Towards Fine-Grained, High-Coverage Internet Monitoring at Scale

Hongyu Wu
Shanghai Jiao Tong University
Shanghai, China
why_36@sjtu.edu.cn

Qi Ling
University of Michigan
Ann Arbor, Michigan, USA
qiling@umich.edu

Penghui Mi
Huawei Cloud Computing Technologies Co., Ltd.
Dongguan, China
mipenghui@huawei.com

Chaoyang Ji
Huawei Cloud Computing Technologies Co., Ltd.
Beijing, China
jichaoyang@huawei.com

Yinliang Hu
Huawei Cloud Computing Technologies Co., Ltd.
Nanjing, China
huyinliang@huawei.com

Yibo Pi
Shanghai Jiao Tong University
Shanghai, China
yibo.pi@sjtu.edu.cn

ABSTRACT

The massiveness of the Internet makes it rather difficult to achieve high-coverage monitoring at scale with reasonable overhead. The traditional wisdom for scalable and high-coverage Internet monitoring is to consider clients in each /24 as a whole and only monitor the representatives, either by active probing or by passive traffic sniffing, such that performance of the rest can be predicted for high coverage. There are two basic assumptions behind this traditional wisdom: 1) clients in the same /24 have similar performance, and 2) tracking all targeted /24s equates to full-coverage monitoring. With the increasing prevalence of load balancing, both assumptions are now questionable. Through large-scale measurements, we evaluate the coverage and predictability issues of current practices, motivate the necessity of link-level fine-grained, high-coverage monitoring, and present new insights on how to achieve it. Our key findings are: 1) the current practices using the representatives of /24s may fail to capture the changes of up to 85% of links in the Internet; 2) the path difference between client flows to the same /24 is both significant and prevalent; 3) it is possible to cover most of the visible links from DCs to both small and large prefixes by carefully choosing client flows; 4) high-coverage monitoring can be achieved with at least three times less overhead than direct link monitoring.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Network monitoring**.

KEYWORDS

Internet Monitoring, Load Balancing

APNET 2023, June 29–30, 2023, Hong Kong, China

ACM ISBN 979-8-4007-0782-7/23/06…$15.00
https://doi.org/10.1145/3600061.3600085

ACM Reference Format:
Hongyu Wu, Qi Ling, Penghui Mi, Chaoyang Ji, Yinliang Hu, and Yibo Pi. 2023. Towards Fine-Grained, High-Coverage Internet Monitoring at Scale. In *7th Asia-Pacific Workshop on Networking (APNET 2023), June 29–30, 2023, Hong Kong, China.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3600061.3600085

1 INTRODUCTION

Due to the rapid growth of cloud services, cloud providers have continued to expand and enhance their datacenters (DCs) and private WANs over the past decade. Tremendous progress has been made in developing highly performant and available intra- and inter-DC networks [3, 13, 26]. In contrast, the public Internet, interconnecting most users and the DCs, evolves slowly with persistent disputes between ISPs over the provision of inter-domain links [18] and severe performance degradation due to network disruptions [23]. As a result, the public Internet often becomes the bottleneck for delivering seamless cloud services, and cloud providers have a strong incentive to closely monitor the public Internet for providing better services to their clients.

Internet monitoring can be conducted using either passive, active, or both types of measurements. Passive measurements could be inferred from client traffic at the server's side [16] or the intermediate routers [11], while active measurements require sending probes to the public Internet and observing the responses. With worldwide clients, major cloud providers can easily collect passive measurements for Internet monitoring. Small-sized cloud providers, however, often need active probing to mitigate coverage issues. For both small- and large-sized cloud providers, active probing can augment the data density of infrequent clients and plays a vital role in Internet debugging, e.g., fault localization [15] and root cause analysis for congestion [8].

In this paper, we focus on using active probing to achieve large-scale high-coverage monitoring at the level of IP links. With granular and wide-ranging knowledge of the Internet status, cloud providers can take not only remedial actions for affected clients, but also precautions against network degradations for future clients yet to be affected, by steering client traffic to a different DC or switching egress points. Furthermore, IP-level link monitoring enables cloud providers to better assist ISPs with fault localization and rate

the quality of their network services. Despite all the benefits above, fine-grained monitoring at the granularity of IP links is deemed incurring too much overhead. For scalable Internet measurement, current practices make two basic assumptions: 1) clients in the same /24 are similar [6, 16], such that it is sufficient to only monitor the representative of each /24 by active network probing [6, 7] or by passive traffic sniffing [16]; 2) tracking the performance to each /24 suffices for full-coverage monitoring. However, both assumptions are challenged by the increasing prevalence of load balancing, which creates a vast number of paths between DCs and client and causes the coverage and predictability issues of current practices. As client flows to the same /24 become less similar and could differ significantly in their paths, current practices lose their strengths in performance prediction.

By conducting large-scale measurements, we evaluate the coverage and predictability issues of current practices and provide insights on how to achieve link-level fine-grained, high-coverage monitoring with reasonable overhead. Specifically, we make the following contributions:

(1) We evaluate the link coverage of two rule-of-thumb practices for scalable Internet measurement from a cloud-centric view, where DCs are used as vantage points in each continent to probe the public Internet (§4.1). We find that up to 85% of links are not covered by current practices, leaving critical links unwatched. This link coverage issue is prevalent for networks in all six continents from the views of two major cloud providers, namely, Amazon and Alibaba.

(2) We evaluate the predictability of performance for client flows to the same /24s and find that the path difference between client flows to the same /24 is both significant and prevalent (§4.2). Such poor path similarity between flows indicates that performance prediction cannot effectively mitigate the coverage issues above for current practices.

(3) We propose to achieve high-coverage monitoring with an end-to-end approach, which covers most of the visible links by carefully selecting probing targets (§2.3). We show with experiments that it is feasible to cover above 80% of visible links by probing each selected target just once, and that this holds for both small and large prefixes (§5.2).

(4) We estimate the overhead for high-coverage monitoring by decomposing it into two determining factors: the scale of visible links and the effectiveness of the end-to-end approach in covering visible links. We find that visible links scale much slower than network size and are, on average, 3 times the scale of /24s (§5.1). Further, the end-to-end approach can cover all visible links with 40% less overhead than direct link monitoring (§5.3). The overhead can be further reduced by eliminating the long-tail effect in full link coverage, where one flow only covers one more link (§5.4).

To promote reproducibility, we publish our tool and dataset at *github.com/SJTU-NMS-Lab/APNet23*.

2 BACKGROUND & MOTIVATION

Internet monitoring with passive measurements is believed to perform well in regions with dense clients and capable of augmenting data coverage with prediction in regions with sparse clients by leveraging client similarities. Both the coverage and predictability of passive measurements are challenged by the prevalence of load balancing.

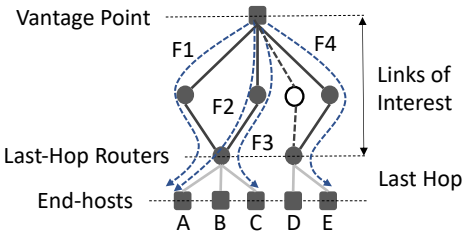

Figure 1: End-to-end approach to high link coverage

2.1 Load Balancing Worsens Coverage

Load balancing has been found prevalent both within and between autonomous systems (ASs) [9], where up to 90% of source-destination pairs were reported to have a load balancer in between [2]. This results in a vast number of paths between DCs and their clients. As passive measurements are collected from client flows taking random load-balanced paths, using passive measurements for Internet monitoring is equivalent to randomly sampling Internet paths. The network coverage is totally determined by the distribution of clients, and a skewed distribution of clients could easily leave some critical links unwatched. These links, if congested or failed, would affect all future client flows passing through.

2.2 Load Balancing Weakens Predictability

When two flows share most of the paths, it is easy to predict the performance of one flow using the other, where similarity leads to predictability. This is the rationale behind data-driven Internet monitoring using passive measurements. However, the prevalence of load balancing reduces the similarity between Internet flows, weakening the power of passive measurements in prediction. Specifically, load balancers can distribute traffic across different paths at the levels of packets, flows, or destinations. Per-destination load balancing may distribute flows across different paths for destinations in the same subnet (e.g., in the same /24), while per-flow load balancing may cause performance to vary greatly even for flows between the same source-destination pair [21], let alone per-packet load balancing. The decreased similarity makes it more difficult to infer the performance of a network without actually probing it.

2.3 High-Coverage Internet Monitoring

The coverage and predictability issues above indicate that it is ideal to monitor all links to the network of interest. However, not all layer-3 links are discoverable by probes, and the overhead of full link coverage is deemed intractable. We thus limit our focus on **high visible link coverage**[1], i.e., monitoring most of the visible IP links to the targeted network to provide the best coverage possible, where the last hops connecting end-hosts to the public Internet are excluded, which can only be covered by abusively probing each end-host.

Instead of directly probing each link, we want to take an end-to-end approach to high link coverage by carefully selecting client flows, such that most of the visible links are traversed by the selected client flows. Since link problems affect the end-to-end latency,

[1] In the following, we will simply refer to it as *high link coverage*.

the latency fluctuations of these flows can be used to detect link problems at scale. As it only takes one probe to measure each flow latency once, this end-to-end approach has great potential to achieve high link coverage with much less overhead than traditional methods (§4.1). Moreover, this approach with active probing can be used together with passive measurements to reduce overhead and increase coverage, where probing mainly targets for uncovered regions.

Figure 1 shows an example of this end-to-end approach, where all visible links (solid lines) can be covered using either flows $F1$, $F2$, and $F4$ or flows $F1$, $F3$, and $F4$. The later is preferred, because each end-host is associated with one flow, while the former requires probing end-host A twice as often as end-host E, causing uneven burdens across end-hosts. Since each probe in a flow traverses (or covers) all links along the flow's path, we say that there exists a **uniform high-coverage realization** if each selected end-host only needs to be probed once to cover most of the visible links.

3 DATASET

We want real-world measurement data for the networks between major cloud providers and their clients to evaluate the link coverage issues of current practices (§4) and to shed light on a new approach to high-coverage link monitoring with reasonable overhead (§5). To our best knowledge, no publicly available dataset targets full link coverage and we need to create our own datasets.

3.1 Vantage Points and Targets Selection

We chose DCs from two major cloud providers (Amazon and Alibaba) as VPs and IPv4 addresses in the public Internet as targets. For wide geographical coverage, we selected one DC in each continent (Alibaba has no DC in Africa and South America). As it is common for providers to establish direct peering with different ISPs, they may reach the same clients through different networks. To compare between providers, we intentionally chose DCs residing in the same metropolitan areas (Frankfurt, Silicon Valley and Sydney) for both providers. In Asia, we chose Alibaba's and Amazon's DCs in Beijing and Tokyo, respectively. In South America and Africa, we chose Amazon's DCs in São Paulo and Cape Town, respectively. As the major traffic volume for DCs is from clients in the same city or continent [6], we selected for each DC the /8 prefix that covers the largest amount of addresses in the host country of the DC as target.

3.2 Measurement Methodology

The major tasks of our datasets are 1) evaluating link coverage issues of current practices as well as 2) verifying our insight for IP-level link monitoring. Specifically, our core insight is that as probes to a target traverse all links along the path to the target, link issues would be reflected on the end-to-end latency and it is thus possible to use latency fluctuations as an indicator for potential link issues. Further, if probing one target can cover all links on a path, we ask that *if it is possible to achieve full link coverage by carefully selecting the probing targets.* Compared to directly probing each link, this indirect approach has potential to incur much less overhead.

To accomplish the tasks above, we create three datasets: 1) ground-truth dataset, which covers all visible links between the selected vantage points and targets and provides the ground truth for evaluating link coverage issues; 2) random-flow dataset, which can be used to mimic client traffic for evaluating the issues of using passive measurements; 3) full-coverage flow dataset, which includes a set of paths covering all visible links to examine the feasibility of IP-level link monitoring. It should be noted that the links in 1) and the paths in 2) and 3) should reflect the same network state for fair comparison. We discuss in detail how each dataset is collected.

Ground-truth dataset: full link coverage. Due to Internet load balancing, there could exist many alternative paths between two hosts. Finding all visible links in between can be done by discovering all the load-balanced paths. Several measurement tools have been designed for this purpose, of which D-Miner is the most recent one capable of enumerating load-balanced paths from a VP to a given prefix at high speed. We thus choose to use D-Miner for link discovery. However, as the pool of targets is very large, it may still take D-Miner days to scan a network of this scale [24]. As mentioned earlier, we want the links discovered here to reflect the same network state as the paths to be found later for datasets 2) and 3). This implies that we need to divide the target pool into small prefixes that can be scanned in minutes by D-Miner to mitigate the impact of route dynamics.

In our measurement campaigns, we divided targets into /16 prefixes and scanned each /16 sequentially with D-Miner at a probing rate of 100,000 pps, the same scanning rate used in [24] for Internet-wide survey. During our campaigns, we scanned each /16 twice in a row to understand how route changes and packet loss might incur link difference between two snapshots. We found less than 2% of link difference between two back-to-back scans to the same /16. The confidence level of D-Miner to discovering all load-balanced paths is set to 0.95. Further improving the confidence level to 0.99 only discovered about a few percent more links at the cost of tens of percent more probes.

Random flow dataset: mimicking client traffic. An Internet flow is identified by its flow ID, <src port, dst port, src addr, dst addr, proto id>, which determines how it will be routed in the Internet. To evaluate the link coverage issues of using passive measurements, we want to mimic client flows with active probing. Specifically, we sent a train of probes with increasing TTLs from the VP to each address, with all probes sharing the same source and destination ports, to discover all visible links in between. The source port was set to 80 and the destination port was randomized between 49152 and 65535, the range for ephemeral ports. For each /16, this process follows immediately after D-Miner finishes scanning, such that it measures the same Internet state as D-Miner did. To speed up this process, we modified ZMap, a stateless prober, to support path discovery to a given address and scan each address in the targeted /16. For stateless operations, the IPID field in the IP header is used to encode the original TTLs of probes, which will later be extracted from the responses. ZMap probed at the same rate as D-Miner for fast scanning and used UDP probes to simulate the downstream traffic flows from DCs to clients.

Full-coverage flow dataset: full link coverage by flows. We are interested to know if it is possible to cover most of the visible links by carefully selecting client flows, i.e., if a uniform high-coverage

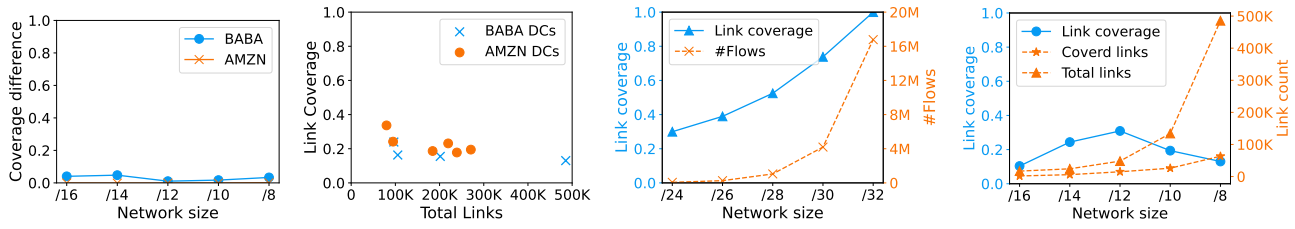

(a) Comparing current practices (b) Coverage vs. total links (c) Coverage vs. prefix size (d) Coverage vs. network size

Figure 2: Coverage issues of current practices for scalable Internet measurement

realization exists between a DC and its clients. To this end, we need to identify a set of flows whose paths cover all visible links discovered by D-Miner. Fortunately, D-Miner accomplishes link discovery by controlling the flow IDs. Simply by using the same set of flow IDs used to discover links, we can discover and collect the paths for a set of flows with full link coverage. We modified ZMap to support path discovery with given source and destination ports and used it to complete this process.

4 COVERAGE & PREDICTABILITY ISSUES

This section presents the coverage issues of current practices for scalable Internet measurement and the predictability issues of flow performance for coverage expansion.

4.1 Coverage Issues of Current Practices

The rule-of-thumb practice for both active and passive Internet measurement is to consider each /24 as a whole and use one end-host as the representative [16, 17]. Active probing techniques commonly target the .1 addresses in /24s for a higher response rate [10, 22], while passive monitoring selects a random client visiting the site as the representative for its /24 [6]. We want to evaluate the link coverage issues of these two popular practices using our random flow dataset, which is collected to include a random path from the selected VPs to each address in the targeted networks. Specifically, we collate links along the paths to the selected addresses and calculate the link coverage, i.e., the percentage of the covered links over all visible links, for the two practices.

Coverage difference between current practices. Recall that we probed from each DC to a /8 prefix in our measurement survey. To study the link coverage under different network sizes, we split the /8 into smaller prefixes of different sizes. Figure 2(a) shows the link coverage difference between the two practices for both Amazon's and Alibaba's DCs, where the link coverage is averaged over all equal-sized prefixes under their respective /8s. It can be seen that there is only a few percent difference in coverage between the two practices. In the following, we simply assume that .1 addresses are used as the representatives.

Low link coverage. Figure 2(b) shows the link coverage from each DC to their respective /8s, where DCs have a wide range of link counts, depending on the link density, hop count and response rate. Except for Amazon Cape Town's DC, all other DCs have link coverage ranging from 0.15 to 0.25, which do not change much with the link count. This means that current practices only cover about one-fifth of all visible links, leaving a majority of links unwatched. Any network events happening to these links will not be observed by current practices, and no precautions for future client flows can be taken. To understand the impact of cloud private networks, we re-calculate coverage for links in public networks. The link coverage almost remains the same, which means that private networks are not the culprit for low link coverage.

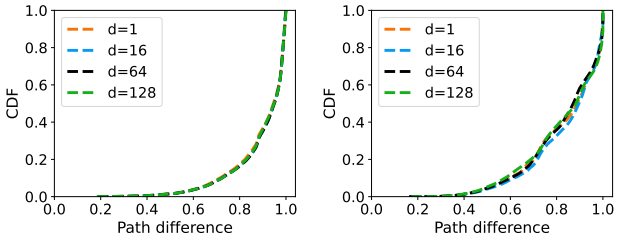

(a) Amazon (b) Alibaba

Figure 3: Path difference between similar flows

Trading off scalability for link coverage. The traditional wisdom to boost link coverage is to increase the granularity of monitoring. We re-do our experiments with prefixes smaller than /24s. Figure 2(c) shows the tradeoff between link coverage and prefix size, where the link coverage increases linearly to 100% when the prefix size decreases to /32, i.e., all addresses are used for monitoring. It is apparent that using traditional wisdom to improve link coverage is not scalable, where exponential increment in monitoring scale only leads to linear increment in coverage.

Impacting factors for link coverage. Link coverage is the proportion of covered links among all links. Both the covered and total links increase with the network size, but at different rates depending on the network topology. Figure 2(d) shows the link coverage for Alibaba Beijing's DC, where the link coverage first increases and decreases afterwards. This is a typical example showing how skewed distribution of link density in the network affects link coverage.

4.2 Predictability Issues for Similar Flows

The basic assumption behind current practices is that clients in the same /24 are similar. It is sufficient to monitor only the representative to predict the performance of the rest. To verify this

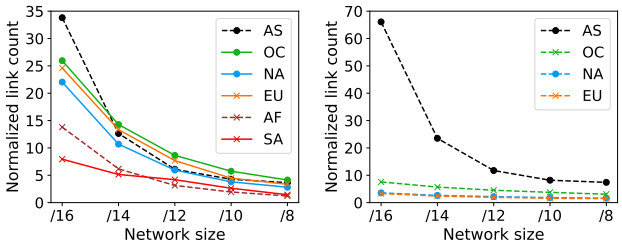

(a) Amazon (b) Alibaba

Figure 4: Links scale much slower than network size

assumption, we calculate the path difference between flows to the same /24s. Denote the sets of links for flows A and B as s_A and s_B. The path difference between flows A and B is defined as $(|s_A \cup s_B| - |s_A \cap s_B|)/(|s_A| + |s_B|)$. It is unlikely for flows with significant path difference to perform similarly. We compute the path difference between two flows respectively to address x and $x + d$, where x is the .0 address in a /24 and $1 \leq d \leq 255$ is the distance between addresses. Figure 3 shows the distribution of path difference between flows to the same /24s under different d's. For both Amazon's and Alibaba's DCs, the path difference is significant and prevalent: 80% of flow pairs have a path difference more than 70%.

5 HIGH-COVERAGE LINK MONITORING

We propose to achieve high-coverage link monitoring with an end-to-end approach that covers most of the links by a set of carefully selected flows to end-hosts. The overhead of this approach depends on the number of flows to be monitored and can be expressed as

$$overhead \propto total\_links \times \frac{1}{avg\_unique\_links\_per\_flow},$$

where the $avg\_unique\_links\_per\_flow$ means the average number of unique links covered by each flow, equal to the total number of unique links divided by the number of flows. To understand the overhead, we first look at how links scale with network size.

5.1 How Do Links Scale with Network Size?

We create networks of different sizes by dividing the targeted /8s into smaller prefixes. Recall that our dataset includes full coverage of visible links from VPs to each /16 of their respective /8s. We can thus collate links for each prefix and calculate the average number of links over all equal-sized prefixes. Figure 4 shows the normalized link count under different network sizes, where the link count is normalized respectively by the number of /24s under each network size. For /16s, the normalized link count is large in Asia for both Amazon's and Alibaba's DCs, because there could exist thousands of links between a VP and a /16, greatly elevating the average normalized link count. The impact of these links is amortized as network size increases. For /8s, the average normalized link counts are below 10 for all continents. Africa and South America even have links at a similar scale as /24s, which may further decrease for larger networks. This implies that if the $avg\_unique\_links\_per\_flow$ is

Table 1: Link Coverage for Uniform High-Coverage Realizations

	Prefix Size				
	/16	/14	/12	/10	/8
Amazon	97.1%	94.6%	89.6%	84.0%	83.1%
Alibaba	95.0%	93.2%	87.2%	86.1%	87.5%

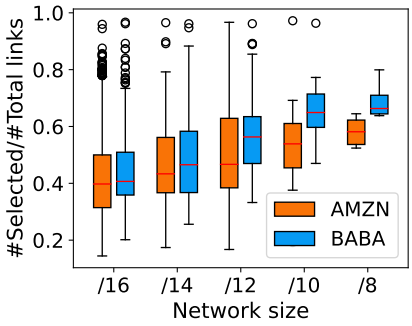

Figure 5: End-to-end approach to high link coverage

large, the end-to-end approach can achieve full link coverage with reasonable overhead.

5.2 Uniform High-Coverage Realization

Since each probe covers all links along the flow's path, we say that a uniform high-coverage realization exists if most of the visible links can be covered by probing each selected flow once. Before jumping into the solutions to high-coverage monitoring, we first want to know if uniform high-coverage realizations exist from the cloud to the public Internet. To this end, we want to initiate a flow from the VP to each address in the targeted prefix such that the link coverage can be maximized. This process assembles what D-Miner does in its first stage of link discovery: generating flows with varying destination addresses from .1 to .255 in each /24 to discover load balancing. If D-Miner discovers most of the links in its first stage, we can achieve high coverage with the same set of flows, i.e., a uniform high-coverage realization exists. Table 1 shows the link coverage for uniform high-coverage realizations, where the average link coverage for /16s is above 95% for all DCs and is still above 80% for /8s. This means that uniform high-coverage realizations exist for both small and large prefixes. After knowing its existence, we next design a greedy algorithm to achieve uniform high coverage.

5.3 A Greedy End-to-End Approach

The goal of the end-to-end approach is to achieve the required coverage with the least number of flows, or equivalently, to maximize the average number of unique links covered by each flow. As a first step, we design a simply greedy algorithm to obtain a rough estimate, which always selects the address to which the flow discovers the most number of new links. When two addresses contribute the same number of new links, we randomly select one to proceed. This process is repeated until all visible links are covered. Figure 5 shows the ratio of the selected flows to the total links. The average

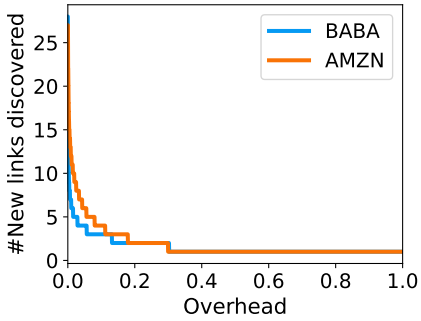

Figure 6: Coverage and overhead tradeoff

ratio increases with the network size, because it becomes more difficult to cover all visible links for larger networks. Nonetheless, on average, the number of flows needed is 60% of the total links for /8s.

5.4 Coverage & Overhead Tradeoff

In the greedy algorithm above, since addresses discovering the highest number of new links are selected first, the efficiency of adding new flows decreases gradually. As shown in Figure 6, only the first 30% of overhead can discover more than one new links and there is a long tail in the distribution. Suppose we are interested in covering 80% of links, the overhead can be reduced by 62% and 68% for Amazon's and Alibaba's DCs, respectively. This further decreases the overhead obtained from the greedy algorithm. For 80% of link coverage, the number of flows needed is only about one-third of the total links.

6 RELATED WORK

For scalable Internet monitoring, past efforts attempt to aggregate similar clients by their attributes (e.g., geolocation [19] and BGP prefix [4]), and only monitor the representative [6, 16]. Tracking the representative can be done either by active probing [1, 5, 20] or by passive traffic sniffing [11]. Major cloud providers can reduce measurement overhead by leveraging passive measurements collected from their client traffic [16]. Passive measurements are commonly used together with active probing for network debugging [14]. Network tomography attempts to infer link-level performance with partial measurements, but runs into scalability issues for the public Internet monitoring [12, 25]. None of the methods above considers IP-level high-coverage Internet monitoring at scale.

7 CONCLUSION AND FUTURE WORK

In this work, we evaluated the issues of current practices in both coverage and predictability. Our results show that current practices fail to monitor the changes of a majority of links in the Internet, leaving critical links unwatched. This motivates IP-level high-coverage Internet monitoring, aiming to capture critical link events by covering most of the visible links. To this end, we propose an end-to-end approach that covers links by carefully selecting probing targets, with great potential to achieve high-coverage monitoring with reasonable overhead. There are still many unaddressed issues for future work, including the temporal variations of link coverage and the optimal flow selection algorithm.

ACKNOWLEDGMENTS

We appreciate the constructive feedback from the anonymous reviewers. This work is partially supported by the Huawei-SJTU ExploreX Funding (SD6040004/052). The corresponding author is Yibo Pi (yibo.pi@sjtu.edu.cn).

REFERENCES

[1] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. 2006. Avoiding traceroute anomalies with Paris traceroute. In *IMC*. 153–158.
[2] B. Augustin, T. Friedman, and R. Teixeira. 2007. Measuring load-balanced paths in the Internet. In *IMC*. 149–160.
[3] R. Basat, S. Ramanathan, Y. Li, G. Antichi, M. Yu, and M. Mitzenmacher. 2020. PINT: Probabilistic in-band network telemetry. In *SIGCOMM*. 662–680.
[4] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye. 2015. Analyzing the performance of an anycast CDN. In *IMC*. 531–537.
[5] M. Calder, R. Gao, M. Schroder, R. Stewart, J. Padhye, R. Mahajan, G. Ananthanarayanan, and E. Katz-Bassett. 2018. Odin: Microsoft's Scalable Fault-Tolerant CDN Measurement System. In *NSDI*. 501–517.
[6] F. Chen, R. K. Sitaraman, and M. Torres. 2015. End-user mapping: Next generation request routing for content delivery. *SIGCOMM CCR* 45, 4 (2015), 167–181.
[7] A. Dainotti, C. Squarcella, E. Aben, kc claffy, M. Chiesa, M. Russo, and A. Pescapé. 2011. Analysis of country-wide internet outages caused by censorship. In *IMC*. 1–18.
[8] A. Dhamdhere, D. D. Clark, A. Gamero-Garrido, M. Luckie, R. KP Mok, G. Akiwate, K. Gogia, V. Bajpai, A. C. Snoeren, and kc claffy. 2018. Inferring persistent interdomain congestion. In *SIGCOMM*. 1–15.
[9] A. Dhamdhere and C. Dovrolis. 2011. Twelve years in the evolution of the Internet ecosystem. *IEEE/ACM Transactions on Networking* 19, 5 (2011), 1420–1433.
[10] X. Fan and J. Heidemann. 2010. Selecting representative IP addresses for Internet topology studies. In *IMC*. 411–423.
[11] S. Gangam, J. Chandrashekar, Í. Cunha, and J. Kurose. 2013. Estimating TCP latency approximately with passive measurements. In *PAM*. 83–93.
[12] D. Ghita, C. Karakus, K. Argyraki, and P. Thiran. 2011. Shifting network tomography toward a practical goal. In *CoNext*. 1–12.
[13] C. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. 2013. Achieving high utilization with software-driven WAN. In *SIGCOMM*. 15–26.
[14] Y. Jin, S. Renganathan, G. Ananthanarayanan, V. N. Padmanabhan J. Jiang, M. Schroder, M. Calder, and A. Krishnamurthy. 2019. Zooming in on wide-area latencies to a global cloud provider. In *SIGCOMM*. 104–116.
[15] R. R. Kompella, J. Yates, A. Greenberg, and Alex C. Snoeren. 2009. Fault localization via risk modeling. *IEEE Transactions on Dependable and Secure Computing* 7, 4 (2009), 396–409.
[16] R. Krishnan, H. V. Madhyastha, S. Srinivasan, S. Jain, A. Krishnamurthy, T. Anderson, and J. Gao. 2009. Moving beyond end-to-end path information to optimize CDN performance. In *IMC*. 190–201.
[17] Y. Lee and N. Spring. 2016. Identifying and aggregating homogeneous ipv4/24 blocks with Hobbit. In *IMC*. 151–165.
[18] M. Luckie, A. Dhamdhere, D. Clark, B. Huffaker, and kc claffy. 2014. Challenges in inferring internet interdomain congestion. In *IMC*. 15–22.
[19] V. N. Padmanabhan and L. Subramanian. 2001. An investigation of geographic mapping techniques for Internet hosts. In *IMC*. 173–185.
[20] C. Pelsser, Luca Cittadini, Stefano Vissicchio, and Randy Bush. 2013. From Paris to Tokyo: On the suitability of ping to measure latency. In *IMC*. 427–432.
[21] Y. Pi, S. Jamin, P. Danzig, and F. Qian. 2020. Latency imbalance among Internet load-balanced paths: A cloud-centric view. *SIGMETRICS* 4, 2 (2020), 1–29.
[22] L. Quan, J. Heidemann, and Y. Pradkin. 2013. Trinocular: Understanding internet reliability through adaptive probing. *SIGCOMM CCR* 43, 4 (2013), 255–266.
[23] A. Schulman and N. Spring. 2011. Pingin'in the rain. In *SIGCOMM*. 19–28.
[24] K. Vermeulen, J. P. Rohrer, O. Fourmaux R. Beverly, and T. Friedman. 2020. Diamond-Miner: Comprehensive Discovery of the Internet's Topology Diamonds. In *NSDI*. 479–493.
[25] L. Xue, M.K. Marina, G. Li, and K. Zheng. 2022. PAINT: Path Aware Iterative Network Tomography for Link Metric Inference. In *ICNP*. 1–12.
[26] Y. Zhao, K. Yang, Z. Liu, T. Yang, L. Chen, S. Liu, N. Zheng, R. Wang, H. Wu, Y. Wang, and N. Zhang. 2021. LightGuardian: A Full-Visibility, Lightweight, In-band Telemetry System Using Sketchlets. In *NSDI*. 991–1010.