Distributed Link Anomaly Detection via Partial Network Tomography

Ting He

Pennsylvania State University, University Park, PA, USA. tzh58@psu.edu [*]

ABSTRACT

We consider the problem of detecting link loss anomalies from end-to-end measurements using network tomography. Network tomography provides an alternative to traditional means of network monitoring by inferring link-level performance characteristics from end-to-end measurements. Existing network tomography solutions, however, insist on characterizing the performance of all the links, which introduces unnecessary delays for anomaly detection due to the need of collecting all the measurements at a central location. We address this problem by developing a distributed detection scheme that integrates detection into the measurement fusion process by testing anomalies at the level of *minimal identifiable link sequences (MILSs)*. We develop efficient methods to configure the proposed detection scheme such that its false alarm probability satisfies a given bound. Meanwhile, we provide analytical bounds on the detection probability and the detection delay. We then extend our solution to further improve the detection performance by designing the probing and fusion process. Our evaluations on real topologies verify that the proposed scheme significantly outperforms both centralized detection based on link parameters inferred by traditional network tomography and distributed detection based on raw end-to-end measurements.

Keywords

Anomaly detection; network tomography; distributed detection; minimal identifiable link sequence.

1. INTRODUCTION

Reliable monitoring of internal network performance (e.g., link delays and loss rates) is crucial to effective network management. The traditional approach of obtaining such information relies on directly collecting performance statistics at internal nodes (routers/switches) using their monitoring functionalities, which has a critical drawback that it requires both *global access* to the internal nodes and *global support* of the monitoring functionalities. This drawback has limited its applicability in networks that are either administratively or functionally heterogeneous (e.g., the Internet, coalition networks, and hybrid networks). In such networks, *network tomography* [33] provides an attractive alternative that infers link-level performance from end-to-end performance measured between special nodes known as *monitors*. Compared with the traditional approach, network tomography has the advantage that it only requires *local access and support* at the endpoints of measurement paths (i.e., monitors). Moreover, whenever possible, network tomography can reduce the measurement overhead by leveraging passive measurements from data traffic on the measurement paths.

Network tomography has received considerable attention in the past decade [9, 5]. Existing works have focused on the estimation problem, which aims at characterizing the performance of individual links. In practice, however, the network administrator may only be interested in detecting the presence of abnormal links. This is the case, for example, when a client monitors its service provider (or a service provider monitors its peers) for *service level agreement (SLA)* violations. Although one can detect link anomalies from link-level performance characteristics, such detailed information is usually unnecessary, as finding anomaly on one path is sufficient for detecting anomaly, while estimating link performance characteristics typically requires measurements on a much larger set of paths. This difference makes it possible to perform anomaly detection more efficiently than solving a typical network tomography problem.

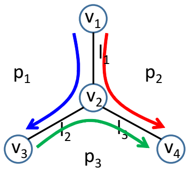

Figure 1: Success probabilities on links l_1, l_2, l_3 are 0.85, 1, 1; per-link threshold is 0.9.

We consider link loss anomaly as a concrete example. Given end-to-end binary measurements (losses/successes) on a set of paths diverse enough to identify all the link success probabilities, we want to detect the presence of abnormal links whose success probabilities are below a given threshold. This problem has broad applications in detecting congestions, transient routing loops, and traffic throttling. Although one can perform anomaly detection on end-to-end measurements, such detection will have limited "resolution" in the sense that it can only detect path anomalies, and

[*] This work was built on research sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-06-3-0001.

© 2017 ACM. ISBN 123-4567-24-567/08/06.

DOI: 10.475/123_4

is ineffective if an abnormal link does not cause any path anomaly after being mixed with normal links. For example, although link l_1 in Fig. 1 is abnormal, none of the paths has an abnormal success probability with respect to the threshold of $0.9^2 = 0.81$. Using network tomography, however, we can detect the anomaly by inferring the link success probabilities from the path success probabilities estimated from sufficiently many measurements.

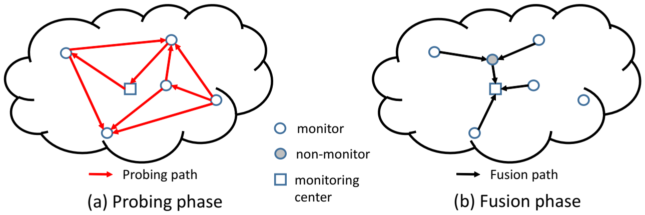

(a) Probing phase

(b) Fusion phase

Figure 2: Process of link anomaly detection.

Inferring the link success probabilities requires all the end-to-end measurements to be available at a single location. As illustrated in Fig. 2, we divide the entire detection process into two phases: (a) *probing phase*, where monitors send probes along probing paths to measure their end-to-end success probabilities, and (b) *fusion phase*, where the (statistics of) end-to-end measurements are fused at a central location, referred to as the *monitoring center*, for processing. Given measurements generated by the probing phase, our goal is to detect link anomalies with the best accuracy and efficiency. We measure the accuracy by the *detection probability*, defined as the probability of detecting anomaly when there is at least one abnormal link, and the *false alarm probability*, defined as the probability of detecting anomaly when there is no abnormal link. We measure the efficiency by the *detection delay*, defined as the time between the end of probing and the instant a decision (anomaly/no anomaly) is made. At one extreme, path-based detection can be performed on raw measurements by the monitors serving as probing destinations, with a small detection delay but a low accuracy (see the example in Fig. 1); at the other extreme, link-based detection can be performed on inferred link success probabilities by the monitoring center after fusing all the measurements, with a high accuracy but a large detection delay. In this work, we aim to combine the strengths of the two via a *distributed detection scheme*, which strives to detect anomaly as soon as possible based on partial views of link performance obtained by *partial network tomography*.

1.1 Related Work

First introduced by Vardi [33], network tomography is a family of network monitoring techniques that infer network characteristics from indirect measurements [5], with three canonical applications: *network performance tomography*, *network topology tomography*, and *traffic matrix tomography* [9]. What we leverage in this paper is network performance tomography, which aims at inferring link performance characteristics from end-to-end measurements.

Early works on network tomography focus on best-effort solutions that aim at extracting the most information from given measurements through estimator design; see [9, 5] and references therein. Due to evidence that end-to-end measurements are frequently insufficient for identifying all the link characteristics [17, 28, 7, 36, 6], later works turn the focus to the design of measurements that guarantee identifiability [15, 21, 22, 23, 1, 2, 8, 25, 24]. For stochastic link metrics, further studies have been performed to minimize estimation error by optimizing the allocation of probes among probing paths/trees according to the principles of optimal experiment design [34, 16, 18]. Existing works, however, have only focused on the estimation problem. To our knowledge, this is the first work that rigorously studies the detection problem in the context of network tomography.

An existing problem closely related to our problem of link anomaly detection is *Boolean network tomography*, where links are associated with binary states (e.g., normal/failed), and the goal is to infer link states from binary end-to-end measurements. Early works assume that at most a single link may fail [3, 19]. In practice, however, multiple links can fail simultaneously [26]. To handle the uncertainty in failure locations, heuristic or Bayesian approaches have been proposed to find the most likely set of failed links [10, 11, 20, 35, 27]. Alternatively, designed measurements can be used to guarantee unique failure localization [1, 2, 8, 25, 24]. Our problem also associates a binary state (normal/abnormal) with each link; however, we are only interested in detecting the presence of at least one abnormal link rather than inferring the states of all the links.

1.2 Summary of Contributions

Our main contributions are five-fold:

1) We develop a distributed detection scheme that integrates anomaly detection into the measurement fusion process by testing anomalies at the level of *minimal identifiable link sequences (MILSs)* (see Section 4.1).

2) We show that in contrast to path-based detection which always has a *uniformly most powerful (UMP)* detector, MILS-based detection generally does not have a UMP detector. Thus, we propose a heuristic detector based on the *maximum likelihood estimator (MLE)* of MILS success probability.

3) We develop efficient methods to configure the proposed detection scheme to satisfy any given false alarm bound. We also give analytical bounds on the detection probability and the expected detection delay.

4) We further study the design of the probing and fusion process. Specifically, we give a polynomial-time algorithm to jointly select the probing paths, their orientations, and the monitoring center to identify all the links while minimizing the maximum fusion path length.

5) We evaluate the proposed solution via extensive simulations based on real topologies. Our results show that under the same false alarm constraint, the proposed detection scheme achieves a much higher detection probability than distributed detection based on raw measurements and a much smaller detection delay than centralized detection based on inferred link parameters, and the proposed measurement design further enhances these improvements.

The rest of the paper is organized as follows. Section 2 formulates the problem. Section 3 reviews preliminaries and two baseline solutions. Section 4 presents the proposed solution and its performance analysis. Section 5 discusses further optimization. Section 6 evaluates the proposed solution against the baselines. Section 7 discusses how to apply our solution when not all the links are identifiable. Section 8 concludes the paper. *All the proofs are in the appendix.*

2. PROBLEM FORMULATION

2.1 Network Model

We model the network as an undirected graph $\mathcal{G} = (N, L)$, where N is the set of nodes and L the set of links. Each link $l \in L$ is associated with an unknown parameter $\theta_l \in (0,\, 1]$, which denotes the success probability (complement of loss probability) for transmissions over l. Let $\boldsymbol{\theta} := (\theta_l)_{l\in L}$. We assume that losses are i.i.d. for different transmissions on the same link and independent across links.

A subset of nodes $Z \subseteq N$ are *monitors*. Monitors actively participate in probing, fusion, and decision making, while non-monitors only forward packets according to the underlying routing mechanism. The monitors measure the success probabilities of a set of paths P by sending probes, where each $p \in P$ is a path starting/ending at monitors that conforms to the underlying routing mechanism. Suppose that n probes are sent on each path $p \in P$, giving measurements $\mathbf{x}_p := (x_{p,1}, \ldots, x_{p,n})$, where $x_{p,i} \in \{0, 1\}$ indicates whether the i-th probe successfully traverses p. Let $\mathbf{x} := (\mathbf{x}_p)_{p\in P}$. Under the independent loss assumption, the measurements follow a conditional distribution

$$f(\mathbf{x}_p;\, \boldsymbol{\theta}) = \alpha_p^{\sum_{i=1}^n x_{p,i}} (1-\alpha_p)^{n-\sum_{i=1}^n x_{p,i}}, \qquad (1)$$

where $\alpha_p := \prod_{l\in p} \theta_l$ is the success probability of path p.

After probing, the monitors receiving probes report the measurements to a *monitoring center* r along fusion paths. Without loss of generality, we assume $r \in Z$. Note that a probing path is used to send probes, and a fusion path is used to send measurements (obtained from probes). We denote the union of the fusion paths by a tree $\mathcal{T} = (N_T,\, L_T)$ rooted at r ($N_T \subseteq N$, $L_T \subseteq L$).

We consider in-band fusion. Assuming that each transmission takes one slot, the per-hop fusion delay d_v to successfully report a measurement (statistic) from node v in $\mathcal{T}$ to its parent along link l_v is geometrically distributed with

$$\Pr\{d_v = k\} = (1-\theta_{l_v})^{k-1} \theta_{l_v}, \quad k \geq 1. \qquad (2)$$

As is clear later, r only needs to fuse the empirical path success probability $\hat{\alpha}_p := \frac{1}{n}\sum_{i=1}^n x_{p,i}$ for each $p \in P$. Let $d_{v,p}$ denote the time after probing that $\hat{\alpha}_p$ is reported to node v, i.e., sum of the per-hop fusion delays from the probing destination of p to v, measured in slots ($d_{v,p} := \infty$ if v is not on the fusion path of $\hat{\alpha}_p$).

2.2 Network Tomography

In our context, network tomography aims at estimating link success probabilities $\boldsymbol{\theta}$ from measurements of path successes $\mathbf{x}$. To achieve consistent estimation, network tomography typically assumes that the probing paths span the link space. Specifically, define a *measurement matrix* $\mathbf{A} = (A_{p,l})_{p\in P,\, l\in L}$, where $A_{p,l}$ denotes the number of times path p traverses link l. The above assumption means that $\mathbf{A}$ has a full column rank, i.e., the rows of $\mathbf{A}$ (viewed as vector representations of paths) have a rank of $|L|$. Since the log path success probabilities are related to the log link success probabilities through[1] $\log \boldsymbol{\alpha} = \mathbf{A} \log \boldsymbol{\theta}$, this assumption implies that one can uniquely determine all the link success probabilities $\boldsymbol{\theta}$ from the (ground truth) path success probabilities $\boldsymbol{\alpha}$ via $\log \boldsymbol{\theta} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T \log \boldsymbol{\alpha}$. We will first present our solution under the full rank assumption, and later discuss how to apply our solution in the case of rank-deficient measurement matrix (Section 7).

2.3 Main Problem: Link Anomaly Detection

Given a minimum acceptable link success probability $\tau \in (0,\, 1)$ and a maximum tolerable false alarm probability $B \in (0,\, 1)$, we want to detect the presence of any abnormal link (i.e., a link with success probability below τ) as soon as possible, with a false alarm probability of at most B.

Specifically, given measurements $\mathbf{x}$ on paths P that span the link space, we want to test the binary hypotheses:

$$H_0:\ \theta_l \geq \tau,\ \forall l \in L \ \text{ vs. } \ H_1:\ \theta_l < \tau,\ \exists l \in L \qquad (3)$$

by a detection scheme $\delta(\mathbf{x}) \in \{0,\, 1\}$ (indicating whether H_0 or H_1 is detected) that minimizes the expected detection delay $\mathbb{E}[D(\delta)|H_1]$ subject to the false alarm constraint $P_F(\delta) \leq B$. Here $D(\delta)$ is the time from the end of probing to the instant that a decision (H_0 or H_1) is made. Any monitor (including the monitoring center) participating in fusion is a possible decision maker, and $D(\delta) := \max_{p\in P'} d_{v,p}$ if the decision is made by node $v \in Z \cap N_T$ based on measurements of paths $P' \subseteq P$. The false alarm probability $P_F(\delta)$ is the probability of detecting H_1 while H_0 is true.

Remark: The above problem is fundamentally different from the classical network tomography problem of estimating the link parameters $\boldsymbol{\theta}$ in that: (i) for a given link l, it suffices to determine the comparison between θ_l and τ, and (ii) for detecting H_1, it suffices to detect $\theta_l < \tau$ for one link. These differences make it possible to design a detection scheme that is more efficient than first estimating $\boldsymbol{\theta}$ and then comparing the estimates against τ.

2.4 Example

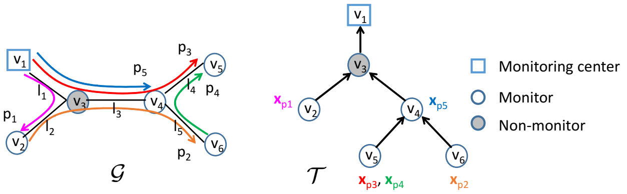

Figure 3: Sample network $\mathcal{G}$ and its fusion tree $\mathcal{T}$.

Consider the example in Fig. 3, where paths $p_1, \ldots, p_5$ are measured to detect if the success probability on any of links $l_1, \ldots, l_5$ is less than $\tau = 0.5$. Suppose that l_4 has a success probability of $\theta_4 = 0.4$, and the other links are lossless. Then neither of the paths traversing l_4 (i.e., p_3, p_4) has an abnormal success probability, but using network tomography, we can detect the anomaly on l_4 by estimating θ_4 using:

$$(\log \hat{\theta}_1, \ldots, \log \hat{\theta}_5)^T = \mathbf{A}^{-1}(\log \hat{\alpha}_1, \ldots, \log \hat{\alpha}_5)^T, \qquad (4)$$

where $\hat{\theta}_i$ is the estimated success probability of link l_i, $\hat{\alpha}_j$ is the empirical success probability of path p_j, and $\mathbf{A}$ is the measurement matrix given by

$$\mathbf{A} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 \end{pmatrix}. \qquad (5)$$

However, we do not have to solve (4). Since $\theta_4 = \alpha_3/\alpha_5$, node v_4 can estimate θ_4 and perform the detection using partial measurements from paths p_3 and p_5, achieving a smaller detection delay compared with the centralized solution of first collecting all the measurements at the monitoring center v_1 and then performing the detection.

[1] For a scalar function $g(\cdot)$ and a vector $\mathbf{x}$, $g(\mathbf{x})$ denotes applying $g(\cdot)$ to each element of $\mathbf{x}$.

3. PRELIMINARIES AND BASELINES

We start by introducing some preliminaries from detection theory and two baseline solutions.

3.1 Preliminaries: Optimal Detector

The optimal detector δ for testing binary hypotheses H_0 vs. H_1 under false alarm constraint β is the one that maximizes the *detection probability* $P_D(\delta) := \Pr\{\delta = 1|H_1\}$ subject to $P_F(\delta) \leq \beta$ ($P_F(\delta) := \Pr\{\delta = 1|H_0\}$). For simple hypotheses (e.g., $H_0 : \theta = \theta_0$ vs. $H_1 : \theta = \theta_1$), the optimal detector always exists. However, for composite hypotheses (e.g., $H_0 : \theta \in \Theta_0$ vs. $H_1 : \theta \in \Theta_1$), the optimal detector may not exist. Here, the optimal detector, called the *uniformly most powerful (UMP) detector*, is a detector that has the maximum $\Pr\{\delta = 1|\theta\}$ for *any* $\theta \in \Theta_1$ subject to $\sup_{\theta \in \Theta_0} \Pr\{\delta = 1|\theta\} \leq \beta$.

If the hypotheses are *one-sided*, e.g., $H_0 : \theta \geq \tau$ vs. $H_1 : \theta < \tau$, then the *Karlin-Rubin theorem* [4] gives a sufficient condition for the existence of UMP as follows. Given a likelihood function $f(x;\, \theta)$, if the likelihood ratio $L(x; \theta', \theta) := f(x; \theta')/f(x; \theta)$ is monotone increasing in $T(x)$ for any $\theta' > \theta$, then the problem is said to have a *monotone likelihood ratio (MLR)* in $T(x)$, and the detector

$$\delta(x) = \begin{cases} 1 & \text{if } T(x) < t, \\ \gamma & \text{if } T(x) = t, \\ 0 & \text{if } T(x) > t, \end{cases} \tag{6}$$

where t and γ are set to satisfy $\mathbb{E}[\delta(x)|\theta = \tau] = \beta$, is the UMP detector under false alarm constraint β. Here, $\delta(x) = \gamma$ ($\gamma \in [0,\, 1]$) means "detecting H_1 with probability γ". Specifically, the parameters are set by: $t = \max\{k : \Pr\{T(x) < k|\theta = \tau\} \leq \beta\}$, and $\gamma = (\beta - \Pr\{T(x) < t|\theta = \tau\})/\Pr\{T(x) = t|\theta = \tau\}$.

3.2 Baseline: Path-based Detection

Suppose that we want to detect anomalies directly based on end-to-end measurements. Given the measurements $\mathbf{x}_p$ of a $|p|$-hop path p, we cannot test the original hypotheses (3), as they may imply the same distribution of $\mathbf{x}_p$. This is because $\alpha_p \in [\tau^{|p|}, 1]$ under H_0 and $\alpha_p \in (0, \tau)$ under H_1, with $[\tau^{|p|}, 1] \cap (0, \tau) \neq \emptyset$ if $|p| > 1$. To ensure that the hypotheses are distinguishable by the measurements, i.e., implying different distributions of the measurements, we modify the hypotheses as

$$H_0 : \alpha_p \geq \tau^{|p|} \quad \text{vs.} \quad H_1 : \alpha_p < \tau^{|p|}, \tag{7}$$

since $\tau^{|p|}$ is the minimum success probability of p without any anomaly. Note that modifying the hypotheses is just one step in solving (3), and the detector derived from the modified hypotheses still needs to be evaluated on the original hypotheses (3).

This problem has an MLR in $\sum_{i=1}^n x_{p,i}$, as for any $\alpha'_p > \alpha_p$, the likelihood ratio $L(\mathbf{x}_p;\, \alpha'_p, \alpha_p)$ equals

$$\begin{aligned} L(\mathbf{x}_p;\, \alpha'_p, \alpha_p) &= \frac{\alpha\prime_p^{\sum_{i=1}^n x_{p,i}} (1 - \alpha'_p)^{n - \sum_{i=1}^n x_{p,i}}}{\alpha_p^{\sum_{i=1}^n x_{p,i}} (1 - \alpha_p)^{n - \sum_{i=1}^n x_{p,i}}} \\ &= \left(\frac{\alpha'_p (1 - \alpha_p)}{\alpha_p (1 - \alpha'_p)} \right)^{\sum_{i=1}^n x_{p,i}} \cdot \left(\frac{1 - \alpha'_p}{1 - \alpha_p} \right)^n, \end{aligned}$$

which is monotone increasing in $\sum_{i=1}^n x_{p,i}$. Therefore, by the Karlin-Rubin theorem [4], the UMP detector exists and equals:

$$\delta_p(\mathbf{x}_p) = \begin{cases} 1 & \text{if } \sum_{i=1}^n x_{p,i} < t, \\ \gamma & \text{if } \sum_{i=1}^n x_{p,i} = t, \\ 0 & \text{if } \sum_{i=1}^n x_{p,i} > t, \end{cases} \tag{8}$$

where the parameters t and γ are set to meet a given (per-path) false alarm probability β. Specifically, since conditioned on $\alpha_p = \tau^{|p|}$, $\sum_{i=1}^n x_{p,i}$ is a binomial random variable with parameters $(n,\, \tau^{|p|})$ (i.e., n trials with success probability $\tau^{|p|}$), $t = \max\{k : \Pr\{\sum_{i=1}^n x_{p,i} < k|\alpha_p = \tau^{|p|}\} \leq \beta\}$ can be calculated numerically by finding the maximum $k \in \{0,\, 1, \ldots, n\}$ such that $F_B(k-1;\, n, \tau^{|p|}) \leq \beta$, where $F_B(k-1;\, n, \tau^{|p|})$ is the *cumulative distribution function (CDF)* of the binomial distribution with parameters $(n,\, \tau^{|p|})$ at $k-1$. Then $\gamma = (\beta - \Pr\{\sum_{i=1}^n x_{p,i} < t|\alpha_p = \tau^{|p|}\})/\Pr\{\sum_{i=1}^n x_{p,i} = t|\alpha_p = \tau^{|p|}\}$, where $\Pr\{\sum_{i=1}^n x_{p,i} < t|\alpha_p = \tau^{|p|}\} = F_B(t-1;\, n, \tau^{|p|})$, and $\Pr\{\sum_{i=1}^n x_{p,i} = t|\alpha_p = \tau^{|p|}\} = \binom{n}{t} \tau^{t|p|}(1 - \tau^{|p|})^{n-t}$. The detector (8) is optimal for (7) in the sense that among all the detectors with false alarm probabilities bounded by β, it has the maximum probability of correctly detecting any $\alpha_p < \tau^{|p|}$.

Remark: This solution has the advantage that it can be implemented in a distributed manner, as the detector for each path p only needs the measurements on p. However, it has poor accuracy in cases where the mixing with good links makes a bad link undetectable, i.e., $\exists l' \in p$ with $\theta_{l'} < \tau$, but $\alpha_p = \prod_{l \in p} \theta_l \geq \tau^{|p|}$ (e.g., p_3 and p_4 in Section 2.4).

3.3 Baseline: Link-based Detection

Alternatively, we can first estimate link parameters by network tomography and then perform detection based on the estimates. Given the estimated link success probabilities $\log \hat{\boldsymbol{\theta}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T \log \hat{\boldsymbol{\alpha}}$, where $\hat{\boldsymbol{\alpha}} := (\hat{\alpha}_p)_{p \in P}$ ($\hat{\alpha}_p := \frac{1}{n}\sum_{i=1}^n x_{p,i}$) is the vector of empirical path success probabilities, the problem is to test the hypotheses

$$H_0 : \theta_l \geq \tau \quad \text{vs.} \quad H_1 : \theta_l < \tau \tag{9}$$

for each link $l \in L$.

In contrast to path-based detection, the UMP for (9) does not exist in general (see Section 4.2.2)[2]. Nevertheless, we can test

$$\delta_l(\mathbf{x}) = \begin{cases} 1 & \text{if } \log \hat{\theta}_l < t, \\ \gamma & \text{if } \log \hat{\theta}_l = t, \\ 0 & \text{if } \log \hat{\theta}_l > t, \end{cases} \tag{10}$$

where the parameters t and γ are set to meet a given (per-link) false alarm probability β. Denote $(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$ by $\mathbf{C} = (C_{l,p})_{l \in L, p \in P}$. Then $\log \hat{\theta}_l = \sum_{p \in P} C_{l,p} \log \hat{\alpha}_p$. Following the parameter setting in (6), we have: $t = \max\{k : \Pr\{\log \hat{\theta}_l < k|\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau\} \leq \beta\}$, and $\gamma = (\beta - \Pr\{\log \hat{\theta}_l < t|\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau\})/\Pr\{\log \hat{\theta}_l = t|\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau\}$, where $\mathbf{1}_{|L|}$ is an $|L|$-dimensional column vector of 1's, and $\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau$ means that $\theta_l = \tau$ for all $l \in L$. Although both parameters depend on the distribution of $\sum_{p \in P} C_{l,p} \log \hat{\alpha}_p$, which is difficult to compute for large $|P|$, we will provide an easier method to set the parameters later (see Section 4.2.5).

Remark: This solution has the advantage that it can accurately detect link anomalies for sufficiently large n. How-

[2] Note that each link (if identifiable) is a special MILS by the definition in Section 4.1.

ever, it incurs a large detection delay due to the need of first collecting all the empirical path success probabilities at the monitoring center. As shown in Section 2.4, it is possible to reduce the delay without sacrificing accuracy via distributed detection at intermediate nodes.

4. DISTRIBUTED DETECTION BASED ON PARTIAL NETWORK TOMOGRAPHY

Limitations of the solutions discussed in Sections 3.2-3.3 motivate us to explore alternatives with a better tradeoff between accuracy and delay. To this end, we propose to perform distributed detection based on partial network tomography.

To ease the presentation, we introduce the following notations: $P_v \subseteq P$ is the set of paths whose measurements will be available at v, $\mathbf{x}_v := (\mathbf{x}_p)_{p\in P_v}$ are the measurements of P_v, $\mathbf{A}_v$ is the sub-measurement matrix consisting of rows corresponding to paths in P_v, and $\boldsymbol{\alpha}_v := (\alpha_p)_{p\in P_v}$ is the vector of success probabilities of paths in P_v.

4.1 Minimal Identifiable Link Sequence (MILS)

After a monitor $v \in N_T$ receives measurements $\mathbf{x}_v$ of paths P_v, it tries to estimate the success probabilities of the links involved in P_v to detect any link anomaly. The challenge is that when $P_v \subset P$, the success probabilities of P_v may not uniquely determine the success probabilities of all the involved links.

To address this challenge, we leverage the concept of *minimal identifiable link sequence (MILS)* [36]. A MILS is a consecutive link sequence of minimal length whose parameter is identifiable (i.e., uniquely determinable) from the path parameters. In our case, each MILS s at node v is a minimal subpath of some $p \in P_v$ such that its success probability ζ_s is identifiable from $\boldsymbol{\alpha}_v$. For example, node v_4 in Fig. 3 will receive measurements on paths $p_2, \ldots, p_5$. It has four MILSs as shown in Fig. 4, where $l_i + l_j$ denotes the concatenation of links l_i and l_j. We have $\log\zeta_{l_1+l_3} = \log\alpha_5$, $\log\zeta_{l_2+l_3} = \log\alpha_2 + \log\alpha_3 - \log\alpha_4 - \log\alpha_5$, $\log\zeta_{l_4} = \log\alpha_3 - \log\alpha_5$, and $\log\zeta_{l_5} = \log\alpha_4 - \log\alpha_3 + \log\alpha_5$. It can also be verified that no subsequence of these MILSs is identifiable.

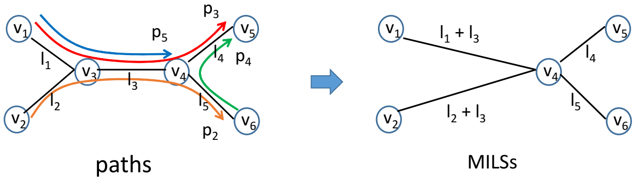

Figure 4: MILSs at node v_4 in Fig. 3.

Given a set of paths P_v, we can compute all the MILSs identifiable from these paths using a polynomial-time algorithm *Seek_MILS* in [36]. These MILSs represent the highest "resolution" at which node v can infer link parameters.

Another nice property of MILSs is that they capture all the information in the paths in the following sense.

Lemma 4.1. Let M_v be the set of all the MILSs identifiable from paths P_v. Then the path success probabilities $\boldsymbol{\alpha}_v$ are uniquely determined by the MILS success probabilities $\boldsymbol{\zeta}_v := (\zeta_s)_{s\in M_v}$.

4.2 MILS-based Detection

Since the MILS success probabilities are the most detailed information an intermediate node can infer about the performance of links, we perform anomaly detection at the level of MILSs. In the sequel, we will develop a concrete detection scheme and analyze its performance.

4.2.1 MILS-based Testing Statistic

The foundation of MILS-based detection is the calculation of (empirical) MILS success probabilities. Given a MILS $s \in M_v$, we know by definition that its log success probability $\log\zeta_s$ equals a linear combination of $\{\log\alpha_p\}_{p\in P_v}$, but it remains to compute the linear coefficients.

Let $\mathbf{s} := (s_l)_{l\in L}$ be the (column) vector representation of MILS s, where $s_l \in \{0, 1\}$ indicates whether link l is contained in s. Without loss of generality, we assume that the sub-measurement matrix $\mathbf{A}_v$ has a full row rank (otherwise, we can find a basis of its rows and consider the submatrix formed by this basis). We can compute the *(thin) QR decomposition* $\mathbf{A}_v^T = \mathbf{Q}_v\mathbf{R}_v$ [14], where $\mathbf{Q}_v$ is an $|L| \times |P_v|$ matrix with unit-length orthogonal columns, and $\mathbf{R}_v$ is a $|P_v| \times |P_v|$ invertible upper-triangular matrix. We have two key observations: (i) since $\mathbf{s}$ falls into the row space of $\mathbf{A}_v$ (by the definition of MILS), $\mathbf{s}^T = \mathbf{s}^T\mathbf{Q}_v\mathbf{Q}_v^T$; (ii) since $\mathbf{A}_v \log\boldsymbol{\theta} = \mathbf{R}_v^T\mathbf{Q}_v^T\log\boldsymbol{\theta} = \log\boldsymbol{\alpha}_v$, $\mathbf{Q}_v^T\log\boldsymbol{\theta} = \mathbf{R}_v^{-T}\log\boldsymbol{\alpha}_v$. Therefore,

$$\log\zeta_s = \mathbf{s}^T\log\boldsymbol{\theta} = \mathbf{s}^T\mathbf{Q}_v\mathbf{Q}_v^T\log\boldsymbol{\theta} = \mathbf{s}^T\mathbf{Q}_v\mathbf{R}_v^{-T}\log\boldsymbol{\alpha}_v. \quad (11)$$

Based on (11), we can estimate ζ_s, or equivalently $\log\zeta_s$. One estimator of particular interest is the *maximum likelihood estimator (MLE)*, which estimates ζ_s by the value that maximizes the likelihood of given measurements. This estimator is "optimal" in many senses, e.g., being the only candidate for efficient estimator (i.e., unbiased and achieving the *Cramér-Rao bound*) at finite sample sizes and asymptotically efficient at large sample sizes [32]. The following lemma states that the empirical version of (11) is actually the MLE of $\log\zeta_s$.

Lemma 4.2. Given the empirical path success probabilities $\hat{\boldsymbol{\alpha}}_v$, the MLE $\log\hat{\zeta}_s$ of the log success probability of MILS s is given by (where $\mathbf{s}$, $\mathbf{Q}_v$, and $\mathbf{R}_v$ are defined as in (11))

$$\log\hat{\zeta}_s = \mathbf{s}^T\mathbf{Q}_v\mathbf{R}_v^{-T}\log\hat{\boldsymbol{\alpha}}_v. \quad (12)$$

Later, we will develop a detector that uses $\log\hat{\zeta}_s$ as the testing statistic (see Section 4.2.3).

4.2.2 Nonexistence of Optimal MILS Detector

Given the measurements $\mathbf{x}_v$, node v cannot directly test the hypotheses (3) for the same reason as in Section 3.2. Since for each MILS $s \in M_v$, $\tau^{|s|}$ is the minimum success probability of s without any anomaly ($|s|$: number of links in s), we modify the hypotheses as

$$H_0 : \zeta_s \geq \tau^{|s|} \quad \text{vs.} \quad H_1 : \zeta_s < \tau^{|s|}. \quad (13)$$

As in Section 3.2, the modification is part of our solution, and the detector derived from the modified hypotheses will still be evaluated on the hypotheses in (3).

Although this problem appears similar to (7), there is a critical difference that the successes/losses on MILSs are not directly observable in general. As shown below, this difference makes the MILS-based detection more difficult.

Theorem 4.3. Generally, the UMP detector for (13) does not exist.

Algorithm 1 Distributed Detection Scheme

```
1: for slot $t = 0, 1, 2, \dots$ do
2:   for each node $v \in N_T \cap Z$ do
3:     for each MILS $s \in M_v$ such that all the statistics in $\{\hat{\alpha}_p\}_{p \in P_s}$ are received by $v$ at the beginning of slot $t$ do
4:       $v$ tests $s$ by the detector in (14)
5:       if the detector returns 1 then
6:         notify the network administrator of decision $H_1$ and stop
7:   if all the statistics in $\{\hat{\alpha}_p\}_{p \in P}$ have been received by the monitoring center $r$ then
8:     notify the network administrator of decision $H_0$ and stop
```

4.2.3 Heuristic MILS Detector

The lack of a (uniformly) optimal detector motivates the search of heuristic detectors with good performance. Given an estimator (12) of MILS success probability, a natural heuristic is to compare the estimated parameter with a threshold, i.e., we test (13) by

$$\delta_s(\mathbf{x}_v) = \begin{cases} 1 & \text{if } \log \hat{\zeta}_s(\mathbf{x}_v) < t, \\ \gamma & \text{if } \log \hat{\zeta}_s(\mathbf{x}_v) = t, \\ 0 & \text{if } \log \hat{\zeta}_s(\mathbf{x}_v) > t, \end{cases} \tag{14}$$

where $\log \hat{\zeta}_s(\mathbf{x}_v)$ is the estimated log success probability of MILS s given by (12), based on empirical path success probabilities $\hat{\boldsymbol{\alpha}}_v$ computed from measurements $\mathbf{x}_v$.

Given a maximum false alarm probability β, we can set t and γ to satisfy β while maximizing the probability of correctly detecting H_1 by:

$$t = \max\{k : \Pr\{\log \hat{\zeta}_s(\mathbf{x}_v) < k | \boldsymbol{\theta} = \mathbf{1}_{|L|}\tau\} \leq \beta\}, \tag{15}$$

$$\gamma = \frac{\beta - \Pr\{\log \hat{\zeta}_s(\mathbf{x}_v) < t | \boldsymbol{\theta} = \mathbf{1}_{|L|}\tau\}}{\Pr\{\log \hat{\zeta}_s(\mathbf{x}_v) = t | \boldsymbol{\theta} = \mathbf{1}_{|L|}\tau\}}, \tag{16}$$

where $\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau$ means $\theta_l = \tau$ for all $l \in L$ as in Section 3.3.

4.2.4 Overall Detection Scheme

Based on the detector in (14), we propose a distributed detection scheme shown in Algorithm 1. Let $P_s \subseteq P$ denote the subset of paths required to identify a MILS s, i.e., the paths that correspond to non-zero coefficients in $\mathbf{s}^T \mathbf{Q}_v \mathbf{R}_v^{-T}$ in (12). Under this scheme, every monitor v participating in fusion independently tests each of the MILSs in M_v by (14) as the required path statistics become available (line 4). If any test returns 1, then the overall decision is H_1 (line 6). Otherwise, if the fusion has completed, then the overall decision is H_0 (line 8).

Under this scheme, different local decisions for a given MILS s have asymmetric impacts on the global decision: $\delta_s(\mathbf{x}_v) = 1$ immediately triggers a global alarm, while $\delta_s(\mathbf{x}_v) = 0$ has no effect. This is due to the asymmetry in the hypotheses (3), where the network is considered abnormal if and only if any link is abnormal. The absence of anomaly in the measurements available at a given node does not preclude the presence of anomaly elsewhere, and hence the detection procedure should continue if $\delta_s(\mathbf{x}_v) = 0$.

Remark: A few remarks are in order regarding this scheme:

- It suffices for nodes to report the empirical path success probabilities, as the detector (14) only depends on $\hat{\boldsymbol{\alpha}}_v$.
- The detection procedure is concurrent with but orthogonal to the fusion procedure. Upon detecting an anomaly on a MILS, the network administrator may stop the fusion, continue it, or take other actions as needed (e.g., confirming the detection by sending more probes).

4.2.5 Parameter Setting

To apply the detection scheme in Algorithm 1, two problems must be resolved: (i) how to configure parameters (t and γ) of the detector in (14), and (ii) how to specify the per-detector false alarm probability β to ensure an acceptable overall false alarm probability. We now address these problems. In the sequel, we will simplify $\hat{\zeta}_s(\mathbf{x}_v)$ into $\hat{\zeta}_s$.

Setting Detection Threshold: We note that it is difficult to directly compute t and γ according to (15, 16). Specifically, let $\mathbf{c}_s := (c_{s,p})_{p \in P_v} = \mathbf{s}^T \mathbf{Q}_v \mathbf{R}_v^{-T}$. By (12), $\log \hat{\zeta}_s = \sum_{p \in P_v} c_{s,p} \log \hat{\alpha}_p$. To compute (15), we need to compute (all the probabilities are conditioned on $\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau$)

$$\begin{aligned} \Pr\{\log \hat{\zeta}_s < k\} &= \Pr\{\sum_{p \in P_v} c_{s,p} \log \hat{\alpha}_p < k\} \\ &= \sum_{\hat{\boldsymbol{\alpha}}_v \in \Phi(\mathbf{c}_s,\, k)} \prod_{p \in P_v} f_B(n\hat{\alpha}_p;\, n, \tau^{|p|}), \end{aligned} \tag{17}$$

where

$$\Phi(\mathbf{c}_s,\, k) := \{(y_p)_{p \in P_v} \in \{0, \frac{1}{n}, \dots, 1\}^{|P_v|} : \sum_{p \in P_v} c_{s,p} \log y_p < k\}$$

is the set of empirical path success probabilities satisfying $\log \hat{\zeta}_s < k$, and $\prod_{p \in P_v} f_B(n\hat{\alpha}_p;\, n, \tau^{|p|})$ is the joint probability for these empirical path success probabilities to equal $\hat{\boldsymbol{\alpha}}_v$ (note that the empirical success probabilities on different paths are independent). Here $f_B(m; n, q) := \binom{n}{m} q^m (1-q)^{n-m}$ is the *probability mass function (PMF)* of the binomial distribution with parameters $(n,\, q)$ at m. Computing (17) requires enumerating all the elements in $\Phi(\mathbf{c}_s, k)$, which is highly complicated for large n and $|P_v|$. Similar arguments will show the complexity of computing (16).

To address this challenge, we introduce two simplifications. First, we avoid computing (16) by simply setting $\gamma = 0$. This is because for $n \gg 1$, $\hat{\zeta}_s$ is nearly continuously distributed, with $\Pr\{\log \hat{\zeta}_s = t\} \approx 0$ (for any t). Thus, we can ignore this case by setting γ to zero without violating the given false alarm constraint. Moreover, instead of computing the exact value of $\Pr\{\log \hat{\zeta}_s < k\}$ as in (17), we bound it as follows.

Lemma 4.4. Given a MILS s with $\log \zeta_s = \sum_{p \in P_v} c_{s,p} \log \alpha_p$, conditioned on $\boldsymbol{\theta} = \mathbf{1}_{|L|}\tau$, we have

$$\Pr\{\log \hat{\zeta}_s < k\} \leq \sum_{p \in P_v : c_{s,p} > 0} F_B(\lceil n e^{\frac{k}{|P_v| c_{s,p}}} \rceil - 1;\, n, \tau^{|p|}) \tag{18}$$

for any $k < 0$, where $F_B(m; n, q)$ is the CDF of the binomial distribution of parameters (n, q) at m.

By Lemma 4.4, we can set t to the largest k such that the righthand side (RHS) of (18) is bounded by β. The result, together with $\gamma = 0$, will guarantee that the false alarm probability of the detector in (14) is bounded by β. Note that requiring $k < 0$ does not lose generality, as it is required to satisfy $\Pr\{\log \hat{\zeta}_s < k\} \leq \beta$ for any $\beta < 1$.

Discussion: We note that the above method of setting the detection threshold can be conservative, i.e., the computed threshold can be much smaller than the targeted threshold according to (15), resulting in a false alarm probability that is much smaller than β. In this case, an alternative is to directly use the empirical CDF of $\log\hat{\zeta}_s$ to approximate $\Pr\{\log\hat{\zeta}_s < k\}$. Specifically, we can draw m samples of the empirical path success probabilities $(\hat{\alpha}_{p,i})_{p\in P_v}$ $(i = 1, \ldots, m)$, where each $n \cdot \hat{\alpha}_{p,i}$ is binomially distributed with parameters $(n, \tau^{|p|})$. Using these samples we can obtain m realizations of $\log\hat{\zeta}_s$ by $\log\hat{\zeta}_{s,i} = \sum_{p\in P_v} c_{s,p} \log\hat{\alpha}_{p,i}$ for $i = 1, \ldots, m$. We sort these realizations into increasing order: $\log\hat{\zeta}_s^{(1)} \leq \ldots \leq \log\hat{\zeta}_s^{(m)}$. Then the 100β-th percentile[3] of the empirical distribution of $\log\hat{\zeta}_s$, given by $\log\hat{\zeta}_s^{(\lfloor\beta m\rfloor+1)}$, gives an approximation of (15). It is easy to see that this approximation becomes accurate as $m \to \infty$.

Setting False Alarm Bound: Given an overall false alarm bound B, we need to translate this bound into a per-MILS false alarm bound β for applying (14). Let $M := \bigcup_{v\in N_{\mathcal{T}}\cap Z} M_v$ be the set of distinct MILSs that will be tested. Let $P_F(M)$ denote the overall false alarm probability, and $P_F(s)$ $(s \in M)$ the false alarm probability for MILS s. We have the following relationship between the two.

Lemma 4.5. *We always have $P_F(M) \leq \sum_{s\in M} P_F(s)$. Moreover, if $c_{s,p} \geq 0$ for all $s \in M$ and $p \in P$, where $\log\zeta_s = \sum_{p\in P} c_{s,p}\log\alpha_p$, then $P_F(M) \leq 1 - \prod_{s\in M}(1 - P_F(s))$.*

By Lemma 4.5, we can set the per-MILS false alarm bound by (assuming the same bound for all the MILSs)

$$\beta = \begin{cases} 1-(1-B)^{\frac{1}{|M|}} & \text{if } c_{s,p} \geq 0,\ \forall s \in M,\ p \in P, \\ B/|M| & \text{o.w.} \end{cases} \tag{19}$$

This setting makes the upper bound on $P_F(M)$ equal to B.

Remark: For $|M| \gg 1$, the Taylor expansion shows that $1-(1-B)^{1/|M|} = -\log(1-B)/|M| + o(1/|M|) \geq B/|M| + o(1/|M|)$, suggesting that $1-(1-B)^{1/|M|}$ is a better choice if the condition $c_{s,p} \geq 0$ ($\forall s \in M$ and $p \in P$) is satisfied.

4.2.6 Performance Analysis

We now analyze the performance of the detection scheme proposed in Section 4.2.4 in terms of false alarm probability, detection probability, and detection delay.

Regarding false alarm probability, the parameter setting in Section 4.2.5 guarantees the following.

Theorem 4.6. *If for each MILS $s \in M$, $\gamma = 0$ and t is set to k such that (18) is bounded by the value of β given by (19), then the false alarm probability of the overall detection scheme in Algorithm 1 satisfies $P_F(M) \leq B$.*

Regarding detection probability (i.e., the probability of detecting H_1 when there exists an anomalous link), we note that the performance usually depends on the specific value of link parameters $\boldsymbol{\theta}$. Let $P_D(M|\boldsymbol{\theta})$ denote the overall detection probability under $\boldsymbol{\theta}$ (M: the set of all the tested MILSs). We bound $P_D(M|\boldsymbol{\theta})$ by analyzing its complement, the *miss probability* for each MILS. Let $P_M(s|\boldsymbol{\theta}) := \Pr\{\log\hat{\zeta}_s \geq t_s|\boldsymbol{\theta}\}$ denote the per-MILS miss probability for a MILS s, i.e., the probability of not detecting anomaly on s using the threshold t_s under link parameters $\boldsymbol{\theta}$. We will need the following constants. Given a MILS s with $\log\zeta_s = \sum_{p\in P} c_{s,p}\log\alpha_p$, let $P_s := \{p \in P:\ c_{s,p} \neq 0\}$ be the set of paths used to identify s (as defined in Section 4.2.4), $\alpha_s := \min_{p\in P_s}\alpha_p$ be the minimum success probability for these paths, $c_s^{\min} := \min_{p\in P_s}|c_{s,p}|$ be the minimum absolute coefficient, and $c_s^{\max} := \max_{p\in P_s}|c_{s,p}|$ be the maximum absolute coefficient.

Lemma 4.7. *If under link parameters $\boldsymbol{\theta}$, a MILS s satisfies $\log\zeta_s \leq t_s - \epsilon$ for some $\epsilon > 0$, then*

$$P_M(s|\boldsymbol{\theta}) \leq |P_s| \exp\left[-2n\alpha_s^2\left(\frac{\exp\left(\frac{\epsilon}{c_s^{\max}|P_s|}\right)-1}{\exp\left(\frac{\epsilon}{c_s^{\min}|P_s|}\right)}\right)^2\right], \tag{20}$$

i.e., the miss probability for s decays exponentially with n, the number of measurements per path.

Based on Lemma 4.7, we can bound $P_M(s|\boldsymbol{\theta})$ for arbitrary MILSs and link parameters by defining the following quantity:

$$\eta_s(\boldsymbol{\theta}) = \begin{cases} \text{RHS of (20) with } \epsilon = t_s - \log\zeta_s & \text{if } \log\zeta_s < t_s, \\ 1 & \text{o.w.} \end{cases} \tag{21}$$

Then we always have $P_M(s|\boldsymbol{\theta}) \leq \eta_s(\boldsymbol{\theta})$. Note that α_s in the RHS of (20) and ζ_s are both functions of $\boldsymbol{\theta}$.

Since our detection scheme detects the anomaly as long as it detects anomaly on one of the MILSs, Lemma 4.7 implies the following bound on $P_D(M|\boldsymbol{\theta})$.

Theorem 4.8. *Under link parameters $\boldsymbol{\theta}$, the detection probability of the overall detection scheme in Algorithm 1 satisfies $P_D(M|\boldsymbol{\theta}) \geq 1 - \min_{s\in M}\eta_s(\boldsymbol{\theta})$ for $\eta_s(\boldsymbol{\theta})$ in (21).*

Remark: Given a set of paths P and a fusion tree $\mathcal{T}$, the MILSs tested at each node can be computed by the *Seek_MILS* algorithm [36]. Moreover, given a detection threshold t_s and link parameters $\boldsymbol{\theta}$, $\eta_s(\boldsymbol{\theta})$ can be computed in closed form. Therefore, the detection probability bound in Theorem 4.8 is easily computable.

Another important performance metric is the detection delay. Since our detection scheme stops as soon as an anomaly is detected on one of the MILSs, the detection delay is closely related to the detection probability (or the miss probability) for each MILS. Using Lemma 4.7, we can bound the expected detection delay as follows.

Theorem 4.9. *Let T denote the time to fuse all the path statistics at r, M_d denote the set of MILSs tested by slot d, and $\mathbf{M} := (M_d)_{d=0}^{T}$. Let $D(\mathbf{M})$ denote the detection delay under $\mathbf{M}$. Under link parameters $\boldsymbol{\theta}$, the expected detection delay of the overall detection scheme in Algorithm 1 satisfies*

$$\mathbb{E}[D(\mathbf{M})|\boldsymbol{\theta}] \leq \mathbb{E}[\sum_{d=0}^{T-1}\min_{s\in M_d}\eta_s(\boldsymbol{\theta})], \tag{22}$$

where the expectation on the RHS of (22) is over $\mathbf{M}$.

Remark: Although the bound in Theorem 4.9 is not in closed form, it sheds light on the asymptotic delay performance of the proposed detection scheme. By Lemma 4.7, $\eta_s(\boldsymbol{\theta})$ decays exponentially with n if $\log\zeta_s < t_s$. For $n \gg 1$, this implies that $\min_{s\in M_d}\eta_s(\boldsymbol{\theta}) \approx 0$ as long as $\exists s \in M_d$ with $\log\zeta_s < t_s$. On the other hand, if $\log\zeta_s \geq t_s$ for all $s \in M_d$, then $\min_{s\in M_d}\eta_s(\boldsymbol{\theta}) = 1$ by definition. Thus, we have

$$\lim_{n\to\infty}\mathbb{E}[\sum_{d=0}^{T-1}\min_{s\in M_d}\eta_s(\boldsymbol{\theta})] = \mathbb{E}[\min(T,\ T')], \tag{23}$$

[3] Precisely, the empirical probability of $\log\hat{\zeta}_s < k$ is bounded by β if and only if $k \leq \log\hat{\zeta}_s^{(\lfloor\beta m\rfloor+1)}$. Thus, $\log\hat{\zeta}_s^{(\lfloor\beta m\rfloor+1)}$ is the solution to (15) based on the empirical distribution.

where $T' := \min\{d :\ \log \zeta_s < t_s,\ \exists s \in M_d\}$ is the first time we test a detectable MILS ($T' := \infty$ if none of the MILSs is detectable). This implies that asymptotically, the proposed detection scheme achieves the minimum detection delay, as $\min(T, T')$ is the earliest time that one can hope to make a correct decision.

5. FURTHER OPTIMIZATION

The detection performance is fundamentally limited by the measurements available to the detectors. Although we have previously assumed that the probing paths P and the fusion tree $\mathcal{T}$ are given, in practice they are usually parameters that can be designed.

5.1 Measurement Design Problem

Assume that the path $p(u, v)$ for each source-destination pair (u, v) ($u,\ v \in N$) is predetermined and symmetric (i.e., $p(u, v) = p(v, u)$). Moreover, assume that we can only probe between monitors and fuse measurements at one of the monitors. The problem of measurement design includes three subproblems: (1) *probing path selection*, which selects a subset P from all the possible probing paths $\mathcal{P} := \{p(u, v) :\ u,\ v \in Z\}$, (2) *probing path orientation*, which determines the source/destination for each $p \in P$, and (3) *monitoring center placement*, which selects a node as the monitoring center to fuse measurements. Here, we assume that this node must be a monitor (i.e., $r \in Z$), but our solution easily extends to other cases. Jointly, these three subproblems uniquely determine the testable MILSs by determining the fusion tree $\mathcal{T}$ and the paths whose measurements will be fused at each node in $\mathcal{T}$.

Ideally, an optimal design should optimize the detection performance. As is evident from Section 4.2.6, however, the detection performance depends on the specific values of link parameters $\boldsymbol{\theta}$, which are unavailable at design time. To address this issue, we use the maximum length (in hop count) of fusion paths as a *proxy* of the detection performance. Let N_D be the probing destinations of paths in P. The problem is then to jointly design P, N_D, and r such that

$$\min \max_{v \in N_D} |p(v, r)| \tag{24a}$$

$$\text{s.t. } P \text{ spans the link space,} \tag{24b}$$

$$P \subseteq \mathcal{P}, \tag{24c}$$

$$r \in Z, \tag{24d}$$

where $|p|$ is the length of path p. Since in the worst case, all the (statistics of) measurements have to be fused at the monitoring center r before a decision can be made, a design by (24) minimizes the proxy of the worst-case detection delay. Note that the delay performance of a design will still be evaluated by the actual detection delay (see Section 6).

The formulation (24) implicitly assumes that we can identify all the links using all the possible probing paths in $\mathcal{P}$ (i.e., $\mathcal{P}$ spans the link space). In cases violating this assumption (due to limitation on monitor locations), we can transform the network topology as in [36] such that links in the transformed topology represent MILSs in the original topology and are thus identifiable.

5.2 Relationship to Matroid Optimization

Clearly, the designs of P, N_D, and r are inter-dependent, which complicates the problem as they need to be solved jointly. Moreover, both P and N_D have solution spaces that are exponentially large in the size of the network. Nevertheless, we will show that (24) can be solved optimally in polynomial time due to the following observations:

1) We observe that r only has a linear-sized solution space, and N_D can be easily optimized given P and r (i.e., orienting each probing path towards the endpoint closer to r).

2) We observe that given r, the subproblem of optimizing P can be solved in polynomial time as a variation of the *matroid optimization problem* defined as follows.

Definition 1 ([12])**.** Given a ground set E and a family of its subsets $\mathcal{I} \subseteq 2^E$, the pair $\mathcal{M} = (E, \mathcal{I})$ is a *matroid* if: (i) $X \subseteq Y$ and $Y \in \mathcal{I}$ imply $X \in \mathcal{I}$, and (ii) $X \in \mathcal{I}$, $Y \in \mathcal{I}$, and $|Y| > |X|$ imply $\exists e \in Y \setminus X$ such that $X \cup \{e\} \in \mathcal{I}$. A set $B \subseteq E$ is called a *basis* of $\mathcal{M}$ if $B \in \mathcal{I}$ and no proper superset of B is in $\mathcal{I}$.

Definition 2 ([12])**.** Given a matroid $\mathcal{M} = (E, \mathcal{I})$ and a weight function $c :\ E \to \mathbb{R}^+$, the *matroid optimization problem* is to compute a basis B^* of $\mathcal{M}$ that maximizes $\sum_{e \in B^*} c(e)$.

It is known that the greedy algorithm is optimal for the matroid optimization problem.

Theorem 5.1 ([12])**.** *The greedy algorithm, which examines elements of E in decreasing order of weights to build a basis, is optimal for the matroid optimization problem in Definition 2 for any matroid and any weight function.*

Matroid generalizes the notion of linear independence in vector space. In particular, when paths are viewed as vectors in the $|L|$-dimensional link space (i.e., rows of the measurement matrix defined in Section 2.2), the set of candidate paths $\mathcal{P}$ and the collection of linearly independent subsets of paths $\mathcal{I}_P$ form a matroid, denoted by $\mathcal{M}_P$. Then the constraint of (24b) is equivalent to "P is a basis of $\mathcal{M}_P$". Moreover, given $r \in Z$, we can define a weight function $c_P(p(u, v); r) := \min(|p(u, r)|, |p(v, r)|)$ to measure the minimum path length for fusing measurements on path $p(u, v)$ to node r. Thus, given a monitoring center r, the optimization in (24) is equivalent to a variation of the matroid optimization problem:

$$\min \max_{p \in P} c_P(p;\ r) \tag{25a}$$

$$\text{s.t. } P \text{ is a basis of } \mathcal{M}_P. \tag{25b}$$

Different from the problem in Definition 2 that optimizes the sum weight, this problem optimizes the worst weight[4]. However, we show that the greedy algorithm remains optimal.

Corollary 5.2. *Given a monitoring center r, the greedy algorithm, which examines candidate paths $p \in \mathcal{P}$ in increasing order of $c_P(p;\ r)$ to build a basis, is optimal for (25).*

5.3 Measurement Design Algorithm

Based on the observations in Section 5.2, we develop an algorithm called *Greedy Measurement Design (GMD)*, shown in Algorithm 2. GMD enumerates all possible locations of the monitoring center r and uses the greedy algorithm to select P for each r (lines 3-9). The overall solution

[4] Both problems can be formulated as equivalent minimization/maximization problems with analogous solutions.

Algorithm 2 Greedy Measurement Design (GMD)

```
1: $r^* \leftarrow \emptyset$, $P^* \leftarrow \emptyset$, $o^* \leftarrow \infty$
2: for each $r \in Z$ do
3:     sort $\mathcal{P}$ into $\{p_1,\ p_2, \ldots\}$ such that $c_P(p_1;\ r) \leq c_P(p_2;\ r) \leq \ldots$
4:     $P \leftarrow \emptyset$
5:     for $i = 1, \ldots, |\mathcal{P}|$ do
6:         if $P \cup \{p_i\}$ are linearly independent then
7:             $P \leftarrow P \cup \{p_i\}$
8:             if $|P| = |L|$ then
9:                 break
10:    if $\max_{p \in P} c_P(p;\ r) < o^*$ then
11:        $r^* \leftarrow r$, $P^* \leftarrow P$, $o^* \leftarrow \max_{p \in P} c_P(p;\ r)$
12: return $r^*$ and $P^*$
```

is then the $(r^*,\ P^*)$ that minimizes the objective function $\max_{p \in P} c_P(p;\ r)$ (lines 10-11). Note that given r^* and P^*, the optimal probing destination for each $p(u, v) \in P^*$ is simply the node $w \in \{u, v\}$ with smaller $|p(w, r^*)|$.

Complexity: There are $|Z|$ candidates of r, and for each of these the greedy algorithm takes $O(|Z|^2 \log |Z|)$ to sort $\mathcal{P}$ and $O(|Z|^2|L|^3)$ to build a basis (where each test of linear independence takes $O(|L|^3)$ by Gaussian elimination). Thus, the overall complexity of GMD is $O(|Z|^3|L|^3)$.

Optimality: Corollary 5.2 directly implies the following.

Theorem 5.3. Algorithm GMD is optimal for the measurement design problem defined in (24).

6. PERFORMANCE EVALUATION

We evaluate the proposed solutions via extensive simulations on real Internet Service Provider (ISP) topologies. Our goals are: (1) comparing the MILS-based detection scheme in Section 4.2 ('MILS') with the path-based detection scheme in Section 3.2 ('path') and the link-based detection scheme in Section 3.3 ('link'), (2) understanding the impacts of various input parameters, including the sample size per path, the false alarm constraint, the network topology, and the type of anomaly, and (3) comparing the basic MILS-based scheme with its variation based on the optimized measurement design in Section 5 ('MILS-O').

6.1 Dataset

We simulate networks according to the Rocketfuel topologies [30], where each node represents a *Point of Presence (POP)*, i.e., a set of co-located backbone/access routers. We select five topologies with various numbers of nodes and average node degrees to represent networks of different sizes/densities. We assume shortest path routing.

For each topology, we select monitors as follows:

1) Select all the nodes with degree 1 or 2 as monitors, as they must be monitors to identify their adjacent links.

2) Assuming all the remaining nodes are monitors, sort them into increasing order of the numbers of links they cover (i.e., shortest paths starting/ending at them cover), and remove each node from the set of monitors if it can be removed without losing identifiability.

The result is a set of monitors that can identify all the links by probing the shortest paths between each other. Note that this method is only used to place monitors for evaluation purpose, and our solution can be used in conjunction with other monitor placement methods.

Table 1: Properties of the Simulated Networks

Network	$\lvert N\rvert$	$\lvert L\rvert$	$\lvert Z\rvert$	$d_{\mathcal{T}}$	$d_{\mathcal{T}^*}$
Exodus	22	51	16	4	3
EUROPEgraph	28	66	20	4	3
Abovenet	22	80	20	3	2
SprintlinkUS	44	106	33	4	4
TiscaliEurope	51	129	43	4	4

We then perform two measurement designs: (i) a *random design*, by randomly selecting a basis of paths from all the monitor-to-monitor shortest paths, randomly selecting a probing destination for each path, and randomly selecting a monitoring center from all the monitors, and (ii) an *optimized design* by GMD (Algorithm 2).

The characteristics of each network are presented in Table 1, where $d_{\mathcal{T}}$ is the maximum length of fusion paths under one realization of the random design, and $d_{\mathcal{T}^*}$ is the same parameter under the optimized design.

6.2 Parameter Setting

Given an overall false alarm bound B, we configure the path-based detectors as in Section 3.2 (with per-path false alarm bound $\beta = B/|P|$), and the link/MILS-based detectors as in Section 4.2.5 (with per-link false alarm bound $\beta = B/|L|$ and per-MILS false alarm bound given by (19)). We set the detection threshold for each link/MILS based on the empirical CDF of its log success probability.

To evaluate the performance under H_0, we measure the false alarm probability when all the links have the minimum acceptable success probability τ. To evaluate the performance under H_1, we measure the detection probability and the expected detection delay when σ-fraction of randomly selected links are abnormal, where the success probabilities of normal links are drawn uniformly from $[\tau,\ 1]$, and those of abnormal links are drawn uniformly from $[\tau_{\min},\ \tau]$. We fix $\tau = 0.9$ and vary σ and $\tau_{\min}$ to study various types of anomaly. All the results are from 1000 Monte Carlo runs.

6.3 Results

6.3.1 Overall Performance Comparison

Fig. 5 shows the comparison between the benchmarks ('path', 'link'), the proposed solution ('MILS'), and the proposed solution with further optimization ('MILS-O'), for each of the networks. For a fair comparison, we set σ differently for different networks such that the number of abnormal links $\sigma|L|$ is the same for all the networks. Here, the false alarm probability is only evaluated to verify validity of the solutions, where a solution is considered valid if its false alarm probability is bounded by B, and the emphasis is on evaluating the detection probability (the larger the better) and the detection delay (the smaller the better). The results verify that all the solutions are valid under our parameter setting (Fig. 5 (a)). Meanwhile, the MILS-based detection scheme significantly outperforms the path-based detection scheme in detection probability (Fig. 5 (b)) and the link-based detection scheme in detection delay (Fig. 5 (c)). Moreover, the improvement in both performance measures increases when we combine the MILS-based detection scheme with the optimized measurement design. While we only show the comparison under one set of parameters, similar comparisons hold under other parameter values. Note that as our definition of detection delay only counts the communication delay in fusing measurements, the path-based scheme always has zero detection delay.

Note that a large fraction of monitors (71%–91%) is used in this evaluation to ensure a full-column-rank measurement matrix as assumed in Section 2.2. When all the nodes are monitors, simply probing all the one-hop paths between monitors and performing path-based detection will achieve the optimal detection accuracy and delay, considerably improving the performance in Fig. 5. Meanwhile, the proposed detection scheme can also be applied in a rank-deficient setting with much fewer monitors; see Section 7.

6.3.2 Impact of Network Topology

We can also use the results in Fig. 5 to evaluate the impact of network topology on detection performance. All the networks except for Abovenet have similar average node degrees, with $2|L|/|N| \in (4.6, 5)$, but different sizes, with $|N|$ ranging from 22 to 51. As the number of abnormal links is fixed, we expect that the larger the network, the harder the detection, because a larger network tends to have longer probing paths that mix the abnormal links with more normal links, as well as longer fusion paths that take longer to fuse the measurements. Indeed, we see from Fig. 5 (b-c) that the larger networks (SprintlinkUS and TiscaliEurope) have relatively worse performance than the smaller ones (Exodus and EUROPEgraph). Meanwhile, the network density also has a significant impact. The first three networks (Exodus, EUROPEgraph, and Abovenet) all have similar sizes in terms of $|N|$, but very different densities, where the average node degree is 4.6, 4.7, and 7.3, respectively. Under shortest path routing, we expect that the denser the topology, the shorter the probing/fusion paths, and hence it becomes easier and quicker to detect the abnormal links from path measurements. Indeed, we see from Fig. 5 (b-c) that Abovenet, which is much denser than the other two networks, has the best performance in both detection probability and delay.

6.3.3 Impact of Sample Size

To evaluate the impact of sample size, we select one of the networks and vary the parameter n. As shown in Fig. 6 (b), when the sample size increases beyond a certain value (here 1000) to allow reasonably accurate estimation of link success probabilities, the link-based detection outperforms the path-based detection in detection probability due to its finer resolution, and the MILS-based detection further improves the detection probability by running a larger set of tests. Meanwhile, Fig. 6 (c) shows that the MILS-based detection can detect anomalies much earlier than the link-based detection, and the improvement increases with n, as a larger sample size makes it easier for the intermediate nodes to detect anomalies using more accurate estimates of the MILS success probabilities. A special remark is required about the superior performance of 'MILS-O'. For this network, the optimized design not only shortens the fusion paths, but also shortens the probing paths, which is why it can reduce the detection delay while increasing the detection probability.

6.3.4 Impact of False Alarm Constraint

To evaluate the impact of false alarm constraint, we select one of the networks and vary the parameter B. Fig. 7 (a) shows that as we relax the false alarm constraint, all the schemes exploit this by raising their false alarm probabilities (via raising their detection thresholds). Accordingly, all the schemes achieve increased detection probabilities (Fig. 7 (b)) and decreased detection delays (Fig. 7 (c)). The amount of improvement is, however, insignificant, suggesting that the detection performance is relatively insensitive to the false alarm constraint.

6.3.5 Impact of Type of Anomaly

To understand how the properties of link anomalies affect the detection performance, we vary the input parameters to evaluate two effects: the *scale of anomaly*, controlled by the ratio of abnormal links σ, and the *extent of anomaly*, controlled by the minimum success probability of abnormal links $\tau_{\min}$. Fig. 8 and 9 show that both parameters have significant impacts on the detection performance. As expected, Fig. 8 shows that a larger scale anomaly (as σ increases) is easier to detect, and Fig. 9 shows that a more subtle anomaly (as $\tau_{\min}$ increases) is harder to detect. We omit the false alarm probabilities from these plots since they do not depend on properties of the anomaly. We have separately evaluated the false alarm probabilities for the four detection schemes under the parameter settings used in these simulations, which equal 0.085, 0.067, 0.067, and 0.085, respectively (all below $B = 0.1$).

7. APPLICATION TO RANK-DEFICIENT MEASUREMENT MATRIX

Although we have assumed that the measurement matrix $\mathbf{A}$ has a full column rank (Section 2.2), our distributed detection scheme proposed in Section 4 can be applied *as is* in scenarios with rank-deficient measurement matrices. Indeed, since each node independently tests anomalies based on locally available measurements, it must handle a possibly rank-deficient sub-measurement matrix, regardless of whether the overall measurement matrix is full-rank or not. Furthermore, our measurement design algorithm Algorithm 2 proposed in Section 5 can also be applied with a minor modification in line 8: we change the condition to $|P| = \text{rank}(\mathcal{P})$ to ensure that the selected paths form a basis of the subspace of the link space spanned by all the candidate paths $\mathcal{P}$.

With a rank-deficient measurement matrix, however, it is no longer possible to perform link-based detection at the monitoring center. Nevertheless, the monitoring center can apply our MILS-based detector to each MILS that is identifiable from all the probing paths, after fusing all the measurements. In the case of full-rank measurement matrix, each MILS at the monitoring center contains a single link, and this scheme reduces to link-based detection. We have verified that the performance advantages of our distributed detection scheme as observed in Section 6 remain valid in rank-deficient scenarios in comparison with the above centralized detection scheme and the (distributed) path-based detection scheme; see the appendix for details.

8. CONCLUSION

We study the problem of detecting links with abnormally high loss probabilities from binary end-to-end measurements. Having shown the limitations of path-based detection and link-based detection, we propose a novel distributed detection scheme, which tests for anomaly at the highest possible resolution (i.e., MILSs) throughout the measurement fusion process. We provide efficient methods to configure the pro-

Figure 5: Rocketfuel topologies ($B = 0.1$, $n = 2000$, $\tau = 0.9$, $\tau_{\min} = 0.8$, $\sigma|L| = 2$)

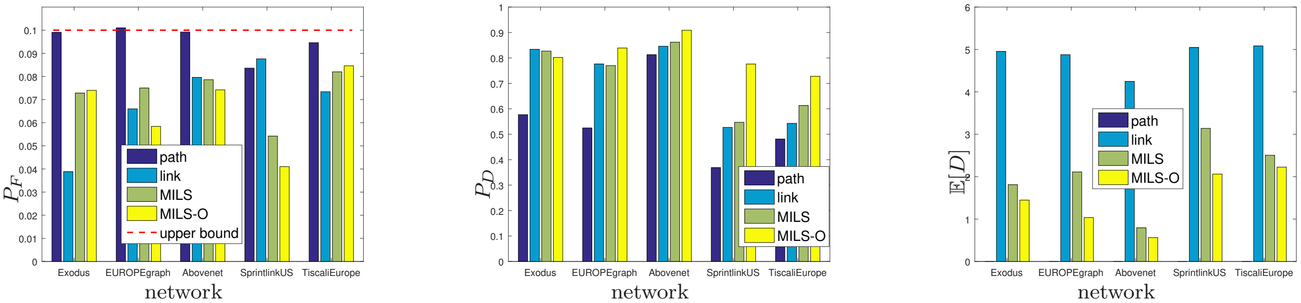

(a) false alarm probability (b) detection probability (c) expected delay

Figure 6: SprintlinkUS: varying number of probes n ($B = 0.1$, $\tau = 0.9$, $\tau_{\min} = 0.8$, $\sigma = 0.01$)

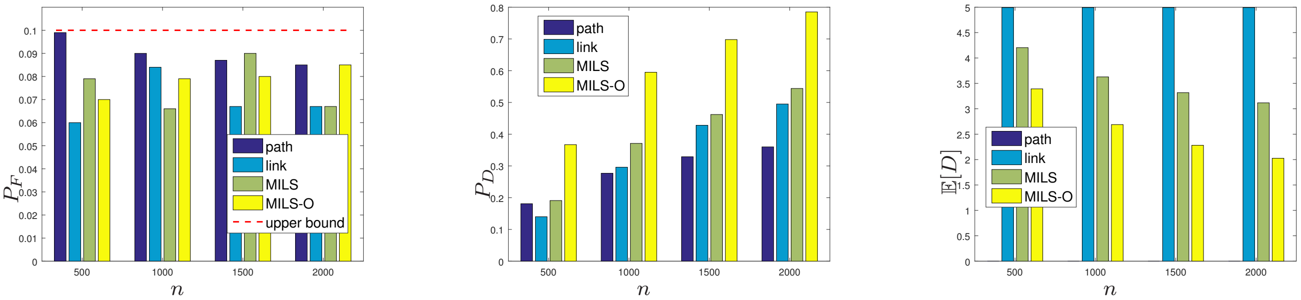

(a) false alarm probability (b) detection probability (c) expected delay

posed scheme to satisfy any given false alarm constraint, while providing analytical bounds on its detection probability and detection delay. We further develop an algorithm to optimize the measurement design. We show via simulations on real topologies that under the same false alarm constraint, our solution significantly outperforms path-based detection in detection probability and link-based detection in detection delay. In contrast to traditional applications of network tomography, our solution allows faster anomaly detection without sacrificing accuracy, and thus has promising applications in scenarios where only the presence of abnormal links is needed (e.g., detection of SLA violations). While we do not model constraints on the number/locations of monitors, our solution remains applicable with rank-deficient measurement matrices induced by such constraints (see Section 7). We leave a detailed study of the anomaly detection problem under these constraints to future work.

9. REFERENCES

[1] S. Ahuja, S. Ramasubramanian, and M. Krunz. SRLG failure localization in all-optical networks using monitoring cycles and paths. In *IEEE INFOCOM*, 2008.

[2] S. Ahuja, S. Ramasubramanian, and M. Krunz. SRLG failure localization in optical networks. *IEEE/ACM Transations on Networking*, 19(4):989–999, Auguest 2011.

[3] Y. Bejerano and R. Rastogi. Robust monitoring of link delays and faults in IP networks. In *IEEE INFOCOM*, 2003.

[4] G. Casella and R. L. Berger. *Statistical Inference*. Duxbury, 2002.

[5] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu. Network tomography: recent developments. *Statistical Science*, 2004.

[6] A. Chen, J. Cao, and T. Bu. Network tomography: Identifiability and Fourier domain estimation. In *IEEE INFOCOM*, 2007.

[7] Y. Chen, D. Bindel, and R. H. Katz. An algebraic approach to practical and scalable overlay network monitoring. In *ACM SIGCOMM*, 2004.

[8] S. Cho and S. Ramasubramanian. Localizing link failures in all-optical networks using monitoring tours. *Elsevier Computer Networks*, 58:2–12, January 2014.

[9] M. Coates, A. O. Hero, R. Nowak, and B. Yu. Internet tomography. *IEEE Signal Processing Magzine*, 19:47–65, 2002.

[10] N. Duffield. Simple network performance tomography. In *ACM SIGCOMM conference on Internet measurement*, 2003.

[11] N. Duffield. Network tomography of binary network performance characteristics. *IEEE Transactions on Information Theory*, 52(12):5373–5388, December 2006.

[12] J. Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1, 1971.

[13] J. D. Esary, F. Proschan, and D. W. Walkup. Association of random variables with applications. *The Annals of Mathematical Statistics*, 38(5), October 1967.

[14] G. H. Golub and C. F. Van-Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore and London, 1996.

[15] A. Gopalan and S. Ramasubramanian. On identifying additive link metrics using linearly independent cycles and paths. *IEEE/ACM Transactions on Networking*,

Figure 7: SprintlinkUS: varying false alarm bound B ($n = 2000$, $\tau = 0.9$, $\tau_{\mathbf{min}} = 0.8$, $\sigma = 0.01$)

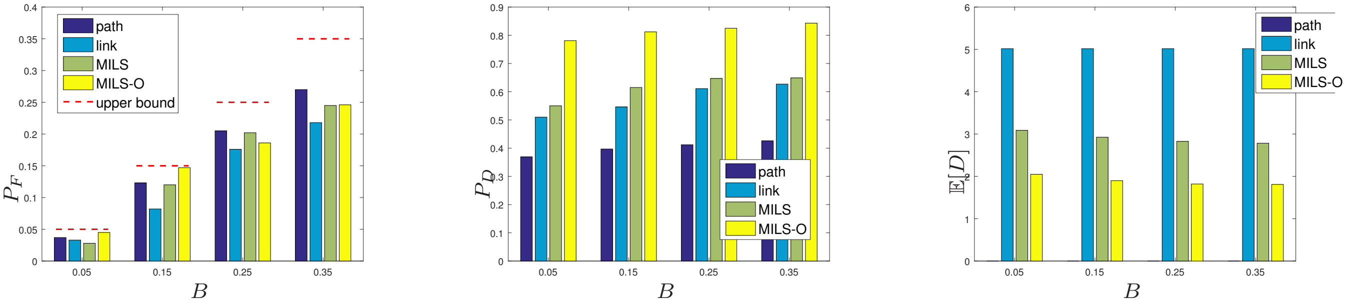

(a) false alarm probability (b) detection probability (c) expected delay

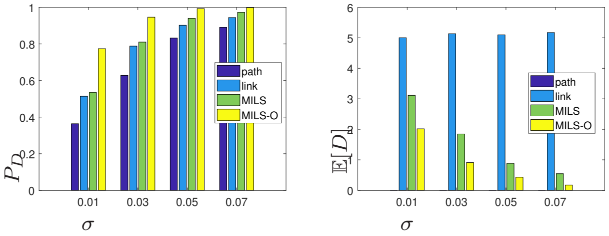

(a) detection probability (b) expected delay

Figure 8: SprintlinkUS: varying ratio of abnormal links σ ($n = 2000$, $B = 0.1$, $\tau = 0.9$, $\tau_{\mathbf{min}} = 0.8$)

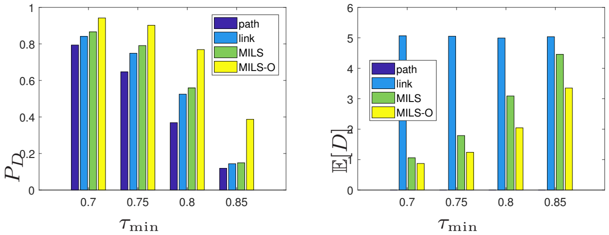

(a) detection probability (b) expected delay

Figure 9: SprintlinkUS: varying minimum link success probability $\tau_{\mathbf{min}}$ ($n = 2000$, $B = 0.1$, $\tau = 0.9$, $\sigma = 0.01$)

20(3), June 2012.

[16] Y. Gu, G. Jiang, V. Singh, and Y. Zhang. Optimal probing for unicast network delay tomography. In *IEEE INFOCOM*, 2010.

[17] O. Gurewitz and M. Sidi. Estimating one-way delays from cyclic-path delay measurements. In *IEEE INFOCOM*, 2001.

[18] T. He, C. Liu, A. Swami, D. Towsley, T. Salonidis, A. Bejan, and P. Yu. Fisher information-based experiment design for network tomography. In *ACM SIGMETRICS*, 2015.

[19] J. D. Horton and A. LÃşpez-Ortiz. On the number of distributed measurement points for network tomography. In *ACM SIGCOMM conference on Internet measurement*, 2003.

[20] R. Kompella, J. Yates, A. G. Greenberg, and A. C. Snoeren. Detection and localization of network black holes. In *IEEE INFOCOM*, 2007.

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

[22] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley. Inferring link metrics from end-to-end path measurements: Identifiability and monitor placement. *IEEE/ACM Transactions on Networking*, 22(4):1351–1368, June 2014.

[23] L. Ma, T. He, K. K. Leung, D. Towsley, and A. Swami. Efficient identification of additive link metrics via network tomography. In *IEEE ICDCS*, 2013.

[24] L. Ma, T. He, A. Swami, D. Towsley, and K. Leung. On optimal monitor placement for localizing node failures via network tomography. *Elsevier Performance Evaluation*, 91:16–37, September 2015.

[25] L. Ma, T. He, A. Swami, D. Towsley, K. Leung, and J. Lowe. Node failure localization via network tomography. In *ACM IMC*, 2014.

[26] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot. Characterization of failures in an IP backbone. In *IEEE INFOCOM*, 2004.

[27] H. X. Nguyen and P. Thiran. The boolean solution to the congested IP link location problem: Theory and practice. In *IEEE INFOCOM*, 2007.

[28] V. N. Padmanabhan, L. Qiu, and H. Wang. Server-based inference of internet link lossiness. In *IEEE INFOCOM*, April 2003.

[29] D. Pollard. *Convergence of Stochastic Processes.* Springer-Verlag, 1984.

[30] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *ACM SIGCOMM*, August 2002.

[31] P. Tan and C. Drossos. Invariance properties of maximum likelihood estimators. *Mathematics Magazine*, 48(1), January 1975.

[32] H. L. V. Trees. *Detection, estimation, and modulation theory.* John Wiley&Sons, 2004.

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

[34] B. Xi, G. Michailidis, and V. Nair. Estimating network loss rates using active tomography. *Journal of the American Statistical Association*, 101(476):1430–1448, December 2006.

[35] H. Zeng, P. Kazemian, G. Varghese, and N. McKeown. Automatic test packet generation. In *ACM CoNEXT*, 2012.

[36] Y. Zhao, Y. Chen, and D. Bindel. Towards unbiased end-to-end network diagnosis. In *ACM SIGCOMM*, 2006.

APPENDIX

Proof of Lemma 4.1: We prove the result by analyzing the algorithm Seek_MILS [36]. For each $p \in P_v$, Seek_MILS examines all the subpaths, and considers a subpath s as a MILS if there is no identifiable subpath with the same starting point as s that is shorter than s. Consider a path $p = (v_1, \ldots, v_m)$ (as a node sequence). Let $s_1 = (v_1, \ldots, v_{m_1})$ be a MILS on p starting from v_1. If $m_1 < m$, then the subpath $(v_{m_1}, \ldots, v_m)$ must be identifiable (since s_1 is identifiable), and must contain a MILS $s_2 = (v_{m_1}, \ldots, v_{m_2})$. Repeating this argument will yield a set of disjoint MILSs $s_1, \ldots, s_k$ ($k \geq 1$) such that their concatenation gives p. Therefore, $\alpha_p = \prod_{i=1}^{k} \zeta_{s_i}$. □

Proof of Lemma 4.2: We prove the claim using the invariance property of MLE [31]. Since the MILS success probabilities $\boldsymbol{\zeta}_v$ are identifiable from the path success probabilities $\boldsymbol{\alpha}_v$ by definition, and the opposite also holds by Lemma 4.1, $\boldsymbol{\zeta}_v$ (and thus $\log \boldsymbol{\zeta}_v$) and $\boldsymbol{\alpha}_v$ form a bijection. Moreover, since the MLE of the mean of a Bernoulli random variable is its empirical mean, we know that the MLE of $\boldsymbol{\alpha}_v$ is $\hat{\boldsymbol{\alpha}}_v$. By the invariance property of MLE, applying the transformation (11) to $\hat{\boldsymbol{\alpha}}_v$ gives the MLE of $\log \zeta_s$. □

Proof of Theorem 4.3: Let $\boldsymbol{\zeta}_{-s}$ denote the success probabilities of all the MILSs in $M_v \setminus \{s\}$. Suppose that we fix the success probabilities for the MILSs in $M_v \setminus \{s\}$, but vary the success probability of MILS s from ζ_s to ζ'_s. Accordingly, the path success probabilities vary from α_p to α'_p ($p \in P_v$). For any $\zeta'_s \neq \zeta_s$, the likelihood ratio $L(\mathbf{x}_v;\, \zeta'_s, \zeta_s | \boldsymbol{\zeta}_{-s}) := f(\mathbf{x}_v;\, \zeta'_s, \boldsymbol{\zeta}_{-s}) / f(\mathbf{x}_v;\, \zeta_s, \boldsymbol{\zeta}_{-s})$ equals

$$\begin{aligned} L(\mathbf{x}_v;\, \zeta'_s, \zeta_s | \boldsymbol{\zeta}_{-s}) &= \frac{\prod_{p: s \in p} {\alpha'_p}^{\sum_{i=1}^n x_{p,i}} (1-\alpha'_p)^{n - \sum_{i=1}^n x_{p,i}}}{\prod_{p: s \in p} \alpha_p^{\sum_{i=1}^n x_{p,i}} (1-\alpha_p)^{n - \sum_{i=1}^n x_{p,i}}} \\ &= \left[\prod_{p: s \in p} \left(\frac{1-\alpha'_p}{1-\alpha_p} \right)^n \right] \\ &\quad \cdot \exp\left[\sum_{p: s \in p} \left(\sum_{i=1}^{n} x_{p,i} \right) \log \frac{\alpha'_p (1-\alpha_p)}{\alpha_p (1-\alpha'_p)} \right], \end{aligned}$$

where $s \in p$ means that MILS s is a subpath of path p. Note that the likelihood function $f(\mathbf{x}_v;\, \zeta_s, \boldsymbol{\zeta}_{-s})$ is well-defined as the path success probabilities $\boldsymbol{\alpha}_v$ are uniquely determined by ζ_s and $\boldsymbol{\zeta}_{-s}$ according to Lemma 4.1.

By the Neyman-Pearson lemma [4], the optimal detector for a given pair of parameters (ζ_s, ζ'_s) satisfying $\zeta_s \geq \tau^{|s|}$ and $\zeta'_s < \tau^{|s|}$ must have the form

$$\delta(\mathbf{x}_v) = \begin{cases} 1 & \text{if } L(\mathbf{x}_v;\, \zeta'_s, \zeta_s | \boldsymbol{\zeta}_{-s}) > t, \\ \gamma & \text{if } L(\mathbf{x}_v;\, \zeta'_s, \zeta_s | \boldsymbol{\zeta}_{-s}) = t, \\ 0 & \text{if } L(\mathbf{x}_v;\, \zeta'_s, \zeta_s | \boldsymbol{\zeta}_{-s}) < t, \end{cases}$$

which is equivalent to comparing the statistic $T(\mathbf{x}_v) := \sum_{p: s \in p} \left(\sum_{i=1}^n x_{p,i} \right) \log \frac{\alpha'_p (1-\alpha_p)}{\alpha_p (1-\alpha'_p)}$ with a threshold t'.

Since $T(\mathbf{x}_v)$ is a function of α_p and α'_p ($p \in P_v$ such that $s \in p$), which in turn depend on ζ_s and ζ'_s, different values of these parameters require different testing statistics. Thus, no single detector is uniformly optimal for all $\zeta_s \geq \tau^{|s|}$ and $\zeta'_s < \tau^{|s|}$. □

Proof of Lemma 4.4: Given $\log \zeta_s = \sum_{p \in P_v} c_{s,p} \log \alpha_p$, we have

$$\Pr\{\log \hat{\zeta}_s < k\} = \Pr\{\sum_{p \in P_v} c_{s,p} \log \hat{\alpha}_p < k\}$$

$$\leq \Pr\{\exists p \in P_v :\, c_{s,p} \log \hat{\alpha}_p < \frac{k}{|P_v|}\} \tag{26}$$

$$\leq \sum_{p \in P_v} \Pr\{c_{s,p} \log \hat{\alpha}_p < \frac{k}{|P_v|}\}, \tag{27}$$

where (26) is by relaxing $\sum_{p \in P_v} c_{s,p} \log \hat{\alpha}_p < k$ to a necessary condition, and (27) is by the union bound. If $c_{s,p} \leq 0$, then $c_{s,p} \log \hat{\alpha}_p \geq 0$, and it is impossible to have $c_{s,p} \log \hat{\alpha}_p < \frac{k}{|P_v|}$ for $k < 0$. If $c_{s,p} > 0$, then

$$\begin{aligned} \Pr\{c_{s,p} \log \hat{\alpha}_p < \frac{k}{|P_v|}\} &= \Pr\{n\hat{\alpha}_p < n e^{\frac{k}{|P_v| c_{s,p}}}\} \\ &= F_B(\lceil n e^{\frac{k}{|P_v| c_{s,p}}} \rceil - 1;\, n, \tau^{|p|}), \end{aligned} \tag{28}$$

as $n\hat{\alpha}_p$ is a binomial random variable with parameters $(n, \tau^{|p|})$. Plugging the above results into (27) yields the bound. □

Proof of Lemma 4.5: The first bound is simply the union bound. To show the second bound, note that (all probabilities are under H_0)

$$\begin{aligned} 1 - P_F(M) &= \Pr\{\log \hat{\zeta}_s \geq t_s,\, \forall s \in M\} \\ &= \Pr\{\sum_{p \in P} c_{s,p} \log \hat{\alpha}_p \geq t_s,\, \forall s \in M\}, \end{aligned} \tag{29}$$

where t_s is the detection threshold for MILS s. Since $\hat{\alpha}_p$'s ($p \in P$) are independent and hence *associated* (by Theorem 2.1 in [13]), and $\sum_{p \in P} c_{s,p} \log \hat{\alpha}_p$ is nondecreasing in $\hat{\alpha}_p$ for $c_{s,p} \geq 0$, by Theorem 5.1 in [13],

$$\begin{aligned} \Pr\{\sum_{p \in P} c_{s,p} \log \hat{\alpha}_p \geq t_s, \forall s \in M\} &\geq \prod_{s \in M} \Pr\{\sum_{p \in P} c_{s,p} \log \hat{\alpha}_p \geq t_s\} \\ &= \prod_{s \in M} (1 - P_F(s)). \end{aligned} \tag{30}$$

Combining (29, 30) gives the result. □

Proof of Theorem 4.6: The result is implied by Lemmas 4.4 and 4.5. For each MILS $s \in M$, the setting of γ and t guarantees that $P_F(s) \leq \beta$ by Lemma 4.4. If $c_{s,p} \geq 0$ for all $s \in M$ and $p \in P$, then by Lemma 4.5 and (19),

$$P_F(M) \leq 1 - \prod_{s \in M} (1 - P_F(s)) \leq 1 - (1-\beta)^{|M|} = B. \tag{31}$$

Otherwise, $P_F(M) \leq \sum_{s \in M} P_F(s) \leq \beta |M| = B$. □

Proof of Lemma 4.7: Since $\log \zeta_s = \sum_{p \in P} c_{s,p} \log \alpha_p \leq t_s - \epsilon$, we have (all probabilities are conditioned on $\boldsymbol{\theta}$)

$$P_M(s|\boldsymbol{\theta}) = \Pr\{\sum_{p \in P} c_{s,p} \log \hat{\alpha}_p \geq t_s\}$$

$$\leq \Pr\{\sum_{p \in P} c_{s,p} (\log \hat{\alpha}_p - \log \alpha_p) \geq \epsilon\}$$

$$\leq \Pr\left\{ \exists p \in P_s,\, c_{s,p} \log \frac{\hat{\alpha}_p}{\alpha_p} \geq \frac{\epsilon}{|P_s|} \right\} \tag{32}$$

$$\leq \sum_{p \in P_s} \Pr\left\{ c_{s,p} \log \frac{\hat{\alpha}_p}{\alpha_p} \geq \frac{\epsilon}{|P_s|} \right\} \tag{33}$$

$$\begin{aligned} &= \sum_{p \in P_s : c_{s,p} > 0} \Pr\left\{ \frac{\hat{\alpha}_p}{\alpha_p} \geq \exp\left(\frac{\epsilon}{c_{s,p} |P_s|} \right) \right\} \\ &\quad + \sum_{p \in P_s : c_{s,p} < 0} \Pr\left\{ \frac{\hat{\alpha}_p}{\alpha_p} \leq \exp\left(\frac{\epsilon}{c_{s,p} |P_s|} \right) \right\}, \end{aligned} \tag{34}$$

where (32) is by relaxing $\sum_{p\in P} c_{s,p}(\log\hat{\alpha}_p - \log\alpha_p) \ge \epsilon$ to a necessary condition, and (33) is by the union bound. We now bound (34) by the Chernoff-Hoeffding bound.

Chernoff-Hoeffding bound [29]: Let $X_1, \ldots, X_n$ be random variables with a common range $[0, 1]$ such that $\mathbb{E}[X_t|X_1, \ldots, X_{t-1}] = \mu$, $\forall 1 \le t \le n$. Let $S_n = \sum_{i=1}^n X_i$. For $a \ge 0$,

$$\Pr\{S_n \ge n\mu + a\} \le e^{-2a^2/n}, \quad \Pr\{S_n \le n\mu - a\} \le e^{-2a^2/n}.$$

Since $n\hat{\alpha}_p = \sum_{i=1}^n x_{p,i}$ for i.i.d. Bernoulli random variables $x_{p,1}, \ldots, x_{p,n}$, with $\mathbb{E}[x_{p,i}] = \alpha_p$ for all $1 \le i \le n$, the Chernoff-Hoeffding bound implies that if $c_{s,p} > 0$,

$$\begin{aligned}\Pr\left\{\frac{\hat{\alpha}_p}{\alpha_p} \ge \exp\left(\frac{\epsilon}{c_{s,p}|P_s|}\right)\right\} &\le \exp\left[-2n\alpha_p^2\left(\exp\left(\frac{\epsilon}{c_{s,p}|P_s|}\right) - 1\right)^2\right]\\ &\le \exp\left[-2n\alpha_s^2\left(\frac{\exp\left(\frac{\epsilon}{c_s^{\max}|P_s|}\right) - 1}{\exp\left(\frac{\epsilon}{c_s^{\min}|P_s|}\right)}\right)^2\right]. \quad (35)\end{aligned}$$

Similarly, if $c_{s,p} < 0$,

$$\begin{aligned}\Pr\left\{\frac{\hat{\alpha}_p}{\alpha_p} \le \exp\left(\frac{\epsilon}{c_{s,p}|P_s|}\right)\right\} &\le \exp\left[-2n\alpha_p^2\left(1 - \exp\left(\frac{\epsilon}{c_{s,p}|P_s|}\right)\right)^2\right]\\ &\le \exp\left[-2n\alpha_s^2\left(\frac{\exp\left(\frac{\epsilon}{c_s^{\max}|P_s|}\right) - 1}{\exp\left(\frac{\epsilon}{c_s^{\min}|P_s|}\right)}\right)^2\right]. \quad (36)\end{aligned}$$

Plugging (35, 36) into (34) yields the result. □

Proof of Theorem 4.8: Let $\delta_s \in \{0, 1\}$ denote the decision for MILS s. Then

$$\begin{aligned}P_D(M|\boldsymbol{\theta}) &= \Pr\{\exists s \in M,\ \delta_s = 1|\boldsymbol{\theta}\}\\ &\ge \max_{s\in M} \Pr\{\delta_s = 1|\boldsymbol{\theta}\}\\ &= 1 - \min_{s\in M} \Pr\{\delta_s = 0|\boldsymbol{\theta}\}. \quad (37)\end{aligned}$$

Note that $\Pr\{\delta_s = 0|\boldsymbol{\theta}\} = P_M(s|\boldsymbol{\theta}) \le \eta_s(\boldsymbol{\theta})$. Plugging this bound into (37) yields the result. □

Proof of Theorem 4.9: Fix a realization of $\mathbf{M}$ (which also specifies a realization of T). Since for $d = 0, \ldots, T-1$, $D(\mathbf{M}) > d$ if and only if no anomaly is detected among the MILSs in M_d, we have (all probabilities and expectations are conditioned on $\mathbf{M}$ and $\boldsymbol{\theta}$)

$$\begin{aligned}\Pr\{D(\mathbf{M}) > d\} &= \Pr\{\log\hat{\zeta}_s \ge t_s,\ \forall s \in M_d\}\\ &\le \min_{s\in M_d} \Pr\{\log\hat{\zeta}_s \ge t_s\} = \min_{s\in M_d} P_M(s|\boldsymbol{\theta}), \quad (38)\end{aligned}$$

By Lemma 4.7, (38) is bounded by $\min_{s\in M_d} \eta_s(\boldsymbol{\theta})$. Then

$$\begin{aligned}\mathbb{E}[D(\mathbf{M})] &= \sum_{i=0}^{T}\sum_{d=0}^{i-1} \Pr\{D(\mathbf{M}) = i\}\\ &= \sum_{d=0}^{T-1} \Pr\{D(\mathbf{M}) > d\} \le \sum_{d=0}^{T-1} \min_{s\in M_d} \eta_s(\boldsymbol{\theta}). \quad (39)\end{aligned}$$

Taking the expectation of (39) over $\mathbf{M}$ yields the bound. □

Table 2: Properties of the Simulated Networks

Network	$\|N\|$	$\|L\|$	$\|Z\|$	rank($\mathbf{A}$)	$d_{\mathcal{T}}$	$d_{\mathcal{T}^*}$
Exodus	22	51	7	15	3	2
EUROPEgraph	28	66	9	26	3	2
Abovenet	22	80	7	19	2	2
SprintlinkUS	44	106	14	39	4	3
TiscaliEurope	51	129	16	54	3	2

Proof of Corollary 5.2: We prove a stronger result: for a generic matroid $\mathcal{M} = (E, \mathcal{I})$ and a generic weight function $c:\ E \to \mathbb{R}^+$, the greedy algorithm is optimal in finding a basis B of $\mathcal{M}$ that minimizes $\max_{e\in B} c(e)$.

Let $G = \{g_1, \ldots, g_k\}$ with $c(g_1) \le \ldots \le c(g_k)$ be the set returned by the greedy algorithm, and $H = \{h_1, \ldots, h_l\}$ with $c(h_1) \le \ldots \le c(h_l)$ be a basis such that $c(h_l) < c(g_k)$. First, G must be a basis (and hence $k = l$), as the greedy algorithm must be able to build a basis (since it examines everything in E) and will stop adding elements after selecting a basis. Moreover, if i is the smallest index such that $c(h_i) < c(g_i)$, then for sets $G_{i-1} = \{g_1, \ldots, g_{i-1}\}$ and $H_i = \{h_1, \ldots, h_i\}$, there must exist an element $h_j \in H_i \setminus G_{i-1}$ such that $G_{i-1} \cup \{h_j\} \in \mathcal{I}$ due to the augmentation property of matroid. We must have $c(h_j) \le c(h_i) < c(g_i)$. This is a contradiction since the greedy algorithm must have examined h_j before g_i and determined that $G_{i-1} \cup \{h_j\} \notin \mathcal{I}$. Therefore, the greedy algorithm must return a basis of $\mathcal{M}$ that minimizes the maximum per-element weight. □

Evaluations in a Rank-deficient Setting: We repeat the evaluations in Section 6 in a new setting, where we only use the first $\lceil 0.3|N|\rceil$ nodes selected by the procedure in Section 6.1 as monitors. As shown in Table 2, all the measurement matrices are rank-deficient. We repeat the evaluation in Fig. 5, with link-based detection ('link') replaced by centralized detection based on MILSs ('centralized'). The results in Fig. 10 show that our scheme still outperforms path-based detection in detection probability and centralized detection in detection delay in most cases. However, the absolute value of the detection probability is much lower than that in Fig. 5 (b), as the abnormal links are not always identifiable (or even covered) by the probing paths.

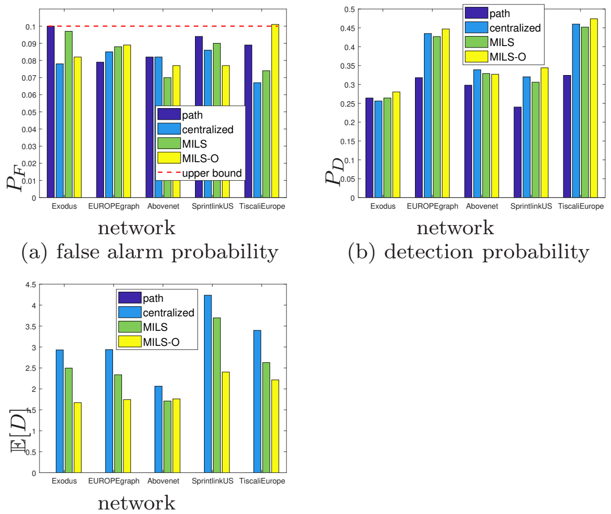

(a) false alarm probability

(b) detection probability

(c) expected delay

Figure 10: Rank-deficient setting (same as Fig. 5)