Evaluating Network Boolean Tomography under Byzantine Attacks

1st Haotian Deng
School of Cyberspace Security
Beijing University of Posts and Telecommunications
Beijing, P. R. China 100876
haotian_deng@bupt.edu.cn

2nd Shengli Pan*
School of Cyberspace Security
Beijing University of Posts and Telecommunications
Beijing, P. R. China 100876
psl@bupt.edu.cn

***Abstract*—It is vital to closely track the operation statuses of network-internal links. Accurate knowledge of the operation statuses of network-internal links is vital for the management of many networks like the Internet, the satellite communication network, etc. Network boolean tomography can identify congested links just using end-to-end path status observations, and is able to work efficiently even without any available cooperation of internal nodes. Nevertheless, it heavily assumes that all the path status observations collected are true while some Byzantine attacks, e.g., the label flip attacks, could violate this assumption. In this paper, we present a performance evaluation of network boolean tomography under Byzantine attacks. Our results against various attacking rates, locations, and scales all show that Byzantine attacks could cause a significant performance degradation of network boolean tomography, suggesting a pressing need of developing the detection and countermeasure techniques.**

***Index Terms*—network boolean tomography, congested link identification, end-to-end measurement, Byzantine attacks**

I. Introduction

Network boolean tomography [1] is a potent tool for localizing traffic congestion or jamming in communication networks. It could be well applied to various types of networks, like the Internet, IoT, satellite, and space communication networks, where it helps identify and understand the effects of signal jamming, transmission interference [2], and bandwidth consumption attacks, etc. Note that though the structure of the satellite network is rather dynamic, its routing topology is technically virtualized as static, making netowrk boolean tomography also an appealing tool for monitoring network performance of satellite networks. To evaluate internal communication jamming is a fundamental aspect of network management. It enables network administrators to quickly detect and troubleshoot issues, plan for future capacity needs, harden network security, optimize costs associated with network usage, and so on. However, almost all of today's communication networks are known vulnerable to a wide range of attacks [3], [4], including Byzantine attacks [5], which can significantly impact the performance of network boolean tomography.

Network boolean tomography first quantifies and gathers the binary status of paths (i.e., "good" or "bad") by comparing their performance observations to a predefined threshold, e.g. in Fig. 1(a), the observed packet loss ratio of a path greater than 1% will indicate that it gets a bad or congestion status; otherwise, its status is good [6]. These observed path statuses then are reasoned by the Bayesian framework of Maximum A Posteriori (MAP) attribution to identify congested links. However, the MAP problem here is normally NP-hard. To circumvent the NP-hardness, [7] proposed the "CLINK" algorithm for a greedy identification, while the "SCFS" algorithm of [6] chose to remove any prerequisite of links' prior congestion probabilities and simplified it as a Maximum Likelihood Estimation (MLE) problem. Besides the NP-hardness, there were also works to address other issues like probing [8], scalability [9], [10], dynamic routing, identifiability [11], sparsity [12], and so on.

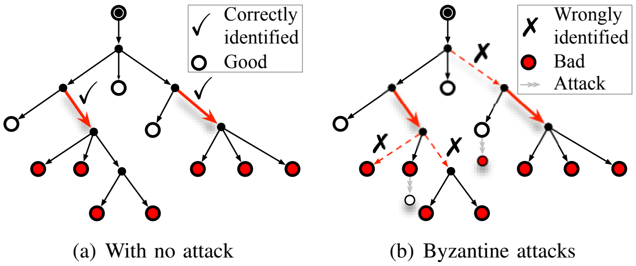

Fig. 1. Illustration of network boolean tomography [6] with and without attacks. There are two congested links depicted in red and bold solid arrows. In (a), both congested links are correctly identified while in (b), not only neither of them are detected, but three other links are even wrongly identified.

Nonetheless, most of the existing works on network boolean tomography assume a benign scenario. More specifically, they tacitly require all the measurements obtained to be dependable. This becomes increasingly demanding in today's Internet, in the context of the growing number and sophistication of both cyber threats and attacks [3]. One commonly-discussed threat would be Byzantine attacks [5], where the compromised participants will not always behave in accordance with the measurement protocol [13], but can report their results dishonestly during the measurement collection procedure. As illustrated in Fig. 1(b), after the third edge node (i.e., the one with its congested/bad status changed to "good") from the left dishonestly tells its path status observation as "good" instead of "bad", the truly congested link on the left side is

* Corresponding author: Shengli Pan was also with the Key Laboratory of Trustworthy Distributed Computing and Service (BUPT), Ministry of Education, P. R. China.

no longer revealed as depicted in Fig. 1(a). While to make matters worse, two more "good" links will be attributed as "congested"/"bad" instead [6]. Such a misidentification also occurs when a "good" path is reported as "bad" by the fourth edge node (i.e., the one with its good status changed to "bad") from the right in Fig. 1(b).

In this work, we attempt to evaluate network boolean tomography in networks, where there are Byzantine attacks. By investigating identification performance degradation across various setups, we find that Byzantine attacks is quite effective even if it is launched at a slow rate with just a single compromised edge node. We also demonstrate that Byzantine attacks launched neither at different rates nor at different edge nodes could be rewarded much differently, indicating a possibility to explore the optimal strategy of Byzantine attack. Our evaluation contributions are summarized as follows:

- ○ The effectiveness of Byzantine attacks is validated by evaluating their impact on network congestion diagnosis performance under various scenarios, including different network topologies, diagnostic methods, attack frequencies, and attack locations.
- ○ It is found that Byzantine attacks could cause different performance degradation to different types of links, while links' congestion probabilities, the location, and the number of paths attacked could also make a difference.
- ○ We explore the key factors like the attacking frequency on the effectiveness of Byzantine attacks, revealing valuable insights into how attackers could possibly optimize their strategies in terms of a limited attack capability.

II. Preliminaries

It is noticed that conventional network tomography approaches, which estimate links' performance parameters (e.g., link latency and link loss ratios) using end-to-end path measurements [1], could also be employed to accurately diagnose links' congestion statuses. However, they in general work at a high cost of both measurement operation complexity and computational complexity, thus greatly rendering the need for network boolean tomography for a better balance between measurement costs and estimation errors. Nevertheless, security threats like Byzantine attacks that could falsify end-to-end measurement results of path statuses, will also no doubt make network boolean tomography a non-trivial task [14]–[16]. In what follows, we introduce network boolean tomography and accordingly define the Byzantine threat model.

A. Boolean Models for Congested Link Identification

The intervening network is modeled as a directed acyclic graph $\mathcal{G}$, where there are m end-to-end paths and n interior links. In terms of measurement coverage, it is assumed that the degree of each node in $\mathcal{G}$ is not equal to two. Specifically, every end node gets a degree of one while the degree of every interior node is greater than two. In each time slot, we first probe a path by sending and receiving a sufficient number of packets at both its end nodes, and then we quantify and obtain a measurement of its boolean status $y \in \{0, 1\}$ through comparing the observed loss ratio to a given threshold, e.g., if the loss ratio is larger than 1%, the path status is congested (i.e., $y = 1$); otherwise it is good (i.e., $y = 0$). All the obtained statuses in the t-th time slot are collected in a vector $\mathbf{y}_t = [y_1, y_2, \cdots, y_m]$, while the unknown staties of all the n interior links are denoted by $\mathbf{x}_t = [x_1, x_2, \cdots, x_n]$.

At each time slot, a path is observed as congested when it traverses at least a single congested link. It is good while all the links it routes through are good. We model this logical reasoning relationship with the following boolean expression,

$$y_i = x_j \cup x_k \cup \cdots \cup x_l, \tag{1}$$

where "$\cup$" denotes the bool operation of "OR" and only those links on the given path are counted in (1). After probing for T total time slots, we will collect T status observations for each path. To express both dimensions of paths and time, we further use the following matrix form,

$$\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_T].$$

Similarly, all the T unknown link status vectors are also expressed as

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_T].$$

We introduce $\mathbf{R}_{m\times n}$ to express the routing relationships between each of the m paths and n link, i.e., $\mathbf{R}(i, j) = 1$ denotes that the i-th path routes through the j-th link; otherwise, $\mathbf{R}(i, j) = 0$. With the information of $\mathbf{R}$, we build the relationship between the observed path statuses and the unknown link statuses over T time slots by

$$\mathbf{Y} = \mathbf{R} \bigcup \mathbf{X}, \tag{2}$$

where the observed status of each path at each time slot is related to its routed links via (1).

However, it is non-trivial to solve (2). Because a network $\mathcal{G}$ normally gets more links than paths, making solving (2) is an ill-posed problem. What's worse, solving (2) often falls into the "integer programming problem" as only boolean status is presented, further resulting in the "NP-hardness". For instance, given the prior congestion probability of each link, solving (2) could be formulated as the following maximum a posteriori (MAP) estimation problem,

$$\widehat{\mathbf{X}} = \arg\max_{\mathbf{X}} p(\mathbf{Y}|\mathbf{X}), \tag{3}$$

where $\mathbf{X}$ and $\mathbf{Y}$ are subject to each other according to (2). Such a MAP problem is equivalent to a maximum vertex cover problem, which is a well-known NP-hard problem. To this end, both "CLINK" [7] and "SCFS" [6] algorithms are proposed. The former solves it with a greedy policy, while the latter assumes no prior knowledge of each link and reduces it as a regular optimization problem for tree networks like Fig. 1.

Though existing literature demonstrates that network boolean tomography could perform quite well for congested link identification, they need to work in a benign network scenario, specifically for implying that all the collected path observations truly reveal their statuses. In the following, we will introduce the Byzantine threat model and argue that such an implicit assumption will not always hold true.

B. Byzantine Threat Model

Among the various types of Byzantine attacks, the label flip attack has emerged as a particularly insidious tactic, where a malicious node modifies the labels of the data being transmitted, causing other nodes within the network to misinterpret the data. The label flip attack assumes that the malicious node has control over the labels of the data packets. For example, the attacker might have exploited vulnerabilities in IoT hardware devices or application software bugs to take control of a node and manipulate the labels of data packets. In this paper, we assume the Byzantine end nodes have the capability to misrepresent the obtained path statuses during the collection of end-to-end path status observations in network boolean tomography. Specifically, a good path status might be reported as "congested" while a "congested" one could be told as "good" to the network monitoring center, changing (1) to

$$y_i' = (x_j \cup x_k \cup \cdots \cup x_l) \oplus a_i, \;\; a_i \in \{0, 1\} \qquad (4)$$

where "$\oplus$" denotes the logical operation of "XOR" and $a_i = 1$ indicates a label flip attack to the status of the i-th path.

As we will show in the next section, network boolean tomography will misperform in the presence of falsified path status observations. While detecting and developing countermeasures for Byzantine label flip attacks is a challenging task and requires ongoing research, we leave the development of detection and potential countermeasure techniques for future work. Nonetheless, a thorough performance evaluation of network boolean tomography under such Byzantine attacks is crucial to develop robust security measures and ensure the practicability of network tomography methodology.

III. Evaluations

We present the results of our evaluation of network boolean tomography under Byzantine attacks, and our evaluation objective is two-fold: the effectiveness of Byzantine attacks and the optimum for its attack strategy. Specifically, we first investigate the performance of network boolean tomography for the network depicted in Fig. 1, where there is a malicious node that attempts to disrupt the accuracy of the boolean identification process; Then, we use real network topologies [17] to explore the optimum for various attack strategies. To make our results reproducible, all of the simulation setups and data in this paper are available as open-source [1].

A. Simulation Setups

In order to gain a deeper insight into the effects of introducing malicious nodes on the boolean identification process, it is crucial to analyze the parameters that can influence the performance of this process. Since its performance can greatly fluctuate depending on the link probability distributions present in the intervening network, we mainly set up two major network scenarios according to the uniform distribution for generating link prior congestion probabilities: the "even-variance" one and the "uneven-variance" one. Specifically, the former would range in $[0.0, 0.1], [0.1, 0.2], \cdots, [0.9, 1.0]$ while the latter is picked from $[0.0, 0.1], [0.0, 0.2], \cdots, [0.0, 1.0]$. Each setup of link prior congestion probabilities is randomly generated for 20 times. We refer to the attack frequency as the proportion of the number of attacks on the network within a specified time frame, the depth of attack locations as the count of links traversed by the attack path, the breadth of attack locations as the number of paths being attacked, and the network's congestion level as the total of congested paths for not being subjected to any attack. We investiagate each network scenario and repeat the attack against each path status observation for 1000 times and 20 times, respectively.

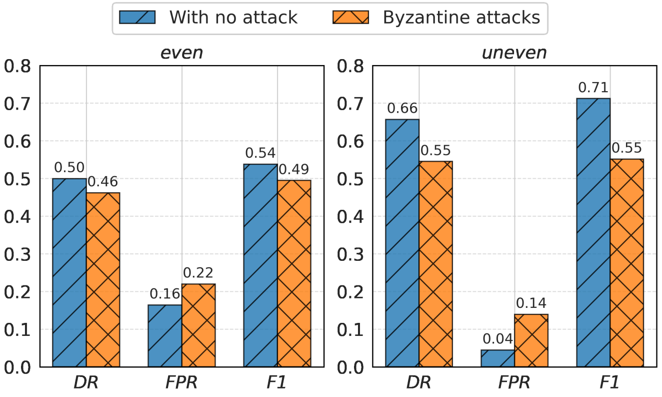

Fig. 2. The identification performances of network boolean tomography with or without Byzantine attacks. "Even" and "uneven" denote the variances of the uniform distribution used for generating link prior congestion probability.

B. Effectiveness of Byzantine Attack

To assess the influence of attacks on diagnostic performance, we conduct attacks with varying frequencies and locations on networks featuring random congestion probabilities, documenting alterations in DR, FPR, and F1-score before and after the attacks. In Fig. 2, it is shown that network boolean tomography does suffer performance decline when there exist attacks. Recalling the illustration of Fig. 1, Byzantine attacks make inaccurate path state information report. This will result in unpredictable diagnostic outcomes and generally lead to a decrease in DR, an increase in FPR, and accordingly a reduction in F1-score.

To further investigate the root cause of performance degradations, we conduct random attacks on networks in different variation distributions and compute the diagnostic performance of different links at various times to observe the recognition situation within the network. We categorize links into three groups based on their location: root links, internal links and edge links. It is found in Fig. 3 that attacks could induce different decline patterns for different link categories. To delve deeper into the effectiveness of attacks on different algorithms, we collect the decrease in diagnostic performance caused by attacks under distinct network scenarios, and ascertain that different algorithms exhibit varying levels of robustness when they are confronted with attacks as demonstrated in Fig. 4.

[1] https://gitee.com/eric-teng/eval_nbt_vs_byzantine

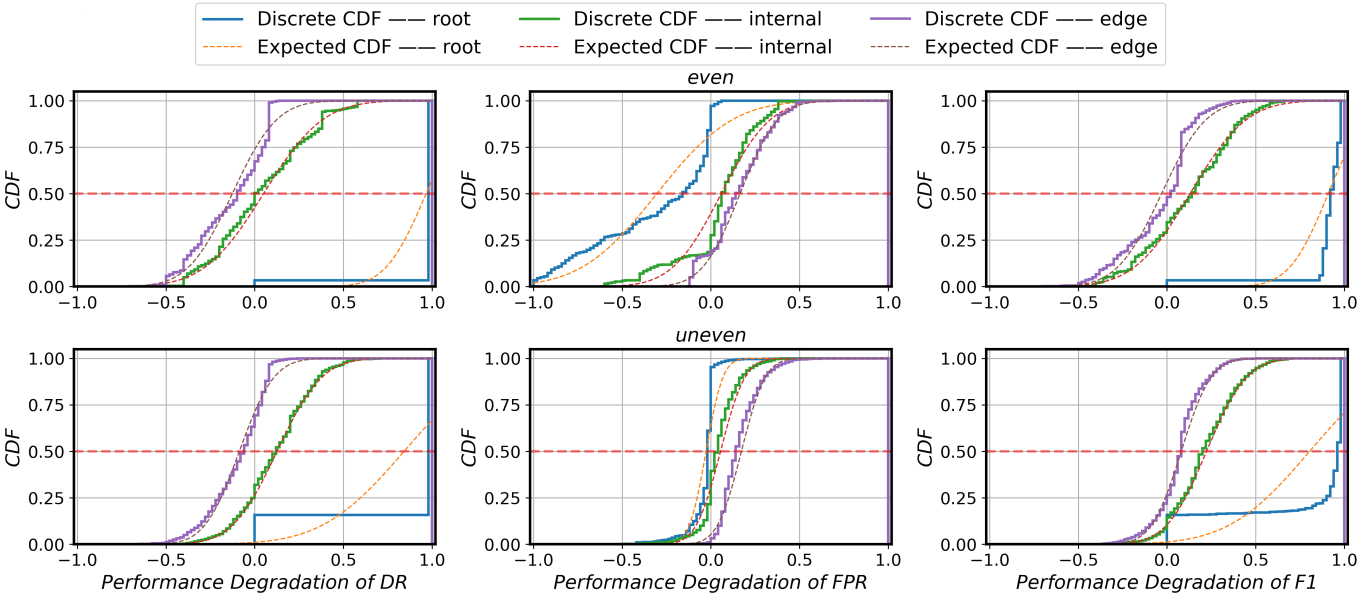

Fig. 3. The CDFs of identification performance degradation of network boolean tomography for different types of congested links against Byzantine attacks. The "expected" curves all refer to the theoretical values that are free of any statistical errors.

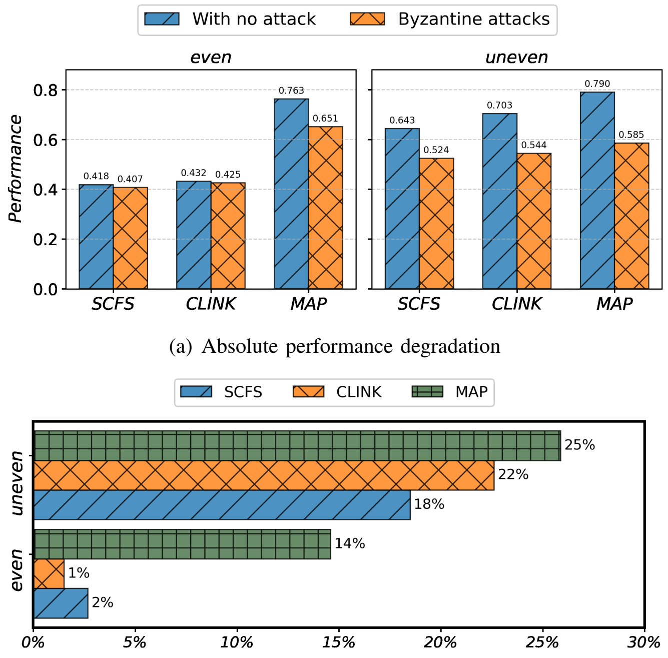

(a) Absolute performance degradation

(b) Relative performance degradation

Fig. 4. The performance degradations of various identification algorithms against Byzantine attacks.

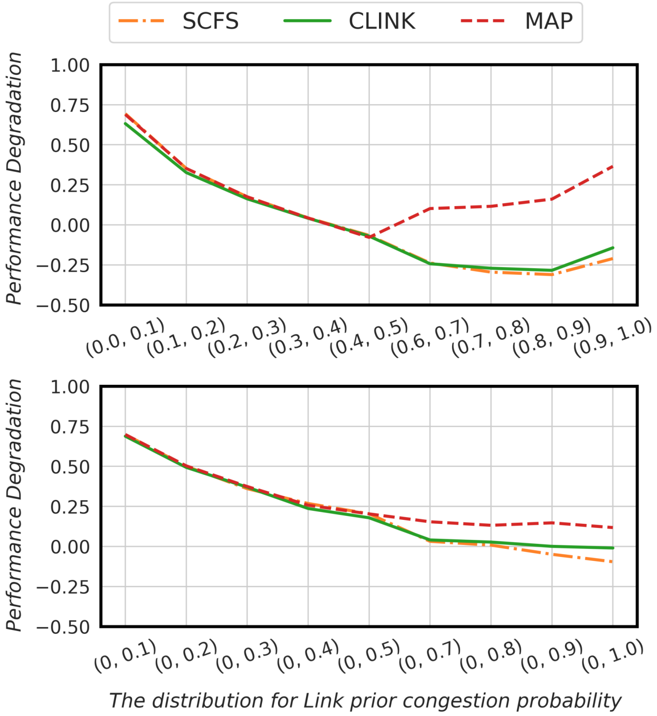

Fig. 5. The performances of various identification algorithms vs different distribution intervals of link prior congestion probability.

In the Fig. 4, where they are drawn with even-variance setups of link prior congestion probability, MAP is the most vulnerable to attacks, particularly in the high-congestion network scenarios; while for uneven-variance ones, it is found SCFS < CLINK < MAP. Similarly, we explore in Fig. 5 the effectiveness of attacks under various congestion conditions. We observe that in even-variance scenarios, the attack effect is more salient when the congestion probability is below 0.4. As the intensity/level of network congestion escalates, the attack effect on MAP grows, while the effect on SCFS and CLINK results in enhanced diagnostic performance. This occurs because SCFS and CLINK believe the maximal bad

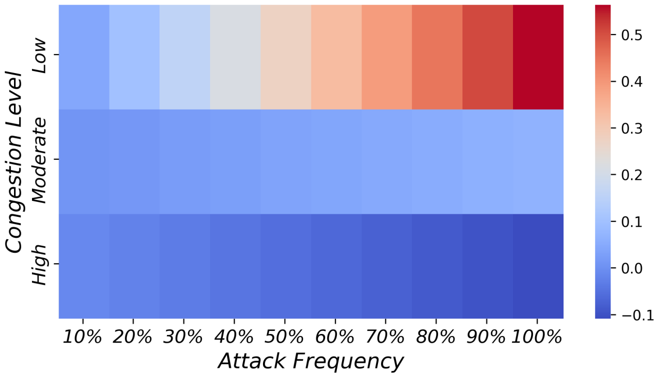

Fig. 6. The performance degradation under various network congestion levels and attack frequencies. The network congestion levels are specified by different uniform (link prior congestion probability) distributions, i.e., the low: $U(0.0, 0.1)$, the Moderate: $U(0.2, 0.3)$, and the high: $U(0.4, 0.5)$.

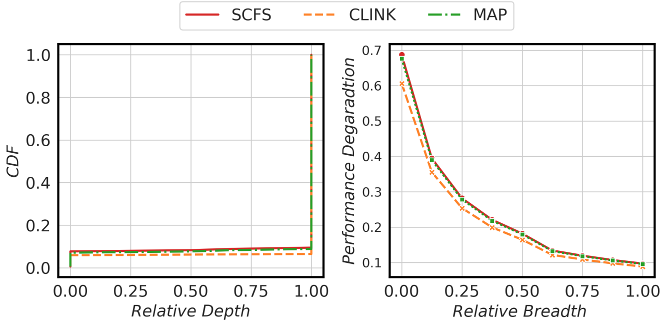

Fig. 7. The performance degradation vs. relative breadth. After normalization, the no. of paths attacked (i.e., the "breadth") becomes the "relative breadth".

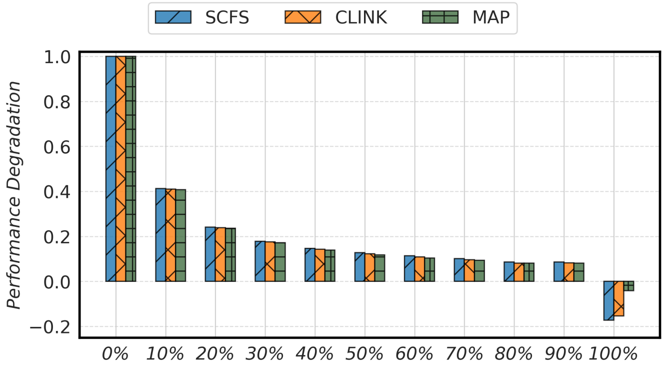

Fig. 8. The identification performance under Byzantine attacks vs. different number of congested paths.

subtree[2] being congested mostly attributes to the root link of it getting congested, while other links within the subtree work well. However, under high congestion level, the greater probability is reverse. Byzantine attacks would break the logic link of the algorithms instead, making the root link more likely be mis-identified as good. And we can observe that in uneven-variance setups, the attack effect diminishes as the link prior congestion probability increases.

The effectiveness of attacks on all links is presented in Fig. 3. It is observed that all types of links suffer a performance decline, while the decline magnitude varies among them. This is because during an attack on the SCFS or MAP identification procedure, the primary links affected are the root link and its child links of the maximal bad subtree formed by the attack path, while other links outside the range will certainly remain unaffected. We leave the theoretical proof in the future.

C. Optimality of Byzantine Attack

To further investigate which key factor can significantly impair the identification performance, we consider the impact of different attack frequencies and locations (depth and breadth) on attack effectiveness using more than 170 real network topologies provided by the Topology Zoo [17]. We employ the setups of fixed attack locations and conduct attacks on the network at varying frequencies and document the performance degradation in Fig. 6. We can observe that in low to moderate congestion, attack effectiveness increases with frequency, especially in low-congestion areas. While in high congestion, such a pattern reverses, showing that more links within the bad subtree would likely to get congested. This will eventually lead both SCFS and CLINK to misattribute end-to-end observations of path congestion.

It reveals that the attack effect amplifies with increased attack frequency, as a higher frequency diminishes the proportion of normal situations, leading to more links being misidentified. The study also demonstrates that internet congestion probability affects the magnitude and direction of the attack effect, but not its overall trend. To explore the impact of attack locations, we measure performance degradations and categorize them into two groups: one to investigate the effects on different paths (i.e., route depth) and the other to analyze the number of paths attacked (i.e., breadth). For discussion purposes, we map all paths from each topology selected from [17] to the 0-1 interval, obtaining the relative depth and breadth, and refer to the congestion probability distribution in the context of low network congestion. Under these conditions, and with a fixed number of attack instances, we execute attacks on various network scale topologies, recording the maximum attack effect based on relative depth. The results in Fig. 7 show that a greater relative depth amplifies the attack effect. This is because the attack is more likely to influence multiple subtrees, increasing the number of affected paths and thereby causing more confusion and errors for the diagnostic algorithm.

We further investigate the breadth aspect by mapping the number of attacked paths to the 0-1 interval. Under low network congestion and with a fixed number of attack instances, we execute attacks on various network scale topologies. The performance degradation under different attack breadths is examined in Fig. 7. The results suggest that the attack effect diminishes with the increase of breadth, likely due to the reduced number of impacted instances under a fixed number of total attacks as breadth increases.

Regarding that the link congestion manifests as the path congestion, it wonders whether the attack effect is closely

[2] If a link ℓ is bad and its parent link is regarded as good, we say that ℓ and all of its descendant links together constitute a "maximal bad subtree".

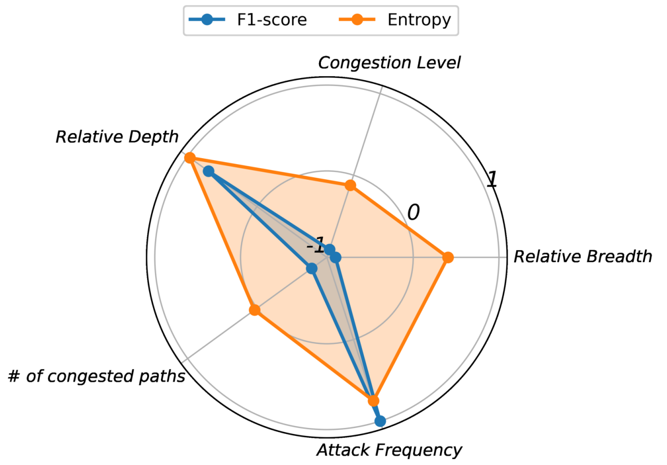

Fig. 9. Various factors for the performance degradation by Byzantine attacks. The contribution of each factor is quantified as the correlation coefficients with either the F1-score or the link congestion entropy.

related to the network's inherent congestion level. We distinguish the congestion level by the proportion of congested paths in the entire network and analyze the attack effects under different congestion levels (Fig. 8). The results there reveal a nuanced relationship between congestion level and attack effectiveness. Specifically, it is observed that the attack effect tends to decrease as the congestion level increases.

Fig. 9 is depicted to gain a comprehensive understanding of how Byzantine attacks perform in various aspects. Specifically it is found that the network congestion level seems to serve as a determining factor. This insight sheds light on the interplay between the attack performance and various factors, providing valuable guidance of both mitigating and designing Byzantine attacks. Moreover, by comparing the link congestion entropy before and after Byzantine attacks, decreased entropy is also observed for each of these factors though with different extents. Clearly, this will suggest a potential employment of an entropy-based detection scheme of Byzantine attacks for network boolean tomography in the future.

IV. Conclusion

Network boolean tomography can efficiently diagnose congested links with fewer end-to-end path measurements, while Byzantine label flip attacks significantly impair performance. These attacks are most effective with high attack frequency or less congested networks, and increased path hops and paths amplify their effectiveness. Additionally, these attacks reduce link congestion entropy, suggesting the potential for an entropy-based attack detection scheme. These evaluation findings could provide valuable insights into the effectiveness of Byzantine attacks on network boolean tomography and might guide network operators to specifically develop their detection and countermeasure strategies. For instance, the game theory of Stackelberg competition could provide a framework for analyzing strategic countermeasures by modeling the interactions between Byzantine attackers and network operators. Meanwhile, recent AI techniques, such as reinforcement learning, might be integrated into the task of detecting Byzantine attacks. Future work will encompass the development of attack schemes in dynamic network scenarios, and the formulation of corresponding attack strategies tailored to address various potential defense mechanisms.

Acknowledgment

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant No. 62271066, the Fundamental Research Funds for the Central Universities (BUPT) under Grant No. 2023RC33, and the High-performance Computing Platform of BUPT.

References

[1] G. Kakkavas, A. Stamou, V. Karyotis, and S. Papavassiliou, "Network tomography for efficient monitoring in sdn-enabled 5g networks and beyond: Challenges and opportunities," *IEEE Communications Magazine*, vol. 59, no. 3, pp. 70–76, 2021.

[2] S. Cho, R. Nithyanand, A. Razaghpanah, and P. Gill, "A churn for the better: Localizing censorship using network-level path churn and network tomography," in *Proceedings of the CoNEXT '17*. New York, NY, USA: ACM, 2017, p. 81–87.

[3] CrowdStrike, "2023 global threat report," https://go.crowdstrike.com/2023-global-threat-report, 2023, accessed on March 15, 2023.

[4] B. Ji, Y. Liu, L. Xing, C. Li, G. Zhang, C. Han, H. Wen, and S. Mumtaz, "Survey of secure communications of internet of things with artificial intelligence," *IEEE Internet of Things Magazine*, vol. 5, no. 3, pp. 92–99, 2022.

[5] S. Li, E. C. H. Ngai, and T. Voigt, "An experimental study of byzantine-robust aggregation schemes in federated learning," *IEEE Transactions on Big Data (Early Access)*, 2023.

[6] N. Duffield, "Network tomography of binary network performance characteristics," *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5373–5388, 2006.

[7] H. X. Nguyen and P. Thiran, "The boolean solution to the congested ip link location problem: Theory and practice," in *The 26th IEEE International Conference on Computer Communications (INFOCOM)*, 2007, pp. 2117–2125.

[8] H. Ikeuchi, H. Saito, and K. Matsuda, "Network tomography based on adaptive measurements in probabilistic routing," in *TheIEEE Conference on Computer Communications (INFOCOM)*, 2022, pp. 2148–2157.

[9] T. He, "Distributed link anomaly detection via partial network tomography," *SIGMETRICS Perform. Eval. Rev.*, vol. 45, no. 3, p. 29–42, 2018.

[10] N. Ogino, T. Kitahara, S. Arakawa, G. Hasegawa, and M. Murata, "Decentralized boolean network tomography based on network partitioning," in *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*, 2016, pp. 162–170.

[11] N. Bartolini, T. He, V. Arrigoni, A. Massini, F. Trombetti, and H. Khamfroush, "On fundamental bounds on failure identifiability by boolean network tomography," *IEEE/ACM Transactions on Networking*, vol. 28, no. 2, pp. 588–601, 2020.

[12] J. Chen, X. Qi, and Y. Wang, "An efficient solution to locate sparsely congested links by network tomography," in *2014 IEEE International Conference on Communications (ICC)*, 2014, pp. 1278–1283.

[13] A. D. Jaggard, S. Kopparty, V. Ramachandran, and R. N. Wright, "The design space of probing algorithms for network-performance measurement," *SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 1, p. 105–116, jun 2013.

[14] S. Zhao, Z. Lu, and C. Wang, "Measurement integrity attacks against network tomography: Feasibility and defense," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 6, pp. 2617–2630, 2021.

[15] C.-C. Chiu and T. He, "Stealthy dgos attack: Degrading of service under the watch of network tomography," *IEEE/ACM Transactions on Networking*, vol. 29, no. 3, pp. 1294–1307, 2021.

[16] T. Hou, T. Wang, Z. Lu, and Y. Liu, "Combating adversarial network topology inference by proactive topology obfuscation," *IEEE/ACM Transactions on Networking*, vol. 29, no. 6, pp. 2779–2792, 2021.

[17] "Topology zoo," http://www.topology-zoo.org/dataset.html, accessed: March 15, 2023.