monitors is the same as that obtained by the virtual monitors. First, it is easy to see that for any $m'_1 \rightarrow m'_2$ measurement $W_{m'_1 m_i \dots m_j m'_2}$ involving links in $\mathcal{G}$, the information relevant for identifying $\mathcal{G}$ is captured by an $m_i \rightarrow m_j$ measurement $W_{m_i \dots m_j}$ ($i, j \in \{1, \dots, \kappa\}, i \neq j$). It remains to show that any measurement between the real monitors $m_1, \dots, m_\kappa$ can be obtained from measurements between m'_1 and m'_2 . To this end, consider a path $m_1 e_1 m_i$ ($i \in \{2, \dots, \kappa\}$) as shown in Fig. 3-b. Four simple paths between m'_1 and m'_2 can be constructed: $\mathcal{P}_A = m'_1 m_i m'_2$, $\mathcal{P}_B = m'_1 m_1 m'_2$, $\mathcal{P}_C = m'_1 m_i e_1 m_1 m'_2$, and $\mathcal{P}_D = m'_1 m_1 e_1 m_i m'_2$. Viewing $m_1 e_1 m_i$ as a “cross-link”, $W_{m_1 e_1 m_i}$ can be computed from measurements on these four paths via (6) (replacing W_y by $W_{m_1 e_1 m_i}$). ■

Furthermore, the special structure of $\mathcal{G}_{ex}$ allows us to consolidate the two conditions ① and ② into a single condition as stated in Theorem III.3, whose proof is given below.

A. Proof of Theorem III.3

Based on the construction of $\mathcal{G}_{ex}$, we can prove (see Proposition II.7–II.8 in [19]) that $\mathcal{G}_{ex}$ satisfies Conditions ① and ② in Theorem III.2 if and only if $\mathcal{G}_{ex}$ is both 3-edge-connected and 3-vertex-connected. According to Proposition 1.4.2 in [18], a 3-vertex-connected graph is also 3-edge-connected. Thus, the necessary and sufficient conditions in Lemma VI.1 can be simplified to $\mathcal{G}_{ex}$ being 3-vertex-connected. ■

VII. TESTING AND ENSURING IDENTIFIABILITY

The conditions we have derived have broader impact than mere theoretical interest. A major benefit of characterizing network identifiability in terms of network topology is that we can leverage existing graph-processing algorithms to efficiently test or ensure the identifiability of a given network. In this section, we present efficient algorithms that can test the identifiability of a given network under given monitor placement.

A. Efficient Identifiability Test

The first question we want to answer is: given a network topology $\mathcal{G}$ and a placement of $\kappa \geq 2$ monitors, how can we efficiently determine if $\mathcal{G}$ is identifiable or not? If $\kappa = 2$, then we know from Theorem III.1 that it is impossible to identify the entire $\mathcal{G}$. Nevertheless, we can test whether the interior graph is identifiable using Conditions ① and ② in Theorem III.2, which transform into multiple tests of edge/vertex connectivity.

The problem of determining whether a given graph is k -edge/vertex-connected has been well studied. Specifically, fast algorithms have been proposed to test if a graph is: (i) 2-edge-connected [20], or (ii) 3-vertex-connected [21], both in time $O(|V| + |L|)$ ($|V|$: number of nodes; $|L|$: number of links). Using these algorithms, we can test the identifiability of the interior graph of $\mathcal{G}$ as follows:

- 1) for each interior link l : apply the 2-edge-connectivity test in [20] to $\mathcal{G} - l$; $\mathcal{G}$ is unidentifiable if the test fails;
- 2) apply the 3-vertex-connectivity test in [21] to $\mathcal{G} + m_1 m_2$; $\mathcal{G}$ is unidentifiable if the test fails;

$\mathcal{G}$ is identifiable if all the tests succeed. The overall complexity is $O(|L(\mathcal{G})|(|V(\mathcal{G})| + |L(\mathcal{G})|))$.

Similarly, if $\kappa \geq 3$, then we can test the identifiability of the entire $\mathcal{G}$ using the condition in Theorem III.3:

- 1) construct the extended graph $\mathcal{G}_{ex}$ as in Fig. 3;
- 2) apply the 3-vertex-connectivity test in [21] to $\mathcal{G}_{ex}$; $\mathcal{G}$ is identifiable if the test succeeds, and unidentifiable otherwise.

The complexity of this algorithm is $O(|V(\mathcal{G}_{ex})| + |L(\mathcal{G}_{ex})|)$, which can be reduced⁹ to $O(|V(\mathcal{G})| + |L(\mathcal{G})|)$.

VIII. CONCLUSION

In this paper, we explore the fundamental conditions on network topology for identifying link metrics using end-to-end measurements along simple paths. We show that it is impossible to identify all the links with only two monitors, irrespective of the network topology and the placement of monitors. In some cases, however, it is possible to identify the interior links that are at least one hop away from the monitors, for which we derive the necessary and sufficient conditions. We further study the case of three or more monitors and develop the corresponding necessary and sufficient conditions for identifying all link metrics. All these conditions can be verified efficiently.

REFERENCES

- [1] F. Lo Presti, N. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Transactions on Networking*, vol. 10, no. 6, pp. 761–775, Dec. 2002.
- [2] A. B. Downey, "Using pathchar to estimate internet link characteristics," in *IEEE SIGCOMM*, 1999.
- [3] Y. Vardi, "Estimating source-destination traffic intensities from link data," *Journal of the American Statistical Assoc.*, pp. 365–377, 1996.
- [4] E. Lawrence and G. Michailidis, "Network tomography: A review and recent developments," *Frontiers and Statistics*, vol. 54, 2006.
- [5] O. Gurewitz and M. Sidi, "Estimating one-way delays from cyclic-path delay measurements," in *IEEE INFOCOM*, 2001.
- [6] Y. Chen, D. Bindel, and R. H. Katz, "An algebraic approach to practical and scalable overlay network monitoring," in *ACM SIGCOMM*, 2004.
- [7] N. Duffield and F. Lo Presti, "Multicast inference of packet delay variance at interior network links," in *IEEE INFOCOM*, 2000.
- [8] Y. Xia and D. Tse, "Inference of link delay in communication networks," *IEEE Journal of Selected Areas in Communications*, 2006.
- [9] A. Chen, J. Cao, and T. Bu, "Network Tomography: Identifiability and Fourier domain estimation," in *IEEE INFOCOM*, 2007.

- [10] A. Adams, T. Bu, T. Friedman, J. Horowitz, D. Towsley, R. Caceres, N. Duffield, F. Presti, and V. Paxson, "The use of end-to-end multicast measurements for characterizing internal network behavior," *IEEE Communications Magazine*, vol. 38, no. 5, pp. 152–159, May 2000.
- [11] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu, "Network tomography: recent developments," *Statistical Science*, 2004.
- [12] M.-F. Shih and A. Hero, "Unicast inference of network link delay distributions from edge measurements," in *IEEE ICASSP*, 2001.
- [13] T. Bu, N. Duffield, and F. Lo Presti, "Network tomography on general topologies," in *ACM SIGMETRICS*, 2002.
- [14] M. Firooz and S. Roy, "Network tomography via compressed sensing," in *IEEE GLOBECOM*, 2010.
- [15] W. Xu, E. Mallada, and A. Tang, "Compressive sensing over graphs," in *IEEE INFOCOM*, 2011.
- [16] S. Ahuja, S. Ramasubramanian, and M. Krunz, "SRLG failure localization in all-optical networks using monitoring cycles and paths," in *IEEE INFOCOM*, 2008.
- [17] A. Gopalan and S. Ramasubramanian, "On identifying additive link metrics using linearly independent cycles and paths," *IEEE/ACM Transactions on Networking*, vol. PP, no. 99, 2011.
- [18] R. Diestel, *Graph theory*. Springer-Verlag Heidelberg, New York, 2005.
- [19] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley, "Identification of additive link metrics: Proof of selected lemmas and propositions," Technical Report, Imperial College, London, UK, May 2012. [Online]. Available: <http://www.commsp.ee.ic.ac.uk/~lm110/pdf/MaTechreport12.pdf>
- [20] R. E. Tarjan, "A note on finding the bridges of a graph," *Information Processing Letters*, pp. 160–161, 1974.
- [21] J. E. Hopcroft and R. E. Tarjan, "Dividing a graph into triconnected components," *SIAM Journal on Computing*, vol. 2, pp. 135–158, 1973.

⁹This is because $|V(\mathcal{G}_{ex})| + |L(\mathcal{G}_{ex})| = |V(\mathcal{G})| + 2 + |L(\mathcal{G})| + 2\kappa \leq |V(\mathcal{G})| + 2 + |L(\mathcal{G})| + 2|V(\mathcal{G})| = 3|V(\mathcal{G})| + |L(\mathcal{G})| + 2$.