

Inferring Link Metrics From End-To-End Path Measurements: Identifiability and Monitor Placement

Liang Ma, *Student Member, IEEE*, Ting He, *Senior Member, IEEE*, Kin K. Leung, *Fellow, IEEE*, Ananthram Swami, *Fellow, IEEE*, and Don Towsley, *Fellow, IEEE, ACM*

Abstract—We investigate the problem of identifying individual link metrics in a communication network from end-to-end path measurements, under the assumption that link metrics are additive and constant. To uniquely identify the link metrics, the number of linearly independent measurement paths must equal the number of links. Our contribution is to characterize this condition in terms of the network topology and the number/placement of monitors, under the constraint that measurement paths must be cycle-free. Our main results are: 1) it is generally impossible to identify all the link metrics by using two monitors; 2) nevertheless, metrics of all the interior links not incident to any monitor are identifiable by two monitors if the topology satisfies a set of necessary and sufficient connectivity conditions; 3) these conditions naturally extend to a necessary and sufficient condition for identifying all the link metrics using three or more monitors. We show that these conditions not only facilitate efficient identifiability tests, but also enable an efficient algorithm to place the minimum number of monitors in order to identify all link metrics. Our evaluations on both random and real topologies show that the proposed algorithm achieves identifiability using a much smaller number of monitors than a baseline solution.

Index Terms—Identifiability condition, linear algebraic model, monitor placement, network tomography.

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)

Manuscript received November 24, 2013; revised April 18, 2014; accepted May 23, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor R. Teixeira. Date of publication June 18, 2014; date of current version August 14, 2014. This work was supported in part by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement No. 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.

L. Ma and K. K. Leung are with the Imperial College London, London SW7 2BT, U.K. (e-mail: l.ma10@imperial.ac.uk; kin.leung@imperial.ac.uk).

T. He is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: the@us.ibm.com).

A. Swami is with the Army Research Laboratory, Adelphi, MD 20783 USA (e-mail: a.swami@ieee.org).

D. Towsley is with the University of Massachusetts Amherst, Amherst, MA 01003 USA (e-mail: towsley@cs.umass.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2014.2328668

is, however, not always feasible due to the traffic overhead of the measurement process and the lack of support at internal network elements for making 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 from a subset of nodes with monitoring capabilities, hereafter referred to as *monitors*.

Depending on the granularity of observations, external approaches can be classified as *hop-by-hop* approaches or *end-to-end* approaches. The former rely on special diagnostic tools such as *traceroute*, *pathchar* [2], and *Network Characterization Service* (NCS) [3] to reveal fine-grained performance metrics of individual links by sending active probes. Traceroute reports delay for each hop on the probed path by gradually increasing the time-to-live (TTL) field of probing packets. Its refinement, *pathchar*, returns hop-by-hop capacities, delays, and loss rates. A later advancement, NCS, also returns available capacities on each link. While providing fine-grained information, the above tools require that Internet Control Message Protocol (ICMP) be supported at each internal node. Even then, they suffer inaccuracies caused by asymmetry in routes and different priorities of ICMP and data packets. Moreover, these tools can generate a large number of probing packets, causing extra load and, potentially, congestion. In risk-sensitive applications, security policies may even block hop-by-hop measurements.

Alternatively, the end-to-end approach provides a solution that does not rely on the cooperation of internal network elements or the equal treatment of control/data packets. It relies 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 [4] refers to the methodology of inferring internal network characteristics through end-to-end measurements. Without requiring special cooperation from internal nodes, network tomography can utilize measurements from data packets to obtain path-level information [5], thus reducing the need for active probes.

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 by using the $\log(\cdot)$ function. For additive metrics, we can model the problem as that of solving 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 along a path. Thus, network tomography essentially solves this linear system of equations.

Existing work on network tomography emphasizes extracting as much information about link metrics as possible from available measurements. However, past experience shows that it is frequently impossible to uniquely identify all link metrics from path measurements [6]–[8]. For example, if two links (not necessarily adjacent) always appear together in 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 two closely related problems: 1) Under what conditions can one uniquely identify all link metrics from end-to-end measurements? 2) Given an arbitrary network topology, how can one place monitors to satisfy the above conditions using the minimum number of monitors? We study both problems in the context of *controllable, cycle-free* measurements, i.e., monitors can direct measurement packets to selected paths as long as they do not contain cycles. Such routing is generally supported in networks under common administration (e.g., single-ISP networks), or overlay networks formed by overlay nodes and physical paths between these nodes (modeled as overlay links), where controllable routing can be achieved by enabling strict (or loose) source routing (an IP option) [9] at all network nodes during measurement. More importantly, the assumption of controllable, cycle-free measurements captures capabilities of a new generation of networks performing *software-defined networking* (SDN) [10]–[12]. In SDN, routing is performed on a per-flow basis in two steps: a route setup phase where a (logically) centralized SDN controller specifies the path of a flow by setting entries in the forwarding tables of switches along the path, and a data forwarding phase where data of the flow are forwarded hop by hop by switches according to their forwarding tables. Therefore, monitors, in cooperation with the SDN controller, can dictate paths of measurement packets during route setup, while the cycle-free constraint precludes formation of endless cycles in the data forwarding phase. In all of the above scenarios,¹ network tomography can assist in accurately monitoring network performance experienced by data traffic while reducing measurement overhead. Throughout this paper, we assume that all link metrics are additive and constant. Our “constant” link metric refers to one that either changes slowly relative to the measurement process, or that is a statistical characteristic (e.g., mean, variance) of the link that remains constant over time.²

Although the linear algebraic answer to the first question is straightforward (the number of linearly independent paths equals the number of links), a useful answer should be expressed in terms of externally verifiable network properties

¹Note that SDN by itself does not dictate reporting of per-link metrics. Furthermore, SDN does not overcome issues with policies in coalition networks that prevent reporting of per-link metrics; thus, network tomography will still be required.

²In this case, end-to-end measurements are also statistical characteristics, e.g., path mean/variance. In the case of variance, we also need the independence between link qualities to make the metric additive.

such as the network topology and the number/placement of monitors. To this end, we establish necessary and sufficient conditions on network topology and monitor placement for uniquely identifying link metrics from controllable, cycle-free measurements between the monitors. We further develop an efficient algorithm that leverages the established conditions to place monitors in a given network such that all link metrics can be uniquely identified by using the minimum number of monitors.

A. Further Discussions on Related Work

Based on the model of link metrics, existing work can be broadly classified into statistical and algebraic approaches. 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 [13], [1], [14]. Algebraic approaches consider link metrics as unknown constants and use linear algebraic techniques to compute link metrics from cumulative path metrics [6], [7].

With link metrics modeled as random variables, multicast, if supported, can be exploited as a measurement method with broad coverage and low overhead [15], [16]. Subtrees and unicast are employed in [5] and [17] as alternatives, due to the inflexibility and impracticality of multicasting to all receivers. Employing multicast, [5] and [18] 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, [19] proposes algorithms to identify the worst performing links. A novel approach proposed in [8] 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 by using the *minimum* number of monitors.

For constant link metrics, [7] shows that it is challenging to solve the inverse problem due to the presence of linearly dependent paths. When all but k link metrics are *zero*, compressive sensing techniques are used to identify the k nonzero link metrics [20], [21]. If all link metrics are *binary* (normal/failed), [22] proves that the network must be $(k + 2)$ -edge-connected to identify up to k failed links by measuring cycles at a single monitor. For arbitrary valued link metrics, few positive results are known. If the network is directed (links have different metrics in different directions), [14] proves that not all link metrics are identifiable unless every nonisolated node is a monitor. Even if every node is a monitor, unique link identification is still impossible if measurement paths are constrained to cycles [6]. If the network is undirected (links have equal metrics in both directions), [23] 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. A similar study in [24] characterizes the minimum number of measurements needed to identify a broader set

of link metrics (including both additive and nonadditive metrics), under the stronger assumption that measurement paths can contain repeated links. Since routing along cycles is typically prohibited in real networks, it remains open as to what the conditions are 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. Once identifiability is confirmed, a follow-up problem is to construct linearly independent paths between monitors to perform measurements, for which we have developed an efficient path construction algorithm in [25].

Little is known when it comes to placing monitors to ensure network identifiability. In [26] and [27], the problem of placing the minimum number of monitors to identify all link metrics under *uncontrollable* routing is proven to be NP-hard, and the NP-hardness persists even if some nodes control their local routing policy [28]. In contrast, we show in this paper that under *controllable* cycle-free routing, this problem can be solved optimally in linear time.

B. Summary of Contributions

We study here for the first time the fundamental topological conditions for identifying additive link metrics by using end-to-end measurements on cycle-free paths. Our contributions are fourfold.

- 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. These conditions are shown to be verifiable in $O(|L|(|V| + |L|))$ time, where $|V|$ is the number of nodes and $|L|$ is the number of links.
- 3) We extend the above result to a necessary and sufficient condition for identifying all link metrics using κ ($\kappa = 3$ is the minimum requirement, but typically many more) monitors. This condition can be verified in $O(|V| + |L|)$ time.
- 4) We propose an algorithm to place monitors in an arbitrary network in $O(|V| + |L|)$ time, which guarantees the identifiability of all link metrics with the minimum number of monitors. Our evaluations on both randomly generated and real topologies show that the proposed algorithm requires a substantially smaller number of monitors than a baseline solution.

We note that our goal is to characterize identifiable scenarios for constant additive link metrics, and questions for identifying nonadditive link metrics (e.g., bandwidth) or estimating link parameters in cases where link metrics are modeled as random variables are beyond the scope of this paper.

The rest of the paper is organized as follows. Section II formulates the problem. Section III summarizes our main results. Sections IV and 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 and deploying monitors are presented in Section VII. Finally, Section VIII concludes the paper.

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, denoted by n)
\mathcal{H}	interior graph (see Definition 1)
$L(v)$	set of links incident to 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 end-points 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 sub-graph 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 cycle
\mathcal{F}	a non-separating cycle (see Definition 4)
m_i	$m_i \in V(\mathcal{G})$ is the i -th monitor in \mathcal{G}
$W_l, C_{\mathcal{P}}$	metric on link l and sum metric on path \mathcal{P}

II. PROBLEM FORMULATION

A. Models and Assumptions

We assume that the network topology is known³ and model it as an undirected graph $\mathcal{G} = (V, L)$, where V and L are the sets of nodes and links⁴, respectively. Note that graph \mathcal{G} can represent a logical topology where each link in \mathcal{G} corresponds to a combination of physical links/nodes, e.g., in an overlay network. Without loss of generality, we assume \mathcal{G} is connected, as different connected components have to be monitored separately. Denote the link incident to nodes u and v by uv ; links uv and vu are assumed to have the same metric. Certain nodes in V are monitors and can initiate/collect measurements. We assume that each link in \mathcal{G} has two distinct endpoints (i.e., no self-loop), and no two links connect to the same pair of nodes. Lastly, no

³Network topology can be obtained through exchanging control packets [29], [30] explicitly. Note that network internal states cannot be inferred from these control packets in an accurate manner due to the unequal treatment of control and data packets.

⁴We use the terms *network* and *graph*, *link* and *edge*, and *node* and *vertex* interchangeably.




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

link metrics in \mathcal{G} are known prior to taking measurements. Table I summarizes all graph theory notations used in this paper (following the convention of [31]).

Let $n := |L|$ denote the number of links in \mathcal{G} , $\{l_i\}_{i=1}^n$ the set of links in \mathcal{G} , $\mathbf{w} = (W_{l_1}, \dots, W_{l_n})^T$ the column vector of all link metrics, and $\mathbf{c} = (C_{\mathcal{P}_1}, \dots, C_{\mathcal{P}_\gamma})^T$ the column vector of all available path measurements, where γ is the number of measurement paths and $C_{\mathcal{P}_i}$ is the sum of link metrics along measurement path \mathcal{P}_i . We assume that monitors can control the routing of measurement packets 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 *nonsimple path* may contain repeated nodes). The path measurements are expressed in terms of the unknown link metrics through the following linear system:

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

where $\mathbf{R} = (R_{ij})$ is a $\gamma \times n$ *measurement matrix*, with each entry $R_{ij} \in \{0, 1\}$ denoting whether link l_j is present on path \mathcal{P}_i . The network tomography problem is to invert this linear system to solve for \mathbf{w} given \mathbf{R} and \mathbf{c} .

A link is *identifiable* if the associated link metric can be uniquely determined from path measurements; network \mathcal{G} is identifiable if all links 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,⁵ 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. We note that within an unidentifiable network, some links may still be identifiable, for which we have built algorithms for determining partial network identifiability; see [32] for details.

B. Objective

Given a network topology \mathcal{G} , the first objective of this paper is to derive necessary and sufficient conditions for identifying all link metrics in \mathcal{G} (or certain subgraphs of \mathcal{G}) under a given placement of κ ($\kappa \geq 2$) monitors by solving the linear system (1). Our second objective is to compute a placement of the minimum number of monitors in \mathcal{G} that enables the identification of all link metrics.

C. Illustrative Example

Fig. 1 displays a sample network with three monitors ($m_1 - m_3$) and 11 links (links $l_1 - l_{11}$). To identify all 11 link metrics, eleven end-to-end paths (one $m_1 \rightarrow m_2$ path, seven $m_1 \rightarrow m_3$

⁵While a full-rank \mathbf{R} is not necessarily robust to invert, we have developed an efficient algorithm to construct \mathbf{R} such that \mathbf{w} can be solved without inverting \mathbf{R} explicitly [25].




Fig. 2. Organizing graph \mathcal{G} into exterior links and interior graph ($\{a_i\}_{i=1}^{k_1}$ and $\{b_j\}_{j=1}^{k_2}$ may have overlap).

paths, and three $m_3 \rightarrow m_2$ paths) are constructed to form the measurement matrix \mathbf{R}

$$\begin{array}{l} m_1 \rightarrow m_2 : \\ l_1 \ l_4 \ l_8 \ l_{11} \\ m_1 \rightarrow m_3 : \\ l_1 \ l_4 \ l_7 \\ l_2 \ l_6 \\ l_2 \ l_3 \ l_4 \ l_7 \\ l_2 \ l_5 \ l_8 \ l_{10} \\ l_2 \ l_5 \ l_7 \\ l_1 \ l_3 \ l_6 \\ l_1 \ l_4 \ l_5 \ l_6 \\ m_3 \rightarrow m_2 : \\ l_9 \\ l_{10} \ l_{11} \\ l_6 \ l_5 \ l_8 \ l_{11} \end{array} \Rightarrow \mathbf{R} = \begin{pmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}$$

where $R_{ij} = 1$ if and only if link l_j is on path \mathcal{P}_i . Then, we have $\mathbf{R}\mathbf{w} = \mathbf{c}$, where \mathbf{c} is the vector of end-to-end measurements taken at the destination monitors. 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 measurement matrix already has full rank. However, if we remove a monitor, say m_3 , then it can be verified that the remaining paths no longer form an invertible measurement matrix. Note that a path such as $l_2 l_5 l_4 l_3 l_6 l_9$ is not allowed 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 (proofs of theorems can be found in Sections IV–VI). To begin, we first establish a negative result that no matter where 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 $n \geq 2$ (n is the number of links), \mathcal{G} is unidentifiable with two monitors, irrespective of their placement.

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 organize \mathcal{G} into two parts⁶ as illustrated in Fig. 2.

⁶An area with a dashed border denotes a subgraph (nodes/links on the dashed border are also part of the subgraph, e.g., $\{a_i\}_{i=1}^{k_1}$ and $\{b_j\}_{j=1}^{k_2}$ in Fig. 2 are part of \mathcal{H}), and a solid line denotes a link/path/cycle.




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

Definition 1: Given a network \mathcal{G} with two monitors m_1 and m_2 , we have the following.

- 1) The *interior graph* \mathcal{H} of \mathcal{G} is the subgraph obtained by removing the monitors and their incident links, i.e., $\mathcal{H} := (V \setminus M, L \setminus L_M)$ where $M = \{m_1, m_2\}$ and $L_M = L(m_1) \cup L(m_2)$.
- 2) We refer to links incident to monitors, i.e., $L(m_1) \cup L(m_2)$, as *exterior links*, and the remaining 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} (with $\|\mathcal{H}\| \geq 1$) of \mathcal{G} under a given placement of two monitors (m_1 and m_2) is connected and direct link $m_1 m_2$ (incident to m_1 and m_2) does not exist in $L(\mathcal{G})$. The necessary and sufficient conditions for identifying all link metrics in \mathcal{H} are the following.

- ① $\mathcal{G} - l$ is 2-edge-connected for every interior link l in \mathcal{H} .
- 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. This condition is based on an *extended graph* \mathcal{G}_{ex} obtained from \mathcal{G} as follows. As illustrated in Fig. 3, given a graph \mathcal{G} with κ monitors, 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 $\kappa (\kappa \geq 3)$ 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} be 3-vertex-connected.

Finally, we develop efficient algorithms that can: 1) test whether a given placement of monitors can identify all link metrics; and 2) compute a placement of the minimum number of monitors needed to identify all link metrics (see Section VII). Both algorithms run in linear time w.r.t. network size (i.e., $O(|V(\mathcal{G})| + |L(\mathcal{G})|)$) and hence are suitable for large networks.

IV. UNIDENTIFIABILITY WITH TWO MONITORS

At least *two* monitors are required to identify link metrics by monitoring simple paths. In this section, we investigate if two monitors suffice to identify all link metrics in the network. Suppose that 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, $\{v_0, \dots, v_k\}$ are the nonmonitors, and $|L| = n$. Let $m_1 m_2$ be a direct link between m_1 and m_2 (if it exists). 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 Sections IV and V, where only two monitors are used.

A. Proof of Theorem III.1

Any \mathcal{G} with $\|\mathcal{G}\| \geq 2$ can be organized as in⁷ Fig. 2. Let $N(v)$ denote the set of neighboring nodes of node v . We define $A := \{a_1, a_2, \dots, a_{k_1}\} = N(m_1)$ and $B := \{b_1, b_2, \dots, b_{k_2}\} = N(m_2)$ to 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 connected 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 (see Table I for notations)

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

for some $a_i \in A$ and $b_j \in B$. This is obtained by rewriting the original equation $W_{m_1 a_i} + C_{\mathcal{P}'_{ij}} + W_{b_j m_2} = C_{\mathcal{P}_{ij}}$ (\mathcal{P}'_{ij} is the segment of \mathcal{P}_{ij} in \mathcal{H}) to place the unknowns on the left-hand side, and setting $\phi_{ij} := C_{\mathcal{P}_{ij}} - C_{\mathcal{P}'_{ij}}$. Thus, we obtain $k_1 \times k_2$ equations from all of the 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)

$$\mathbf{R} = \begin{pmatrix} W_{m_1 a_1} & \dots & W_{m_1 a_{k_1}} & W_{b_1 m_2} & \dots & W_{b_{k_2} m_2} \\ \hline 1 & & & 1 & & \\ 1 & & & & 1 & \\ \vdots & & & & & \ddots \\ 1 & & & & & 1 \\ \hline & 1 & & 1 & & \\ & 1 & & & 1 & \\ & \vdots & & & & \ddots \\ & 1 & & & & 1 \\ \hline & & \ddots & & \vdots & \\ & & & 1 & & \\ & & & 1 & & \\ & & & \vdots & & \\ & & & 1 & & 1 \end{pmatrix} \quad (3)$$

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 row $(qk_2 + i)$ by row $(qk_2 + i) - \text{row}(i) -$

⁷If 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} .

$\text{row}(qk_2 + 1) + \text{row}(1)$; it can be verified that the result is a row of zeros. Ignoring rows of zeros, \mathbf{R} transforms into

$$\mathbf{R}' = \left(\begin{array}{c|c} W_{m_1 a_1} \cdots W_{m_1 a_{k_1}} & W_{b_1 m_2} \cdots W_{b_{k_2} m_2} \\ \hline \left. \begin{array}{l} 1 \\ 1 \\ \vdots \\ 1 \end{array} \right\} \begin{array}{l} k_2 \\ \text{rows} \end{array} & \left. \begin{array}{l} 1 \\ 1 \\ \vdots \\ 1 \end{array} \right\} \begin{array}{l} k_1 - 1 \\ \text{rows} \end{array} \end{array} \right) \quad (4)$$

where the rows are linearly independent, and the number of rows equals $k_1 + k_2 - 1$. This transformation preserves the rank of \mathbf{R} because \mathbf{R} can be reconstructed from \mathbf{R}' as follows: The first k_2 rows in \mathbf{R} and \mathbf{R}' are identical; $\text{row}(qk_2 + 1)$ in \mathbf{R} equals $\text{row}(k_2 + q)$ in \mathbf{R}' ($q = 1, \dots, k_1 - 1$); $\text{row}(qk_2 + i)$ in \mathbf{R} equals $\text{row}(i) + \text{row}(k_2 + q) - \text{row}(1)$ in \mathbf{R}' ($q = 1, \dots, k_1 - 1, i = 2, \dots, k_2$). Hence, the rank of \mathbf{R} equals $k_1 + k_2 - 1$. When \mathcal{H} is not connected, some rows in \mathbf{R} 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 ($W_{m_1 a_i}^{k_1}_{i=1}$ and $W_{b_j m_2}^{k_2}_{j=1}$), 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 that none of the exterior links is identifiable because the metric of any one exterior link will add a linearly independent row to \mathbf{R}' in (4), thus yielding full column rank of \mathbf{R}' (i.e., the metrics of the remaining exterior links are also determinable).

Corollary IV.1: None of the exterior links (except $m_1 m_2$) can be identified with two monitors.

Proof: Assume that all the interior link metrics are known. We prove Corollary IV.1 by contradiction. From the proof of Theorem III.1, we see that the transformed measurement matrix \mathbf{R}' 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 $\{W_{b_j m_2}\}_{j=1}^{k_2}$. Suppose there exists one exterior link metric, say $W_{m_1 a_i}$, which is identifiable. Since each row in \mathbf{R}' only involves two exterior link metrics, it is easy to see that knowing the metric of any exterior link (i.e., $W_{m_1 a_i}$ according to the assumption) will allow unique identification of all the other exterior link metrics, contradicting the fact that \mathbf{R}' is rank-deficient and hence not able to identify all the metrics. Note the direct link $m_1 m_2$ (if exists) is always identifiable. Therefore, none of the exterior links (except $m_1 m_2$) can be identified with two monitors. ■

B. Discussion of Paths With Cycles

At the end of [23, Sec. III-B], the authors raise the question whether or not monitoring nonsimple paths (i.e., paths that may contain cycles) between two monitors suffices to identify all link




Fig. 4. Illustration of Condition ①: bridge l_1 not incident to monitors; (b) bridge l_2 incident to monitor m_1 ; (c) edge cut $\{l_1, l_2\}$ separating \mathcal{G}_2 and monitors; (d) edge cut $\{l_1, l_2\}$ separating m_1 and m_2 .

metrics in the network.⁸ 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 [23] is that monitoring (simple or nonsimple) paths between two monitors is *not* sufficient to identify all link metrics.

V. IDENTIFIABILITY OF INTERIOR LINKS WITH TWO MONITORS

Network administrators are more interested in using end-to-end measurements to infer the qualities of links that are at least *one hop away* as network administrators can employ other methods to estimate metrics of the links that are directly incident to monitors. 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 links in \mathcal{H} using two monitors (m_1 and m_2), under the earlier assumption that no link metrics (including those of links incident to m_1 and m_2) are known ahead of time.

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 subgraphs $\mathcal{G}_i := \mathcal{H}_i + m_1 + m_2$, 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 decoupled. Therefore, in the rest of this section, we assume \mathcal{H} is connected with $\|\mathcal{H}\| \geq 1$. Our result can be applied to each \mathcal{G}_i separately when \mathcal{H} is disconnected.

A. Proof of Theorem III.2: Necessary Part

Suppose all links in \mathcal{H} are identifiable. We prove the necessity of Conditions ① and ② in Theorem III.2 by contradiction.

- Let $l_1 \in L(\mathcal{H})$ be an arbitrary interior link. If $\mathcal{G} - l_1$ is disconnected, then l_1 is a bridge⁹ in \mathcal{G} [shown in Fig. 4(a)]. If \mathcal{G}_1 and \mathcal{G}_2 each contain a monitor, then l_1 is unidentifiable by Lemma A.1 (Appendix). If m_1 and m_2 are both in \mathcal{G}_1 (or \mathcal{G}_2), then l_1 cannot be included in any measurement path (otherwise, l_1 will be used more than once), and is thus unidentifiable. Both cases contradict the assumption that all interior links are identifiable.
- 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 A.1 (Appendix), its adjacent interior links $ra_i \in L(\mathcal{H})$ are

⁸We restrict the paths from containing repeated monitors to exclude measurement of cycles, as is already solved in [23].

⁹A link whose removal will disconnect the graph is a bridge [31].

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) and (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 every $m_1 \rightarrow m_2$ path must traverse either l_1 or l_2 , but not both. Assuming that \mathcal{G}_1 and \mathcal{G}_2 in Fig. 4(d) are connected and all link metrics in them are known, then the resulting measurement matrix \mathbf{R}_1 is similar to (3), except that each row in \mathbf{R}_1 has a new entry associated with W_{l_1} or W_{l_2}

$$\mathbf{R}_1 = \begin{pmatrix} \text{Exterior} & & & \\ \text{Links} & & & \\ \mathbf{R} & | & W_{l_1} & | & W_{l_2} \\ \mathbf{R} & | & \mathbf{e}_1 & | & \mathbf{e}_1 \end{pmatrix}.$$

Here, blank entries correspond to zeroes, \mathbf{R} is given by (3), and \mathbf{e}_1 is a $(k_1 \times k_2)$ -element column vector of all ones. Let \mathbf{R}_2 and \mathbf{R}_3 denote the first and last $k_1 \times k_2$ rows in \mathbf{R}_1 , respectively. Then, applying the same linear transformations as applied to (3) to both \mathbf{R}_2 and \mathbf{R}_3 , we can transform \mathbf{R}_1 into \mathbf{R}'_1 (ignoring rows of zeros)

$$\mathbf{R}'_1 = \begin{pmatrix} \text{Exterior} & & & \\ \text{Links} & & & \\ \mathbf{R}' & | & W_{l_1} & | & W_{l_2} \\ \mathbf{R}' & | & \mathbf{e}_2 & | & \mathbf{e}_2 \end{pmatrix}.$$

where \mathbf{R}' is given by (4) and \mathbf{e}_2 is a $(k_1 + k_2 - 1)$ -element column vector of all ones. In \mathbf{R}'_1 , for each $q = k_1 + k_2, \dots, 2k_1 + 2k_2 - 2$, replace row(q) with row(q) - row($q - k_1 - k_2 + 1$). Ignoring duplicate rows, \mathbf{R}'_1 transforms into

$$\mathbf{R}''_1 = \begin{pmatrix} \text{Exterior} & & & \\ \text{Links} & & & \\ \mathbf{R}' & | & W_{l_1} & | & W_{l_2} \\ \mathbf{R}' & | & \mathbf{e}_2 & | & \mathbf{e}_2 \\ & & -1 & | & 1 \end{pmatrix}.$$

where all rows are linearly independent. Since any subset of K ($k \geq 1$) equations in \mathbf{R}''_1 contains more than k unknown variables, none of these variables can be identified. When \mathcal{G}_1 and \mathcal{G}_2 are not connected, the rank of \mathbf{R}''_1 can be even smaller. Thus, W_{l_1} and W_{l_2} are unidentifiable, contradicting 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 A.2 (Appendix). ■

B. Proof of Theorem III.2: Sufficient Part

Given Conditions ① and ②, we need to show that all links in \mathcal{H} are identifiable. We first introduce two types of identifiable links. The argument then consists of showing that every interior link belongs to one of these two types.




Fig. 5. Two types of identifiable links in \mathcal{H} (a) cross-link y (b) shortcut y .

1) *Cross-Link, Shortcut, and Nonseparating Cycle*: We begin by introducing a few notions essential to our proof.

The first notion is a cross-link that connects nodes on two simple paths between the monitors.

Definition 2: As illustrated 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 from simple paths $\mathcal{P}_1, \dots, \mathcal{P}_4$ by

$$\begin{cases} \mathcal{P}_A = \mathcal{P}_1 \cup \mathcal{P}_2 \\ \mathcal{P}_B = \mathcal{P}_3 \cup \mathcal{P}_4 \end{cases} \quad \begin{cases} \mathcal{P}_C = \mathcal{P}_1 \cup y \cup \mathcal{P}_4 \\ \mathcal{P}_D = \mathcal{P}_3 \cup y \cup \mathcal{P}_2 \end{cases} \quad (5)$$

such that

$$\begin{cases} |\mathcal{P}_1 \cap \mathcal{P}_2| = 1 \\ |\mathcal{P}_3 \cap \mathcal{P}_4| = 1 \end{cases} \quad \begin{cases} |\mathcal{P}_2 \cap \mathcal{P}_3| = 0 \\ |\mathcal{P}_1 \cap \mathcal{P}_4| = 0. \end{cases} \quad (6)$$

See Table I for definitions of graph union/intersection and $|\cdot|$; note that paths are also graphs. The constraints in (6) 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 and \mathcal{P}_2 have no common node other than v_1) ensures that no cycles exist in \mathcal{P}_A . However, this does not require $\mathcal{P}_1 - \mathcal{P}_4$ to be node disjoint, e.g., \mathcal{P}_1 and \mathcal{P}_3 can have common nodes. A cross-link y can then be identified by

$$W_y = \frac{1}{2}(C_{\mathcal{P}_C} + C_{\mathcal{P}_D} - C_{\mathcal{P}_A} - C_{\mathcal{P}_B}). \quad (7)$$

The second notion is a shortcut that connects the endpoints of a simple path whose metric is known.

Definition 3: As illustrated 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 $m_1 \rightarrow m_2$ simple paths can be formed:

$$\mathcal{P}_A = \mathcal{P}_1 \cup y \cup \mathcal{P}_2 \quad \mathcal{P}_B = \mathcal{P}_1 \cup \mathcal{P}_3 \cup \mathcal{P}_2, \quad (8)$$

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 and \mathcal{P}_B are simple paths. A shortcut y can be identified by

$$W_y = C_{\mathcal{P}_A} - C_{\mathcal{P}_B} + C_{\mathcal{P}_3}. \quad (9)$$

The third notion is a special kind of cycle defined as follows.

Definition 4: A *nonseparating cycle* in \mathcal{G} , denoted by \mathcal{F} , is an induced subgraph¹⁰ such that: 1) \mathcal{F} is a cycle (see definition in Table I); and 2) \mathcal{F} does not separate any node from monitors, i.e., each connected component in $\mathcal{G} \setminus \mathcal{F}$ contains at least one monitor.

For example, there are four nonseparating cycles in Fig. 6: $v_1 v_2 v_3 v_1, v_4 v_3 v_2 v_5 v_4, m_1 v_1 v_3 v_4 m_1,$ and $v_5 v_2 m_2 v_5$. Cycle $v_4 v_3 v_1 v_2 v_5 v_4$ is not a nonseparating cycle as it is not induced (due to link $v_2 v_3$), neither is $v_4 m_1 v_1 v_2 v_5 v_4$ as it separates v_3 from monitors.

¹⁰An *induced subgraph* \mathcal{G}' of \mathcal{G} is a subgraph such that for any pair of vertices v and w in \mathcal{G}' , vw is an edge in \mathcal{G}' if and only if vw is an edge in \mathcal{G} .





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.

2) *Proof*: The key to the sufficiency proof is to show that each interior link can be categorized as either a cross-link or a shortcut when the network satisfies Conditions ① and ②. The proof proceeds in three steps.

The *first step* is to show that under Conditions ① and ②, every interior link vw satisfies one of the three cases as shown in Fig. 7. Specifically, we have the following lemma.

Lemma V.1: If graph \mathcal{G} satisfies Conditions ① and ②, then for any interior link vw , there exists a nonseparating cycle \mathcal{F} with $vw \in L(\mathcal{F})$, a cycle \mathcal{C} with $vw \in L(\mathcal{C})$, a simple path \mathcal{P}_1 connecting one monitor with a node on $\mathcal{F} - v - w$, and a simple path \mathcal{P}_2 connecting the other monitor with a node on $\mathcal{C} - v - w$ such that we have the following.

- \mathcal{F} and \mathcal{C} have at most one common node other than v, w (i.e., $|V(\mathcal{F}) \cap V(\mathcal{C})| \leq 3$).
- \mathcal{P}_1 and \mathcal{P}_2 are disjoint, neither going through v nor w (i.e., $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset, v, w \notin V(\mathcal{P}_1)$ and $v, w \notin V(\mathcal{P}_2)$).
- $|V(\mathcal{P}_1) \cap V(\mathcal{F})| = 1, |V(\mathcal{P}_2) \cap V(\mathcal{C})| = 1$.

Proof: See [33]. ■

Lemma V.1 states that there must exist two cycles sharing link vw , a nonseparating cycle \mathcal{F} and a (not necessarily nonseparating) cycle \mathcal{C} that satisfy conditions a)–c) in the lemma. These conditions imply three possible cases.

- Case A* [Fig. 7(a)]: \mathcal{C} is also a nonseparating cycle, \mathcal{F} and \mathcal{C} have no common node other than v and w , and each of \mathcal{F} and \mathcal{C} connects to a different monitor by a disjoint simple path.
- Case B-1* [Fig. 7(b)]: Any path from any node in $\mathcal{F} - v - w$ to monitors must have a common node with $\mathcal{C} - v - w$.
- Case B-2* [Fig. 7(c)]: \mathcal{F} and \mathcal{C} must have at least one common node (denoted by r) other than v and w .

These cases are guaranteed to be complete by Lemma V.1. We will show later that Case-A links are cross-links and Case-B (including B-1 and B-2) links are shortcuts.

The *second step* is to recognize that in Case A [Fig. 7(a)], we can construct four measurement paths $\mathcal{P}_A - \mathcal{P}_D$ by (5), using $\mathcal{P}_1 - \mathcal{P}_4$ marked in the figure. Thus, a Case-A link vw is a cross-link and can be identified as in (7). The challenge is that it is impossible to construct measurement paths as in (5) for the other two cases [Fig. 7(b) and (c)] since the conditions in (6) cannot be satisfied. This motivates us to explore the possibility of identifying W_{vw} in these two cases by (9), i.e., to investigate whether vw in Fig. 7(b) and (c) can be characterized as a shortcut.

The *third step* is to show that any link of Case B-1 or B-2 is indeed a shortcut. We prove this by showing that for any Case-B link, we can find a detour path connecting its endpoints such that all the links in this detour path are cross-links, and thus can be identified by (7). To this end, we show the following lemma.

Lemma V.2: Let $m_1^*, m_2^* \in \{m_1, m_2\}$ with $m_1^* \neq m_2^*$. If \mathcal{G} satisfies Conditions ① and ②, then we have the following.

- For any nonseparating cycle in \mathcal{G} , there is at most *one* Case-B link in this nonseparating cycle.
- For any Case-B link vw in the interior graph of \mathcal{G} , there exists a nonseparating cycle \mathcal{F}_{vw} with $vw \in L(\mathcal{F}_{vw})$ and $m_1, m_2 \notin V(\mathcal{F}_{vw})$. For this nonseparating cycle \mathcal{F}_{vw} , there exist disjoint simple paths $\mathcal{P}(m_1^*, v)$ and $\mathcal{P}(m_2^*, w)$, each intersecting with \mathcal{F}_{vw} only at the endpoint, i.e., $V(\mathcal{P}(m_1^*, v) \cap \mathcal{F}_{vw}) = \{v\}$ and $V(\mathcal{P}(m_2^*, w) \cap \mathcal{F}_{vw}) = \{w\}$.

Proof: See [33]. ■

Lemma V.2(b) implies that a Case-B link y must reside on a nonseparating cycle \mathcal{F} within the interior graph. Meanwhile, Lemma V.2 a) implies that all the other links on cycle \mathcal{F} are cross-links, as there is at most one Case-B link on each nonseparating cycle. Furthermore, we show in Lemma V.2(b) that there exist disjoint simple paths \mathcal{P}_1 and \mathcal{P}_2 connecting the endpoints of link y to different monitors, each sharing only one common node with \mathcal{F} at the endpoints of link y . Together, these allow us to construct two measurement paths \mathcal{P}_A and \mathcal{P}_B as in (8), where $\mathcal{P}_3 = \mathcal{F} - y$ consists of only cross-links whose metrics are known from (7). Thus, a Case-B link y is a shortcut and can be identified by (9).

Therefore, under Conditions ① and ②, every interior link is identifiable. ■

As an example, Fig. 6 displays a network satisfying Conditions ① and ②, where each interior link can be characterized as either a cross-link or a shortcut.

VI. IDENTIFIABILITY WITH THREE OR MORE MONITORS

Since two monitors are not sufficient to identify all link metrics in \mathcal{G} , we explore in this section the case where the network contains three or more monitors.

A. Conversion Into Two-Monitor Problem

Section V suggests that it is easier to identify links that are one hop away from the monitors. This observation motivates 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}) are at least one hop away from 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 two monitors (again we have no prior knowledge of link metrics in \mathcal{G} or \mathcal{G}_{ex}). Therefore, we can apply Theorem III.2 to obtain the following result.

Lemma VI.1: Employing κ 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 attainable by the real monitors m_1, \dots, m_κ is the same as the information attainable by the

virtual monitors m'_1 and m'_2 , if the virtual monitors can make end-to-end measurements along simple paths in \mathcal{G}_{ex} .

First, we show that any measurement between the real monitors can be obtained from measurements between m'_1 and m'_2 . To this end, consider a path $m_i\mathcal{P}m_j$ ($i, j \in \{1, \dots, \kappa\}, i \neq j$) in \mathcal{G} , as shown in Fig. 3(b). Four simple paths between m'_1 and m'_2 can be constructed

$$\begin{cases} \mathcal{P}_A = m'_1 m_i m'_2 \\ \mathcal{P}_B = m'_1 m_j m'_2 \\ \mathcal{P}_C = m'_1 m_i \mathcal{P} m_j m'_2 \\ \mathcal{P}_D = m'_1 m_j \mathcal{P} m_i m'_2. \end{cases} \quad (10)$$

Viewing $m_i\mathcal{P}m_j$ as a ‘‘cross-link,’’ we can compute $W_{m_i\mathcal{P}m_j}$ from the measurements on these four paths via (7) (replacing W_y by $W_{m_i\mathcal{P}m_j}$).

Second, we show that measurements between m'_1 and m'_2 in \mathcal{G}_{ex} do not provide extra information for identifying links in \mathcal{G} compared to measurements attainable by the real monitors. This is proved by observing that for any $m'_1 \rightarrow m'_2$ simple path $m'_1 m_i \mathcal{P} m_j m'_2$ ($i, j \in \{1, \dots, \kappa\}, i \neq j$) containing at least one link in \mathcal{G} , the information relevant for identifying links in \mathcal{G} can be obtained by measuring its $m_i \rightarrow m_j$ subpath $m_i\mathcal{P}m_j$, which must also be a simple path. ■

B. Sketch of Proof for Theorem III.3

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, based on the following arguments.

From the structure of \mathcal{G}_{ex} (see Fig. 3), we prove in Propositions A3 and A4 (Appendix) 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 [31, Proposition 1.4.2], a 3-vertex-connected graph is also 3-edge-connected. Thus, the necessary and sufficient conditions in Lemma VI.1 can be simplified to a single condition that \mathcal{G}_{ex} be 3-vertex-connected. ■

Remark: Note \mathcal{G}_{ex} cannot be 3-vertex-connected if $\kappa < 3$. Therefore, 3 is the absolute lower bound of minimum monitor requirement for ensuring identifiability of an arbitrary nontrivial network topology; however, the minimum number of monitors required for achieving identifiability of a given network is typically greater than 3 (see Section VII-C).

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 for or ensure the identifiability of a given network. In this section, we present efficient algorithms that can: 1) test for the identifiability of a given network with a given monitor placement; and 2) place the minimum number of monitors in a given network to identify all its link metrics.

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 do 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: 1) 2-edge-connected [34], or 2) 3-vertex-connected [35], both in time $O(|V| + |L|)$ ($|V|$: number of nodes; $|L|$: number of links). Using these algorithms, we can test for the identifiability of the *interior graph* of \mathcal{G} as follows.

- 1) For each interior link l , apply the 2-edge-connectivity test in [34] to $\mathcal{G} - l$. \mathcal{G} is unidentifiable if the test fails.
- 2) Apply the 3-vertex-connectivity test in [35] to $\mathcal{G} + m_1 m_2$. \mathcal{G} is unidentifiable if the test fails.

The interior graph of \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 for 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 [35] to \mathcal{G}_{ex} . \mathcal{G} is identifiable if the test succeeds, and unidentifiable otherwise.

The complexity of this algorithm is $O(|V(\mathcal{G}_{\text{ex}})| + |L(\mathcal{G}_{\text{ex}})|)$, which is the same as $O(|V(\mathcal{G})| + |L(\mathcal{G})|)$.

B. Optimal Monitor Placement

The next question we want to answer is the following: Given an arbitrary topology \mathcal{G} , what is the minimum number of monitors needed and where should they be placed to identify all link metrics in \mathcal{G} ? While one can enumerate all possible placements for $\kappa = 3, 4, \dots$ monitors and test for identifiability until an identifiable placement is found, what we really want is an efficient algorithm to achieve the same outcome. This requires a deeper understanding of the structure of an identifiable graph.

We illustrate our idea by an example in Fig. 8(a). We first consider the minimum deployment: Using fewer monitors certainly renders the network unidentifiable. This yields the following rules.

- 1) *Dangling* node (e.g., h) must be a monitor, as otherwise its adjacent link cannot be measured using simple paths.
- 2) A node on a *tandem* of links (e.g., j) must be a monitor, as otherwise we can only identify the sum of its adjacent link metrics (W_{dj} and W_{jk}) and not the individual metrics.
- 3) For a subgraph with *two cut-vertices*¹¹ (e.g., \mathcal{G}_1) or a *2-vertex cut*¹² (e.g., \mathcal{G}_2), at least one node other than those cuts must be a monitor, as otherwise even if all links outside this subgraph have been identified, the vertices in the cuts (w and d for \mathcal{G}_1 , f and i for \mathcal{G}_2) are effectively the two ‘‘monitors’’ for this subgraph, and Theorem III.1 states that this subgraph cannot be identified.
- 4) Similarly, for a subgraph with *one cut-vertex* (e.g., \mathcal{G}_3), at least two nodes other than the cut-vertex must be monitors.

Our strategy is to use the above four rules to deploy the necessary monitors. If we can prove that these necessary monitors are also sufficient to identify all links, then our deployment is optimal. To formally present the algorithm, we introduce the following definitions.

¹¹A *cut-vertex* is a vertex whose removal will disconnect the graph.

¹²A *2-vertex cut* is a set of two vertices $\{v_1, v_2\}$ such that removing v_1 or v_2 alone does not disconnect \mathcal{G} , but removing both disconnects \mathcal{G} .

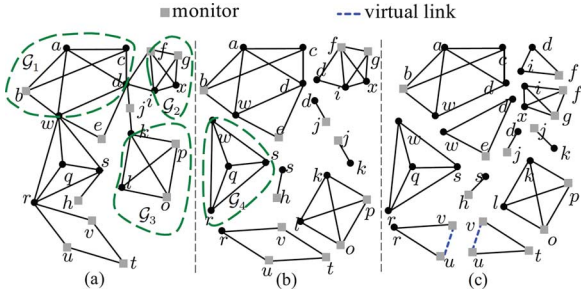


Fig. 8. (a) Original graph. (b) Biconnected components. (c) Triconnected components.

Definition 5: A k -connected component of \mathcal{G} is a maximal subgraph of \mathcal{G} that is either: 1) k -vertex-connected, or 2) a complete graph with up to k vertices. The case of $k = 2$ is also called a *biconnected component*, and $k = 3$ a *triconnected component*.

Intuitively, a biconnected component is a subgraph connected to the rest of the graph by cut-vertices, and a triconnected component (within a biconnected component) is a subgraph connected to the rest by 2-vertex cuts. For instance, Fig. 8(b) shows the biconnected components of Fig. 8(a), separated by cut-vertices d, j, k, w, s , and r . Fig. 8(c) shows the triconnected components, separated by the above cut-vertices and 2-vertex cuts $\{w, d\}$, $\{f, i\}$, and $\{u, v\}$. To identify the triconnected components, we process the graph by adding *virtual links* as follows: If \exists a *minimal* 2-vertex cut¹³ whose vertices are not neighbors (e.g., $\{u, v\}$), connect them by a virtual link; repeat this on the resulting graph until no such cut exists. In the sequel, nodes that are cut-vertices or part of 2-vertex cuts are called *separation vertices* (e.g., $w, d, f, i, j, k, s, r, u$, and v).

Monitor Placement Algorithm: There exist fast algorithms to partition an arbitrary graph \mathcal{G} into biconnected components [36] and then into triconnected components [35]. Note that the output of [35] is the set of triconnected components of \mathcal{G} augmented with the virtual links.¹⁴ Based on these algorithms, we propose a master algorithm, *Minimum Monitor Placement (MMP)*, to place the minimum number of monitors needed to identify \mathcal{G} . As shown in Algorithm 1, MMP first applies rules 1) and 2) to select all the dangling vertices and vertices on tandems as monitors (line 1), and then applies rules 3) and 4) to select additional monitors in each¹⁵ triconnected/biconnected component. For a component \mathcal{D} , let $s_{\mathcal{D}}$ denote the number of separation vertices, $c_{\mathcal{D}}$ the number of cut-vertices, and $M_{\mathcal{D}}$ the number of (already selected) monitors in \mathcal{D} . MMP goes through each triconnected and then biconnected component that contains three or more nodes to ensure that: 1) each triconnected component has at least three nodes that are either separation vertices or monitors (lines 6–8); and 2) each biconnected component has at least three nodes that are either cut-vertices or monitors (lines 10–12). Finally, it selects additional monitors as needed to ensure that the total number of monitors is at least three (lines 14–16).

Optimality: It is easy to see from rules 1)–4) that MMP only deploys monitors when needed, and thus no algorithm can achieve identifiability with fewer monitors; on the other hand,

¹³That is, neither of the vertices are cut-vertices.

¹⁴The algorithm in [35] does not require a separate step to add virtual links; it adds virtual links as needed in the process of graph partitioning.

¹⁵MMP does not depend on the order of biconnected/triconnected components being considered for monitor selection.

Algorithm 1: Minimum Monitor Placement (MMP)

input: Connected graph \mathcal{G}
output: A subset of nodes in \mathcal{G} as monitors

- 1 choose all the nodes with degree less than 3 as monitors;
- 2 partition \mathcal{G} into biconnected components $\mathcal{B}_1, \mathcal{B}_2, \dots$;
- 3 **foreach** biconnected component \mathcal{B}_i with $|\mathcal{B}_i| \geq 3$ **do**
- 4 partition \mathcal{B}_i into triconnected components $\mathcal{T}_1, \mathcal{T}_2, \dots$;
- 5 **foreach** triconnected component \mathcal{T}_j of \mathcal{B}_i with $|\mathcal{T}_j| \geq 3$ **do**
- 6 **if** $0 < s_{\mathcal{T}_j} < 3$ **and** $s_{\mathcal{T}_j} + M_{\mathcal{T}_j} < 3$ **then**
- 7 randomly choose $3 - s_{\mathcal{T}_j} - M_{\mathcal{T}_j}$ nodes in \mathcal{T}_j that are neither separation vertices nor monitors as monitors;
- 8 **end**
- 9 **end**
- 10 **if** $0 < c_{\mathcal{B}_i} < 3$ **and** $c_{\mathcal{B}_i} + M_{\mathcal{B}_i} < 3$ **then**
- 11 randomly choose $3 - c_{\mathcal{B}_i} - M_{\mathcal{B}_i}$ nodes in \mathcal{B}_i that are neither cut-vertices nor monitors as monitors;
- 12 **end**
- 13 **end**
- 14 **if** the total number of monitors $\kappa < 3$ **then**
- 15 randomly choose $3 - \kappa$ nonmonitor nodes as monitors;
- 16 **end**

we show that the monitor placement by MMP is also sufficient, i.e., all link metrics can be identified from end-to-end measurements between the selected monitors. Thus, MMP is optimal as it places the *minimum* number of monitors to identify all link metrics in \mathcal{G} , as stated in the following theorem.

Theorem VII.1: For an arbitrary connected network \mathcal{G} , Algorithm 1 (MMP) generates the *optimal* monitor placement in the sense that: 1) all link metrics in \mathcal{G} are identifiable under this placement; and 2) no placement can identify all link metrics in \mathcal{G} with a smaller number of monitors.

Proof: See the Appendix. ■

Complexity: In Algorithm 1, lines 1 and 14–16 take $O(|V(\mathcal{G})|)$ time. Splitting \mathcal{G} into biconnected (line 2) and then triconnected components (line 4) takes $O(V(\mathcal{G}) + L(\mathcal{G}))$ time [35], [36]. Selecting monitors takes $O(1)$ time per component, and the counters ($s_{\mathcal{T}_j}, M_{\mathcal{T}_j}, c_{\mathcal{B}_i}$, and $M_{\mathcal{B}_i}$) can be computed during the splitting/selecting process. Therefore, the entire algorithm has $O(|V(\mathcal{G})| + |L(\mathcal{G})|)$ time complexity.

Example: Given the graph in Fig. 8(a) as input (total of 22 nodes), MMP selects 11 monitors, where nodes $\{h, u, v, t, j, e\}$ are selected by lines 6–8, $\{b, g, o, p\}$ by lines 10–12, and f by lines 10–12. It can be verified that the resulting graph satisfies the identifiability condition in Theorem III.3.

C. Evaluation of MMP

We evaluate MMP through a set of simulations on both randomly generated and real network topologies. We use the following algorithm, referred to as *Random Monitor Placement (RMP)*, as a benchmark for comparison¹⁶: Given network \mathcal{G} , randomly select κ ($\kappa = 2, \dots, |V|$) nodes as monitors and test the identifiability of the resulting network using the algorithm in

¹⁶To our knowledge, MMP is the first monitor placement algorithm for identifying additive link metrics by measuring controllable, cycle-free paths. Thus, we use random placement to represent the average performance of an arbitrary monitor placement for comparison.

Section VII-A. Generally, RMP cannot guarantee network identifiability for arbitrary \mathcal{G} and κ . Therefore, we measure its performance by the fraction of random placements achieving network identifiability over multiple Monte Carlo runs.

1) *Random Topologies*: We first consider synthetic topologies generated according to four widely used random graph models: Erdős–Rényi (ER) graphs, Random Geometric (RG) graphs, Barabási–Albert (BA) graphs, and Random Power Law (PL) graphs. We randomly generate 100 graph realizations of each model,¹⁷ with each realization containing 150 nodes (i.e., $|V| = 150$). The generated graphs are then input to the monitor placement algorithms. We now explain the models and the corresponding results separately.

ER Graph: The ER graph is a simple random graph generated by independently connecting each pair of nodes by a link with a fixed probability p . The result is a purely random topology where all graphs with an equal number of links are equally likely to be selected. It is known [37] that $p_0 = \log |V|/|V|$ is a sharp threshold for the graph to be connected with high probability, which implies a minimum value of $p = 0.0334$ for $|V| = 150$.

RG Graph: The RG graph is frequently used to model the topology of wireless ad hoc networks. It generates a random graph by first randomly distributing nodes in a unit square, and then connecting each pair of nodes by a link if their distance is no larger than a threshold d_c , which denotes node communication range. The resulting topology contains well-connected subgraphs in densely populated areas and poorly connected subgraphs in sparsely populated areas. It is known that $d_c \geq \sqrt{\log |V|/(\pi|V|)}$ ensures a connected graph with high probability [38], which gives a minimum range of $d_c = 0.1031$ for $|V| = 150$.

BA Graphs: The BA model [39] provides a random power-law graph generated by the following preferential attachment mechanism. We begin with a small connected graph $\mathcal{G}_0 := (\{v_1, v_2, v_3, v_4\}, \{v_1v_2, v_1v_3, v_1v_4\})$ and add nodes sequentially. For each new node v , we connect v to n_{\min} existing nodes, where n_{\min} specifies (a lower bound on) the minimum node degree, such that the probability of connecting to node w is proportional to the degree of w . If the number of existing nodes is smaller than n_{\min} , then v connects to all existing nodes. The BA graph has been used to model many naturally occurring networks, e.g., Internet, citation networks, and social networks.

PL Graphs: The BA model introduces an artifact that all node degrees are lower-bounded by n_{\min} . Alternatively, the PL graph [40] provides another way of generating power-law graphs by directly specifying a sequence of expected node degrees $(d_1, \dots, d_{|V|})$ according to the power law, i.e., $d_i = i^\alpha$ ($\alpha > 0$). The generation of a PL graph is similar to that of an ER graph, except that instead of connecting each pair of nodes with the same probability, nodes i and j in a PL graph are connected by a link with probability $p_{ij} = d_i d_j / \sum_{k=1}^{|V|} d_k$.

We conduct simulations in two scenarios, densely connected graphs (Fig. 9 with parameter configurations: $p = 0.039$ for ER, $d_c = 0.11943$ for RG, $n_{\min} = 3$ for BA, and $\alpha = 0.42$ for PL) and sparsely connected graphs (Fig. 10 with parameter configurations: $p = 0.0253$ for ER, $n_{\min} = 2$ for BA, and

¹⁷All realizations are guaranteed to be connected, as we discard disconnected realizations in the generation process.




Fig. 9. Comparison between RMP and MMP: densely connected graphs ($\bar{n} = 441$, $\bar{\kappa}_{\text{MMP}} = 3$ for BA, $\bar{n} = 437$, $\bar{\kappa}_{\text{MMP}} = 9.36$ for ER, $\bar{n} = 437$, $\bar{\kappa}_{\text{MMP}} = 19.42$ for PL, and $\bar{n} = 451$, $\bar{\kappa}_{\text{MMP}} = 14.52$ for RG).

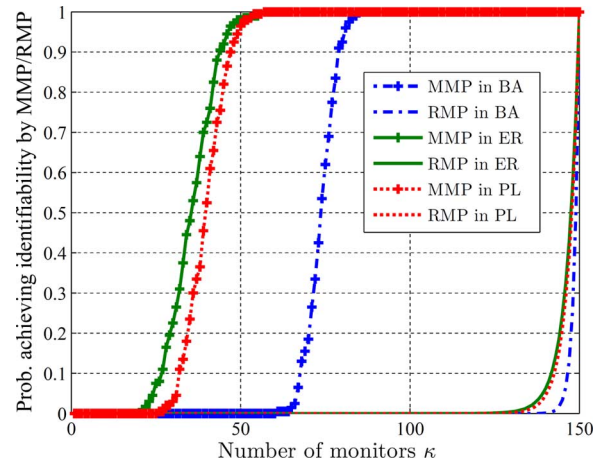


Fig. 10. Comparison between RMP and MMP: sparsely connected graphs ($\bar{n} = 295$, $\bar{\kappa}_{\text{MMP}} = 73.51$ for BA, $\bar{n} = 293$, $\bar{\kappa}_{\text{MMP}} = 36.76$ for ER, and $\bar{n} = 297$, $\bar{\kappa}_{\text{MMP}} = 40.24$ for PL).

$\alpha = 0.32$ for PL), to evaluate how graph density affects the performance of MMP and RMP. Since the number of links n and the minimum number of monitors κ_{MMP} (computed by MMP) vary across graph realizations, we present the average values denoted by \bar{n} and $\bar{\kappa}_{\text{MMP}}$ for each graph model, shown in the captions of Figs. 9 and 10. In each scenario, we have tuned parameters of each model to generate roughly the same average number of links. As RMP is a randomized algorithm, we repeat it for 2000 Monte Carlo runs to obtain the average performance for each graph realization. We then average the results over the 100 graph realizations to obtain the final results,¹⁸ shown in Figs. 9 and 10.

For densely connected graphs (Fig. 9), the probability that RMP is able to identify all the links increases with the number of monitors κ . However, fewer than 20% of the ER, RG, and PL graphs are identifiable when $\kappa \leq 120$ (recall $|V| = 150$), whereas a careful monitor deployment in the same graphs by

¹⁸For each graph realization, MMP achieves identifiability with probability one for $\kappa \geq \kappa_{\text{MMP}}$ and zero for $\kappa < \kappa_{\text{MMP}}$. Therefore, the overall probability for MMP to achieve identifiability using κ monitors is computed as the fraction of graph realizations with $\kappa_{\text{MMP}} \leq \kappa$.

MMP ensures identifiability with significantly fewer monitors ($\kappa < 25$). Therefore, in most cases, the proposed algorithm (MMP) substantially outperforms a randomized scheme (RMP) in the required number of monitors. One exception is the BA graphs, where RMP achieves similar performance to that of MMP. This can be explained as follows: When $n_{\min} = 3$, further simulations show that 87.8% of the generated BA graphs are 3-vertex-connected as each node (except the initial nodes v_2, v_3 , and v_4) has at least three neighbors; according to MMP, this implies that an arbitrary placement of three monitors will achieve identifiability, and there is no need for a sophisticated placement algorithm. Meanwhile, a comparison of the BA and the PL models implies that the exceptional performance of RMP in the BA model is due to the requirement that all node degrees be lower-bounded by three rather than the power-law distribution of node degrees since it does not achieve the same superior performance on PL graphs. For ER, RG, and PL graphs, it is easy to form sparsely connected subgraphs, thus requiring more monitors in these sparsely connected subgraphs to guarantee network identifiability. Among the four models, we observe that RMP exhibits the worst performance in the model with the largest $\bar{\kappa}_{\text{MMP}}$. This is because large $\bar{\kappa}_{\text{MMP}}$ is resulted from poor connectivity within certain subgraphs, where a considerable number of nodes have to be selected as monitors to achieve identifiability. Thus, RMP fails to achieve identifiability if it does not select a sufficient number of nodes in any of these subgraphs.

We perform similar simulations for sparsely connected graphs by adjusting the parameters,¹⁹ shown in Fig. 10. In comparison to Fig. 9, both MMP and RMP perform worse in sparsely connected graphs, requiring a larger number of monitors to achieve the same probability of identifiability. This is because sparser graphs impose more constraints on candidate measurement paths between pairs of monitors, and thus more monitors are needed to identify all links. As we compare the performance of the same graph models in Figs. 9 and 10, we observe that all models exhibit smooth transitions when we vary the average number of links \bar{n} except for the BA model, which has extremely good performance for $n_{\min} = 3$ but poor performance (except for RG) for $n_{\min} = 2$. This can again be explained by examining the 3-vertex-connectivity of the generated graph realizations. We have verified that as we reduce n_{\min} from 3 to 2, the probability for BA graphs to be 3-vertex-connected drops from 87.8% to 0%, making it unlikely for an arbitrary placement to achieve identifiability. In fact, when $n_{\min} = 2$, on average 49.2% of the nodes in BA graphs have degrees less than 3, which requires them to be selected as monitors by rules 1) and 2) in MMP. Therefore, RMP fails to achieve identifiability if it misses one of these nodes, resulting in its poor performance.

2) *Autonomous System Topologies*: We now test MMP and RMP on real network topologies. We use the *autonomous system* (AS) topologies from both the Rocketfuel [29] and the CAIDA [30] projects, which represent IP-level connections between backbone/gateway routers of several ASs from major *Internet service providers* (ISPs) around the globe. The

¹⁹To generate roughly the same number of links as in the other models, d_c for the RG graph becomes much smaller than $\sqrt{\log |V| / (\pi |V|)}$, making the probability of generating a connected RG graph very small. Thus, the RG model is omitted from Fig. 10.

TABLE II
PARAMETERS OF AS TOPOLOGIES IN ROCKETFUEL

AS	ISP Name	$ L $	$ V $	κ_{MMP}	r_{MMP}
6461	Abovenet (US)	294	182	117	0.64
1755	Ebone (Europe)	381	172	55	0.32
3257	Tiscali (Europe)	404	240	138	0.58
3967	Exodus (US)	434	201	85	0.42
1221	Telstra (Australia)	758	318	164	0.52
7018	AT&T (US)	2078	631	208	0.33
1239	Sprintlink (US)	2268	604	163	0.27
2914	Verio (US)	2821	960	408	0.43
3356	Level3 (US)	5298	624	94	0.15

TABLE III
PARAMETERS OF AS TOPOLOGIES IN CAIDA

AS	$ L $	$ V $	κ_{MMP}	r_{MMP}
15706	874	325	276	0.84
9167	1590	769	483	0.62
8717	3755	1778	1266	0.71
4761	3760	969	624	0.64
20965	8283	968	110	0.11

parameters of selected networks obtained from these two projects are listed in Tables II and III, where we sort the networks according to their numbers of links, and $r_{\text{MMP}} := \kappa_{\text{MMP}}/|V|$ denotes the minimum fraction of monitors computed by MMP in a network with $|V|$ nodes.

a) *AS Topologies From Rocketfuel*: As shown in Table II, each AS in Rocketfuel corresponds to an ISP. To identify the entire network, we observe that most ISPs need a significant fraction of nodes to be monitors, ranging from around 30% (Ebone, AT&T, Sprintlink) to more than 60% (Abovenet). This is because ISP networks contain a large number of gateway²⁰ routers to connect to customer networks or other ISPs, which appear as dangling nodes that have to be selected as monitors [see rule 1) in Section VII-B]. We repeat RMP for $15 \cdot |V|$ Monte Carlo runs for each ISP to evaluate its average performance, measured by the fraction of Monte Carlo runs achieving identifiability, as shown in Fig. 11. To facilitate comparison, we mark the fraction of monitors r_{MMP} needed by MMP in the legends of the same plot; note that MMP guarantees identifiability for each network so long as $\kappa/|V| \geq r_{\text{MMP}}$. We normalize the number of monitors by the total number of nodes in order to compare networks of different sizes. Moreover, we only plot the results of RMP for $\kappa/|V| \geq 95\%$ since RMP fails to achieve identifiability in almost all of the simulations when $\kappa/|V| < 95\%$.

As in the case of synthetic graphs, we again observe a significant improvement of MMP over RMP. Specifically, RMP has at most 50% probability of identifying all the links even if 99% of the nodes are monitors, whereas MMP guarantees identifiability using at most 64% of nodes as monitors. The poor performance of RMP is due to the heterogeneous connectivity

²⁰In real networks, monitor selection may be constrained to a subset of nodes, e.g., gateways. Under such a constraint, the issue of achievable number of identifiable links (i.e., partial network identifiability) is left for future work.




Fig. 11. Comparison between RMP and MMP: ISP topologies in Rocketfuel ($15 \cdot |V|$ Monte Carlo runs, $r_{MMP} := \kappa_{MMP}/|V|$).




Fig. 12. Comparison between RMP and MMP: AS topologies in CAIDA (3000 Monte Carlo runs, $r_{MMP} := \kappa_{MMP}/|V|$).

within ISP networks, which contain poorly connected subnetworks that need a large fraction of monitors, a requirement unlikely to be fulfilled by random placement. We also observe that the relative performance of MMP and RMP varies for different networks, e.g., Level3 and Exodus experience similar performance under RMP, whereas their minimum fractions of monitors computed by MMP differ significantly (0.15 for Level3 and 0.42 for Exodus). Intuitively, this is because RMP performance is determined by the relative number of valid placements, i.e., the fraction of all the $\binom{|V|}{k}$ candidate placements that achieve identifiability, whereas MMP performance is determined by the smallest κ for which this fraction is nonzero.

b) AS Topologies From CAIDA: ISP topologies have evolved since the Rocketfuel project; consequently, we repeat the above evaluation on a recent data set obtained by the CAIDA project; see results in Table III and Fig. 12. Compared to the ASes in Rocketfuel data set, we observe that ASes with similar average node degrees (i.e., $2|L|/|V|$) in the CAIDA data set require a larger fraction of monitors for complete identification, e.g., $r_{MMP} = 0.71$ for AS8717 with average node degree 4.2 in Table III, whereas $r_{MMP} = 0.32$ for AS1755 (Ebone) with average node degree 4.4 in Table II.




Fig. 13. Two cases of bridge link l : (a) interior bridge; (b) exterior bridge.

This is because ASes in the CAIDA data set tend to be more skewed in connectivity, with more densely connected cores and a larger number of dangling nodes (likely the gateways for peer/customer connections). As all dangling nodes have to be selected as monitors, such skewed connectivity leads to a higher ratio of monitors, although the average node degree is not necessarily smaller. Meanwhile, comparison to RMP (3000 Monte Carlo runs), presented in Fig. 12, shows that the improvement of MMP over RMP remains significant. As in Fig. 11, RMP again exhibits poor performance for all the ASes considered in Fig. 12, which has less than 35% probability of identifying all links even if the fraction of monitors $\kappa/|V|$ is as large as 0.99. In particular, for almost all the networks (except for AS20965) in Fig. 12, RMP fails to identify all the links in more than 60% of the simulations even if *all but one* node are monitors. In contrast, carefully selected monitors by MMP can guarantee complete identification while substantially reducing the required number of monitors.

VIII. CONCLUSION

We study the fundamental conditions on network topology and placement of monitors for identifying additive link metrics using end-to-end measurements along simple paths between monitors. We show that with two monitors, it is impossible to identify all the link metrics, but is possible to identify the metrics of interior links that are at least one hop away from the monitors, for which we derive the necessary and sufficient conditions in terms of edge/vertex connectivity of the network topology. We further study the case of three or more monitors and derive the corresponding necessary and sufficient conditions for identifying all the link metrics. We show that these conditions are not only useful for testing network identifiability under a given monitor placement, but also enable an efficient monitor placement algorithm that guarantees identifiability using the minimum number of monitors. Our evaluations on both random and real network topologies verify that the proposed algorithm significantly outperforms a baseline solution.

APPENDIX

Lemma A.1:

Suppose two monitors are deployed in \mathcal{G} to measure simple paths. If link l is a bridge in \mathcal{G} with one monitor on each side, as illustrated in Fig. 13, then neither l nor its adjacent links are identifiable.

Proof: Consider the case of Fig. 13(a). Assume all links except link l and its adjacent links $a_i r$ ($i = 1, \dots, k_1$) and sb_j ($j = 1, \dots, k_2$) are identified. We can then reduce the linear equation associated with any $m_1 \rightarrow m_2$ path to the form

$$W_{a_i r} + W_{sb_j} + W_l = \phi_{ij}, \quad i = 1, \dots, k_1, j = 1, \dots, k_2. \quad (11)$$

Writing these equations in matrix form and applying the linear transform in the proof of Theorem III.1 yield a transformed measurement matrix (blank entries are zero)

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

This matrix corresponds to the maximum set of linearly independent equations involving the unknown variables W_l , $(W_{a_i r})_{i=1}^{k_1}$, and $(W_{sb_j})_{j=1}^{k_2}$. Since any subset of k equations contains more than k unknown variables, none of these variables can be identified. Therefore, l and its adjacent links are all unidentifiable.

In the case of Fig. 13(b), similar argument applies, except that $W_{a_i r}$ is replaced by W_{ra_i} and W_{sb_j} is replaced by $W_{b_j m_2}$. ■

Proposition A.2: Using two monitors measuring simple paths, if all link metrics in the interior graph of \mathcal{G} are identifiable, then $\mathcal{G} + m_1 m_2$ is 3-vertex-connected.

Proof: Suppose that all the interior links of \mathcal{G} are identifiable, and $\mathcal{G} + m_1 m_2$ is not 3-vertex-connected. Then, the connectivity²¹ of $\mathcal{G} + m_1 m_2$ must be 1 or 2.

- 1) Suppose that the connectivity of $\mathcal{G} + m_1 m_2$ is 1. Then, it must have a cut-vertex, denoted by r . There are three possible cases as illustrated in Fig. 14. First, if $r \in \{m_1, m_2\}$, as shown in Fig. 14(a), then removing r does not disconnect $\mathcal{G} + m_1 m_2$ since \mathcal{H} has been assumed to be connected. Thus, r must be in \mathcal{H} . Let $\mathcal{G}_1 - r$ and $\mathcal{G}_2 - r$ denote two of the connected components separated by r . If each of \mathcal{G}_1 and \mathcal{G}_2 contains a monitor, as shown in Fig. 14(b), then removing r does not disconnect $\mathcal{G} + m_1 m_2$ because link $m_1 m_2$ still connects \mathcal{G}_1 and \mathcal{G}_2 ; if one of the components, say $\mathcal{G}_2 - r$, has no monitor, as shown in Fig. 14(c), then any $m_1 \rightarrow m_2$ path employing links in \mathcal{G}_2 must both enter and leave \mathcal{G}_2 through r , forming a cycle (which is forbidden). Therefore, the connectivity of $\mathcal{G} + m_1 m_2$ must be greater than 1.
- 2) Suppose that the connectivity of $\mathcal{G} + m_1 m_2$ is 2. Thus, there must be a 2-vertex cut, denoted by $\{r, s\}$. There are three possibilities. First, if $\{r, s\} = \{m_1, m_2\}$ [Fig. 14(a)], then removing $\{r, s\}$ does not disconnect $\mathcal{G} + m_1 m_2$ because the remaining graph \mathcal{H} is still connected. Second, if \exists one monitor in $\{r, s\}$, say $s = m_2$ as shown in Fig. 15(a),

²¹The greatest integer k such that \mathcal{G} is k -vertex-connected is the *connectivity* of \mathcal{G} .




Fig. 14. Possible scenarios of a cut-vertex r : (a) r is a monitor; (b) the removal of r separates m_1 and m_2 ; (c) \mathcal{G}_2 contains no monitors when r is deleted.




Fig. 15. Possible scenarios of a 2-vertex cut $\{r, s\}$: (a) $\{r, s\}$ contains a monitor; (b) $\{r, s\}$ contains no monitors.

then any $m_1 \rightarrow m_2$ path employing links in \mathcal{G}_2 must enter \mathcal{G}_2 through r and exit through m_2 . We can effectively view r and m_2 as the new “monitors” and their adjacent links $\{ra_i\}_{i=1}^{k_1}$ and $\{m_2 b_j\}_{j=1}^{k_2}$ as the new “exterior links” for \mathcal{G}_2 . Applying Corollary IV.1 yields that $\{ra_i\}_{i=1}^{k_1}$ and $\{m_2 b_j\}_{j=1}^{k_2}$ are all unidentifiable, contradicting the assumption that all the interior links are identifiable (because $\{ra_i\}_{i=1}^{k_1}$ are interior links). Finally, if both r and s are in \mathcal{H} , as shown in Fig. 15(b), then any $m_1 \rightarrow m_2$ path employing links in \mathcal{G}_2 must enter/exit \mathcal{G}_2 via r and s . Note that there must be a component separated by $\{r, s\}$ that has no monitor (e.g., \mathcal{G}_2) because otherwise link $m_1 m_2$ will keep the components connected after removing r and s . For links in \mathcal{G}_2 , r and s are effectively the new “monitors,” and thus by Corollary IV.1, the links $\{ra_i\}_{i=1}^{k_1}$ and $\{sb_j\}_{j=1}^{k_2}$ are unidentifiable. This also contradicts the assumption that all the interior links are identifiable.

Thus, the connectivity of $\mathcal{G} + m_1 m_2$ must be greater than 2, i.e., $\mathcal{G} + m_1 m_2$ is 3-vertex-connected. ■

Proposition A.3: Given a graph \mathcal{G} employing κ ($\kappa \geq 3$) monitors, the extended graph \mathcal{G}_{ex} of \mathcal{G} satisfies Conditions ① (i.e., $\mathcal{G}_{\text{ex}} - l$ is 2-edge-connected for each link l in \mathcal{G}) if and only if \mathcal{G}_{ex} is 3-edge-connected.

Proof: Necessary Part: Suppose $\mathcal{G}_{\text{ex}} - l$ is 2-edge-connected for all l in \mathcal{G} . Consider removing two links in \mathcal{G}_{ex} , denoted by l_1 and l_2 .

- 1) If at least one of these links, say l_1 , is in $L(\mathcal{G})$, then by assumption $\mathcal{G}_{\text{ex}} - l_1$ is 2-edge-connected. Thus, $\mathcal{G}_{\text{ex}} - l_1 - l_2$ is connected.
- 2) Suppose none of these links is in \mathcal{G} , i.e., both l_1 and l_2 are virtual links. Since the virtual monitors m'_1 and m'_2 each connect to all actual monitors in \mathcal{G} , and there are at least three actual monitors, m'_1 and m'_2 are each connected to \mathcal{G} via at least three virtual links. Therefore, m'_1 and m'_2 are still connected with \mathcal{G} after l_1 and l_2 are deleted. Since we




Fig. 16. $\mathcal{G}_{\text{ex}} - v_1 - v_2$ is disconnected, where $v_1, v_2 \in V(\mathcal{G})$: (a) $\{v_1, v_2\}$ separates \mathcal{G}_2 and monitors; (b) $\{v_1, v_2\}$ separates monitors.

have assumed \mathcal{G} to be a connected graph, $\mathcal{G}_{\text{ex}} - l_1 - l_2$ is connected.

We have shown that \mathcal{G}_{ex} remains connected after removing any two links. Therefore, \mathcal{G}_{ex} is 3-edge-connected when $\mathcal{G}_{\text{ex}} - l$ ($l \in L(\mathcal{G})$) is 2-edge-connected.

Sufficient Part: Suppose \mathcal{G}_{ex} is 3-edge-connected. Then, obviously, $\mathcal{G}_{\text{ex}} - l$ is 2-edge-connected for each $l \in L(\mathcal{G})$. ■

Proposition A.4: Given a graph \mathcal{G} employing κ ($\kappa \geq 3$) monitors, the extended graph \mathcal{G}_{ex} of \mathcal{G} satisfies Condition ② (i.e., $\mathcal{G}_{\text{ex}} + m'_1 m'_2$ is 3-vertex-connected) if and only if \mathcal{G}_{ex} is 3-vertex-connected.

Proof: Necessary Part: We prove the necessary part by contradiction. Suppose \mathcal{G}_{ex} is not 3-vertex-connected, but $\mathcal{G}_{\text{ex}} + m'_1 m'_2$ is 3-vertex-connected, then the connectivity of \mathcal{G}_{ex} must be 2, because removing one link will decrease connectivity by at most 1. Thus, there must exist two nodes, denoted by v_1 and v_2 , whose removal will disconnect \mathcal{G}_{ex} . There are three possibilities for v_1 and v_2 .

- 1) If v_1, v_2 are m'_1, m'_2 , then after their removal, the remaining graph (\mathcal{G}) is still connected.
- 2) If v_1 is a virtual monitor (m'_1 or m'_2) and v_2 is a node in \mathcal{G} , then $\mathcal{G}_{\text{ex}} - v_1 - v_2$ being disconnected will imply $\mathcal{G}_{\text{ex}} + m'_1 m'_2 - v_1 - v_2$ also being disconnected (as the remaining graphs of \mathcal{G}_{ex} and $\mathcal{G}_{\text{ex}} + m'_1 m'_2$ are the same), contradicting the assumption that $\mathcal{G}_{\text{ex}} + m'_1 m'_2$ is 3-vertex-connected.
- 3) If v_1, v_2 are both in \mathcal{G} (can be real monitors), then two cases may occur after removing v_1 and v_2 : a) \exists a connected component that does not contain any real monitor; b) each connected component contains at least one real monitor, as illustrated in Fig. 16. In the case of Fig. 16(a), $\mathcal{G}_{\text{ex}} + m'_1 m'_2 - v_1 - v_2$ is disconnected as well, contradicting the 3-vertex-connectivity of $\mathcal{G}_{\text{ex}} + m'_1 m'_2$. In the case of Fig. 16(b), different components in $\mathcal{G}_{\text{ex}} - v_1 - v_2$ can still connect via virtual links and virtual monitors, thus contradicting the assumption that $\mathcal{G}_{\text{ex}} - v_1 - v_2$ is disconnected.

Hence, when $\mathcal{G}_{\text{ex}} + m'_1 m'_2$ is 3-vertex-connected, the connectivity of \mathcal{G}_{ex} cannot be less than 3, i.e., \mathcal{G}_{ex} is also 3-vertex-connected.

Sufficient Part: If \mathcal{G}_{ex} is 3-vertex-connected, then after adding one link $m'_1 m'_2$, $\mathcal{G}_{\text{ex}} + m'_1 m'_2$ is also 3-vertex-connected. ■

Proof of Theorem VII.1: If \mathcal{G} is already 3-vertex-connected, then \mathcal{G}_{ex} is always 3-vertex-connected as long as there are at least three monitors, no matter how they are placed. This case is handled by lines 14–16 in Algorithm 1. We will show that




Fig. 17. Possible scenarios for a triconnected component \mathcal{T} with $s_{\mathcal{T}} = 1$.

even if \mathcal{G} is not 3-vertex-connected, Algorithm 11 still guarantees that \mathcal{G}_{ex} is 3-vertex-connected. We prove this statement by showing that after removing any two nodes u_1, u_2 in \mathcal{G} , each remaining node in $\mathcal{G} - u_1 - u_2$ is connected to at least one monitor. There are three possible cases: 1) u_1, u_2 belong to the same triconnected component; 2) u_1, u_2 belong to different triconnected components within the same biconnected component; 3) u_1, u_2 belong to different biconnected components. We now analyze these cases separately.

- 1) Consider deleting two nodes u_1, u_2 in a triconnected component \mathcal{T} . If $|\mathcal{T}| = 2$, then \mathcal{T} is a bridge. According to rules 1)–4) (Section VII-B), each of the neighboring components of \mathcal{T} must contain at least one monitor other than u_1 and u_2 . Thus, after deleting u_1 and u_2 , each connected component must contain at least one monitor.

Now consider the case of $|\mathcal{T}| \geq 3$. Since $\mathcal{T} - u_1 - u_2$ must be connected (due to its triconnectivity), it suffices to show that $\mathcal{T} - u_1 - u_2$ contains or is connected to at least one monitor. There are four possibilities depending on the number of separation vertices in \mathcal{T} , denoted by $s_{\mathcal{T}}$.

- a) If $s_{\mathcal{T}} = 0$, then \mathcal{G} must be 3-vertex-connected or a triangle. After removing two vertices, the remaining graph is either connected with at least one monitor or a degenerate graph of a single node (which is a monitor).
- b) If $s_{\mathcal{T}} = 1$, then there exists one cut-vertex v_c in \mathcal{T} (shown²² in Fig. 17). If $v_c \notin \{u_1, u_2\}$, then $\mathcal{T} - u_1 - u_2$ is still connected to monitors in neighboring components via v_c ; if $v_c \in \{u_1, u_2\}$, then $\mathcal{T} - u_1 - u_2$ must contain a monitor because two of the nonseparation vertices in \mathcal{T} must be monitors by rule 4).
- c) If $s_{\mathcal{T}} = 2$, then there are two possible scenarios (shown in Fig. 18), i.e., a 2-vertex-cut $\{v_1, v_2\}$, or \mathcal{T} has two cut-vertices v_1 and v_2 . If $\{v_1, v_2\} = \{u_1, u_2\}$, then $\mathcal{T} - u_1 - u_2$ must contain at least one monitor because \mathcal{T} contains at least one monitor that is not a separation vertex [by rule 3)]. If $v_2 \in \{u_1, u_2\}$ and $v_1 \notin \{u_1, u_2\}$ (or the other way), then $\mathcal{T} - u_1 - u_2$ is still connected to monitors in neighboring components via v_1 . Similarly, if $v_1, v_2 \notin \{u_1, u_2\}$, then $\mathcal{T} - u_1 - u_2$ is still connected to monitors in neighboring components via v_1 and v_2 .
- d) If $s_{\mathcal{T}} \geq 3$ (three possible scenarios are shown in Fig. 19), then Algorithm 1 will not place any monitor in \mathcal{T} . Nevertheless, after removing two nodes u_1 and u_2 , at least one separation vertex must remain, and thus $\mathcal{T} - u_1 - u_2$ is still connected to monitors in

²²Note that there may be more than one triconnected subgraph connecting at v_c , but we only illustrate the minimum for clarity; the same applies to subsequent illustrations in this proof.




Fig. 18. Possible scenarios for a triconnected component \mathcal{T} with $s_{\mathcal{T}} = 2$: (a) \mathcal{T} contains a 2-vertex cut $\{v_1, v_2\}$; (b) \mathcal{T} contains two cut-vertices v_1 and v_2 .

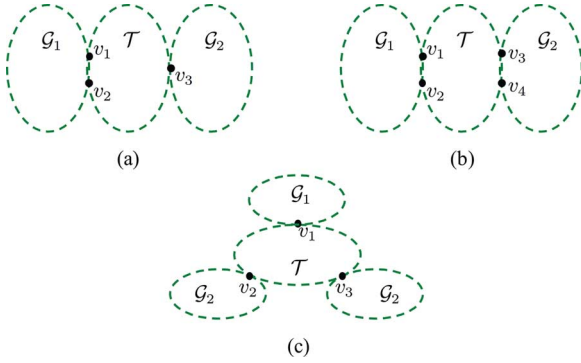


Fig. 19. Possible scenarios for a triconnected component \mathcal{T} with $s_{\mathcal{T}} \geq 3$: (a) \mathcal{T} contains a 2-vertex cut $\{v_1, v_2\}$ and a cut-vertex v_3 ; (b) \mathcal{T} contains two 2-vertex cuts $\{v_1, v_2\}$ and $\{v_3, v_4\}$; (c) \mathcal{T} contains three cut-vertices v_1, v_2 , and v_3 .

at least one neighboring component via the remaining separation vertex.

The above shows that remaining nodes in \mathcal{T} are still connected to monitors after removing two nodes u_1 and u_2 within the same triconnected component \mathcal{T} .

- 2) Now consider deleting two nodes u_1, u_2 in two different triconnected components $\mathcal{T}_i, \mathcal{T}_j$ ($i \neq j$) that are in the same biconnected component \mathcal{B} . Note that the remaining graph $\mathcal{B}' = \mathcal{B} - u_1 - u_2$ must be connected because the only scenario for \mathcal{B}' to be disconnected is when u_1, u_2 belong to the same 2-vertex cut of \mathcal{B} , which implies that u_1, u_2 must belong to the same triconnected component, contradicting our assumption. Similar to Case 1), we have the following four possibilities depending on the number of cut-vertices in \mathcal{B} , denoted by $c_{\mathcal{B}}$.

- If $c_{\mathcal{B}} = 0$, then \mathcal{G} is biconnected ($\mathcal{B} = \mathcal{G}$), and $\mathcal{G} - u_1 - u_2$ must contain at least one monitor because the total number of monitors is at least three (see lines 14–16).
- If $c_{\mathcal{B}} = 1$, then rule 4) implies that $\mathcal{B}(|\mathcal{B}| \geq 3)$ contains two monitors other than the cut-vertex v_c . If $v_c \notin \{u_1, u_2\}$, then $\mathcal{B} - u_1 - u_2$ is still connected to monitors in neighboring components via v_c ; if $v_c \in \{u_1, u_2\}$, then $\mathcal{B} - u_1 - u_2$ itself contains a monitor.
- If $c_{\mathcal{B}} = 2$, then \mathcal{B} connects to neighboring biconnected components via two cut-vertices v_1 and v_2 . If $v_1, v_2 \in \{u_1, u_2\}$, then $\mathcal{B} - u_1 - u_2$ itself contains a monitor, as rule 3) requires \mathcal{B} to contain a monitor that is not a cut-vertex; if $v_1 \notin \{u_1, u_2\}$ (or $v_2 \notin \{u_1, u_2\}$), then $\mathcal{B} - u_1 - u_2$ is connected to monitors in at least one neighboring component via v_1 (or v_2).

- If $c_{\mathcal{B}} \geq 3$, then \mathcal{B} connects to neighboring biconnected components via at least three cut-vertices v_1, v_2 , and v_3 . There may not be any monitor in \mathcal{B} in this case. However, $\mathcal{B} - u_1 - u_2$ must be connected to monitors in at least one neighboring component via the remaining cut-vertex.

The above shows that nodes in $\mathcal{B} - u_1 - u_2$ are still connected to monitors after removing nodes u_1 and u_2 in two different triconnected components but within the same biconnected component \mathcal{B} .

Although our argument has focused on the (triconnected or biconnected) component that contains both u_1 and u_2 , it is easy to see that nodes in the neighboring components that share u_1 or u_2 are also connected to monitors after these two nodes are removed, as this is a special case of the above argument where the number of removed nodes is at most two, and the removed nodes have to be separation vertices. *Remark:* The above argument assumes that virtual links have been added (see explanations after Definition 5). However, the construction of virtual links implies that if a node in $\mathcal{B} - u_1 - u_2$ will be disconnected from monitors without virtual links, then this node must have only two neighbors, and line 1 in Algorithm 1 would have selected this node as a monitor. Thus, our argument holds without added virtual links.

- Finally, consider deleting two nodes u_1 and u_2 from two different biconnected components \mathcal{B}_1 and \mathcal{B}_2 , respectively. Then one of the following scenarios will occur.
 - If neither u_1 nor u_2 is a cut-vertex,²³ then $\mathcal{G} - u_1 - u_2$ is still connected. Moreover, it contains at least one monitor as there are at least three monitors in \mathcal{G} .
 - If u_1 is a cut-vertex but u_2 is not (or the other way), then we leverage a result from Cases 1) and 2): Removing two nodes from a biconnected component (whether in the same triconnected component or not), all remaining nodes are still connected to monitors. Applying this result, we see that removing one node from a biconnected component also does not disconnect any remaining node from monitors, i.e., each connected component in $\mathcal{G} - u_1$ has at least one monitor. Consider the connected component \mathcal{G}' containing u_2 . Removing u_2 cannot disconnect \mathcal{G}' , as otherwise u_2 will have to be in the same biconnected component as u_1 or a cut-vertex itself, contradicting our assumptions; furthermore, \mathcal{G}' must contain at least two monitors by rule 4). Thus, each connected component in $\mathcal{G} - u_1 - u_2$, including $\mathcal{G} - u_2'$, has at least one monitor.
 - If both u_1 and u_2 are cut-vertices, then each connected component in $\mathcal{G} - u_1 - u_2$ must connect to the rest of the graph in the original \mathcal{G} through one or two cut-vertices (i.e., u_1 and u_2). By rules 3) and 4), each of these connected components contains at least one monitor other than u_1 or u_2 .

In summary, the arguments in Cases 1)–3) imply that, for any \mathcal{G} (which may not be 3-vertex-connected), 1) guarantees that all nodes in $\mathcal{G} - u_1 - u_2$ are still connected to monitors after deleting two arbitrary nodes u_1, u_2 in \mathcal{G} . Since all the monitors are connected through virtual links and virtual monitors in the

²³ u_1 and u_2 cannot be a 2-vertex cut because otherwise u_1 and u_2 are in the same triconnected component, contradicting our assumption.

extended graph \mathcal{G}_{ex} , $\mathcal{G}_{\text{ex}} - u_1 - u_2$ must be connected. Moreover, if the virtual monitors m'_1, m'_2 are deleted from \mathcal{G}_{ex} , then the remaining graph (i.e., \mathcal{G}) is still connected since \mathcal{G} is assumed to be connected. Finally, if m'_1 (or m'_2) and $u \in V(\mathcal{G})$ are deleted, then we have shown that all the nodes in $\mathcal{G} - u$ are connected to monitors, and since these monitors are still connected via m'_2 , the entire graph $\mathcal{G}_{\text{ex}} - m'_1 - u$ is connected. Therefore, employing Algorithm 1 to place monitors, the resulting extended graph \mathcal{G}_{ex} is always 3-vertex-connected. By Theorem III.3, this implies that the placed monitors can uniquely identify all the link metrics in \mathcal{G} . ■

REFERENCES

- [1] F. L. Presti, N. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Trans. Netw.*, vol. 10, no. 6, pp. 761–775, Dec. 2002.
- [2] A. B. Downey, "Using pathchar to estimate Internet link characteristics," in *Proc. ACM SIGCOMM*, 1999, pp. 241–250.
- [3] G. Jin, G. Yang, B. Crowley, and D. Agarwal, "Network characterization service (NCS)," in *Proc. IEEE HPDC*, 2001, pp. 289–299.
- [4] M. Coates, A. O. Hero, R. Nowak, and B. Yu, "Internet tomography," *IEEE Signal Process. Mag.*, vol. 19, no. 3, pp. 47–65, May 2002.
- [5] E. Lawrence and G. Michailidis, "Network tomography: A review and recent developments," in *Proc. Frontiers Statist.*, 2006, vol. 54, pp. 345–364.
- [6] O. Gurewitz and M. Sidi, "Estimating one-way delays from cyclic-path delay measurements," in *Proc. IEEE INFOCOM*, 2001, pp. 1038–1044.
- [7] Y. Chen, D. Bindel, and R. H. Katz, "An algebraic approach to practical and scalable overlay network monitoring," in *Proc. ACM SIGCOMM*, 2004, pp. 55–66.
- [8] A. Chen, J. Cao, and T. Bu, "Network tomography: Identifiability and Fourier domain estimation," in *Proc. IEEE INFOCOM*, 2007, pp. 1875–1883.
- [9] Information Sciences Institute, University of Southern California, Marina del Rey, CA, USA, "Internet protocol, DARPA Internet program protocol specification," RFC: 791, 1981 [Online]. Available: <http://www.ietf.org/rfc/rfc0791.txt>
- [10] Open Networking Foundation, Palo Alto, CA, USA, "Open Networking Foundation," [Online]. Available: <http://www.opennetworkingfoundation.org>
- [11] S. Jain et al., "B4: Experience with a globally-deployed software defined WAN," in *Proc. ACM SIGCOMM*, 2013, pp. 3–14.
- [12] D. Mcdysan, "Software defined networking opportunities for transport," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 28–31, Mar. 2013.
- [13] N. Duffield and F. Lo Presti, "Multicast inference of packet delay variance at interior network links," in *Proc. IEEE INFOCOM*, 2000, pp. 1351–1360.
- [14] Y. Xia and D. Tse, "Inference of link delay in communication networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 12, pp. 2235–2248, Dec. 2006.
- [15] A. Adams et al., "The use of end-to-end multicast measurements for characterizing internal network behavior," *IEEE Commun. Mag.*, vol. 38, no. 5, pp. 152–159, May 2000.
- [16] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu, "Network tomography: Recent developments," *Statist. Sci.*, vol. 19, no. 3, pp. 499–517, 2004.
- [17] M.-F. Shih and A. Hero, "Unicast inference of network link delay distributions from edge measurements," in *Proc. IEEE ICASSP*, 2001, pp. 3421–3424.
- [18] T. Bu, N. Duffield, and F. Lo Presti, "Network tomography on general topologies," in *Proc. ACM SIGMETRICS*, 2002, pp. 21–30.
- [19] N. Duffield, "Simple network performance tomography," in *Proc. ACM/USENIX IMC*, 2003, pp. 210–215.
- [20] M. Firooz and S. Roy, "Network tomography via compressed sensing," in *Proc. IEEE GLOBECOM*, 2010, pp. 1–5.
- [21] W. Xu, E. Mallada, and A. Tang, "Compressive sensing over graphs," in *Proc. IEEE INFOCOM*, 2011, pp. 2087–2095.
- [22] S. Ahuja, S. Ramasubramanian, and M. Krunz, "SRLG failure localization in all-optical networks using monitoring cycles and paths," in *Proc. IEEE INFOCOM*, 2008, pp. 700–708.
- [23] A. Gopalan and S. Ramasubramanian, "On identifying additive link metrics using linearly independent cycles and paths," *IEEE/ACM Trans. Netw.*, vol. 20, no. 3, pp. 906–916, Jun. 2012.
- [24] N. Alon, Y. Emek, M. Feldman, and M. Tennenholtz, "Economical graph discovery," in *Proc. Symp. Innov. Comput. Sci.*, 2011, pp. 476–486.
- [25] L. Ma, T. He, K. Leung, D. Towsley, and A. Swami, "Efficient identification of additive link metrics via network tomography," in *Proc. IEEE ICDCS*, 2013, pp. 581–590.
- [26] Y. Bejerano and R. Rastogi, "Robust monitoring of link delays and faults in IP networks," in *Proc. IEEE INFOCOM*, 2003, pp. 134–144.
- [27] R. Kumar and J. Kaur, "Practical beacon placement for link monitoring using network tomography," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 12, pp. 2196–2209, Dec. 2006.
- [28] J. D. Horton and A. Lopez-Ortiz, "On the number of distributed measurement points for network tomography," in *Proc. ACM/USENIX IMC*, 2003, pp. 204–209.
- [29] University of Washington, Seattle, WA, USA, "Rocketfuel: An ISP topology mapping engine," 2002 [Online]. Available: <http://www.cs.washington.edu/research/networking/rocketfuel/>
- [30] The Cooperative Assoc. for Internet Data Anal. (CAIDA), La Jolla, CA, USA, "Macroscopic Internet Topology Data Kit (ITDK)," Apr. 2013 [Online]. Available: <http://www.caida.org/data/active/internet-topology-data-kit/>
- [31] R. Diestel, *Graph Theory*. New York, NY, USA: Springer-Verlag, 2005.
- [32] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley, "Monitor placement for maximal identifiability in network tomography," in *Proc. IEEE INFOCOM*, 2014, to be published.
- [33] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley, "Identification of additive link metrics: Proof of selected theorems," Imperial College, London, U.K., Tech. Rep., 2012 [Online]. Available: <http://www.commsp.ee.ic.ac.uk/%7elm110/pdf/MaTechreportJul12b.pdf>
- [34] R. E. Tarjan, "A note on finding the bridges of a graph," *Inf. Process. Lett.*, vol. 2, pp. 160–161, 1974.
- [35] J. E. Hopcroft and R. E. Tarjan, "Dividing a graph into triconnected components," *SIAM J. Comput.*, vol. 2, pp. 135–158, 1973.
- [36] R. Tarjan, "Depth-first search and linear graph algorithms," *SIAM J. Comput.*, vol. 1, pp. 146–160, 1972.
- [37] P. Erdős and A. Rényi, "On the evolution of random graphs," *Pub. Math. Inst. Hungarian Acad. Sci.*, vol. 5, pp. 17–61, 1960.
- [38] P. Gupta and P. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications*. New York, NY, USA: Springer, 1999, pp. 547–566.
- [39] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Rev. Modern Phys.*, vol. 74, pp. 47–97, Jan. 2002.
- [40] F. Chung and L. Lu, *Complex Graphs and Networks*. Providence, RI, USA: Amer. Math. Soc., 2006.



Liang Ma (S'09) received the B.Sc. and M.Sc. degrees in electrical engineering with distinction from the Beijing University of Posts and Telecommunications (BUPT), Beijing, China, in 2007 and 2010, respectively, and is currently pursuing the Ph.D. degree in electrical and electronic engineering at Imperial College London, London, U.K.

Before joining Imperial College, he worked as a Research Intern with NTT DoCoMo Beijing Labs, Ericsson (China) Communications, and Microsoft Research Asia, all in Beijing, China, where he was involved in WLAN medium access control, high-speed switching system, and software radio-based gigabit multiantenna communications, respectively. Since 2011, he has been working with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA; Army Research Lab, Adelphi, MD, USA; and the University of Massachusetts, Amherst, MA, USA, on network measurement and algorithm design.

Mr. Ma has served as a peer reviewer in a range of journals and conferences, including *IEEE Network*, *Journal of Wireless Networks*, *INFOCOM*, *WCNC*, etc. He was the recipient of the International Conference on Distributed Computing System (ICDCS 2013) Best Paper Award, IBM Invention 1st File Regular Patent Award, Chatschik Bisdikian Memorial Best Student Paper Award of ITA in Network & Information Sciences 2013, ACM SIGCOMM Internet Measurement Conference (IMC 2013) Best Paper Award Nomination, IEEE INFOCOM 2014 Student Travel Grant, and the winner of the Outstanding Graduate Student 2008 and Excellent Student Awards four times during 2003–2006.



Ting He (S'04–M'07–SM'13) received the B.S. degree in computer science from Peking University, Beijing, China, in 2003, and the Ph.D. degree in electrical and computer engineering from Cornell University, Ithaca, NY, USA, in 2007.

She has worked as an Undergraduate Research Assistant with the Micro Processor Research and Development Center, Peking University, from 2001 to 2003. From 2003 to 2007, she was a Graduate Research Assistant with the Adaptive Communications and Signal Processing Group (ACSP), Cornell University, under the supervision of Prof. L. Tong. In 2007, she joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, where she is currently a Research Staff Member with the Network Analytics Research Group. At IBM, she has worked as a primary researcher and task lead in several research programs including the International Technology Alliance (ITA) program funded by the U.S. ARL and U.K. MoD, the ARRA program funded by NIST, the Social Media in Strategic Communication (SMISC) program funded by DARPA, etc. Her work is in the broad areas of networking, statistical inference, and information theory.

Dr. He has served as the Membership Co-Chair of ACM N2Women and the GHC PhD Forum committee. She has served on the TPC of a range of communications and networking conferences, including IEEE INFOCOM, IEEE SECON, IEEE/ACM IWQoS, IEEE MILCOM, IEEE ICNC, IFIP Networking, etc. She received the Best Paper Award at the 2013 International Conference on Distributed Computing Systems (ICDCS), a Best Paper Nomination at the 2013 Internet Measurement Conference (IMC), and the Best Student Paper Award at the 2005 International Conference on Acoustic, Speech and Signal Processing (ICASSP). Her student received the Best Student Paper Award at the 2013 ITA Fall Annual Meeting (previously Annual Conference of ITA). In school, she was an Outstanding College Graduate of Beijing Area and an Outstanding Graduate of Peking University in 2003, a winner of the Excellent Student Award of Peking University from 1999 to 2002, and a recipient of the Canon, Sony, and Yang-Wang Academician scholarships from 1999 to 2002.



Kin K. Leung (S'78–M'86–SM'93–F'01) received the B.S. degree in electronics from the Chinese University of Hong Kong, Hong Kong, in 1980, and the M.S. and Ph.D. degrees in computer science from University of California, Los Angeles, CA, USA, in 1982 and 1985, respectively.

He joined AT&T Bell Labs, Holmdel, NJ, USA, in 1986 and worked at its successor companies, AT&T Labs and Bell Labs of Lucent Technologies, until 2004. Since then, he has been the Tanaka Chair Professor in Internet Technology with Imperial

College London, London, U.K. He serves as the head of the Communications and Signal Processing Group and as the Deputy Director for the University Defense Research Center in Signal Processing, Electrical Engineering Department, Imperial College. His research interests include networking, protocols, optimization, and modeling issues for wireless broadband, sensor, and ad hoc networks.

Prof. Leung has actively served on many conference committees. He was a member of the IEEE Fellow Evaluation Committee for Communications Society (2009–). He was a Guest Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, *IEEE Wireless Communications*, and *Mobile Networks and Applications*, and was an Editor for the IEEE JOURNAL

ON SELECTED AREAS IN COMMUNICATIONS Wireless Series and the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS. Currently, he is an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the *International Journal on Sensor Networks*, and *ACM Computing Surveys*. He received the Distinguished Member of Technical Staff Award from AT&T Bell Labs in 1994 and was a co-recipient of the 1997 Lanchester Prize Honorable Mention Award. He received the Royal Society Wolfson Research Merits Award from 2004 to 2009.



Ananthram Swami (S'79–M'84–SM'96–F'08) received the B.Tech. degree from the Indian Institute of Technology (IIT), Bombay; India, the M.S. degree from Rice University, Houston, TX, USA, and the Ph.D. degree from the University of Southern California (USC), Los Angeles, CA, USA, all in electrical engineering.

He has held positions with Unocal Corporation, USC, CS-3 and Malgudi Systems. He was a Statistical Consultant to the California Lottery, developed HOSAT, a MATLAB-based toolbox for non-Gaussian signal processing, and has held visiting faculty positions with INP, Toulouse, France. He is the ST for Network Science with the U.S. Army Research Laboratory, Adelphi, MD, USA. He is co-editor of the book *Wireless Sensor Networks: Signal Processing & Communications Perspectives* (Wiley, 2007). His work is in the broad areas of signal processing, wireless communications, and sensor and mobile ad hoc networks.

Dr. Swami is a member of the IEEE Signal Processing Society's (SPS) Technical Committee on Sensor Array and Multichannel Systems and serves on the Senior Editorial Board of the IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING. He was a tutorial speaker on Networking Cognitive Radios for Dynamic Spectrum Access at IEEE ICC 2010, co-chair of IEEE SPAWC'10, and has served on the IEEE SPS BoG.



Don Towsley (M'78–SM'93–F'95) received the B.A. degree in physics and Ph.D. degree in computer science from the University of Texas at Austin, Austin, TX, USA, in 1971 and 1975, respectively.

He is currently a Distinguished Professor with the Department of Computer Science, University of Massachusetts, Amherst, MA, USA. He has held visiting positions at numerous universities and research labs. His research interests include networks and performance evaluation.

Prof. Towsley is a Fellow of the Association for Computing Machinery (ACM) and a member of ORSA. He currently serves on the Editorial Board of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS and previously served as Editor-in-Chief of the IEEE/ACM TRANSACTIONS ON NETWORKING and on numerous other editorial boards. He has served as a Program Co-Chair of several conferences including IEEE INFOCOM 2009. He has received several achievement awards including the 2007 IEEE Koji Kobayashi Award and the 2011 INFOCOM Achievement Award. He has received numerous paper awards including a 2008 SIGCOMM Test-of-Time Paper Award and the 1998 IEEE Communications Society William Bennett Best Paper Award.