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

Hongyu Wu[1], Qi Ling[2], Penghui Mi[3], Chaoyang Ji[3], Yinliang Hu[3], Yibo Pi[1]

[1]Shanghai Jiao Tong University

[2]University of Michigan

[3]Huawei Cloud Computing Technologies Co., Ltd.

JOINT INSTITUTE
交大密西根学院

Public Internet: Bottleneck of Cloud Services

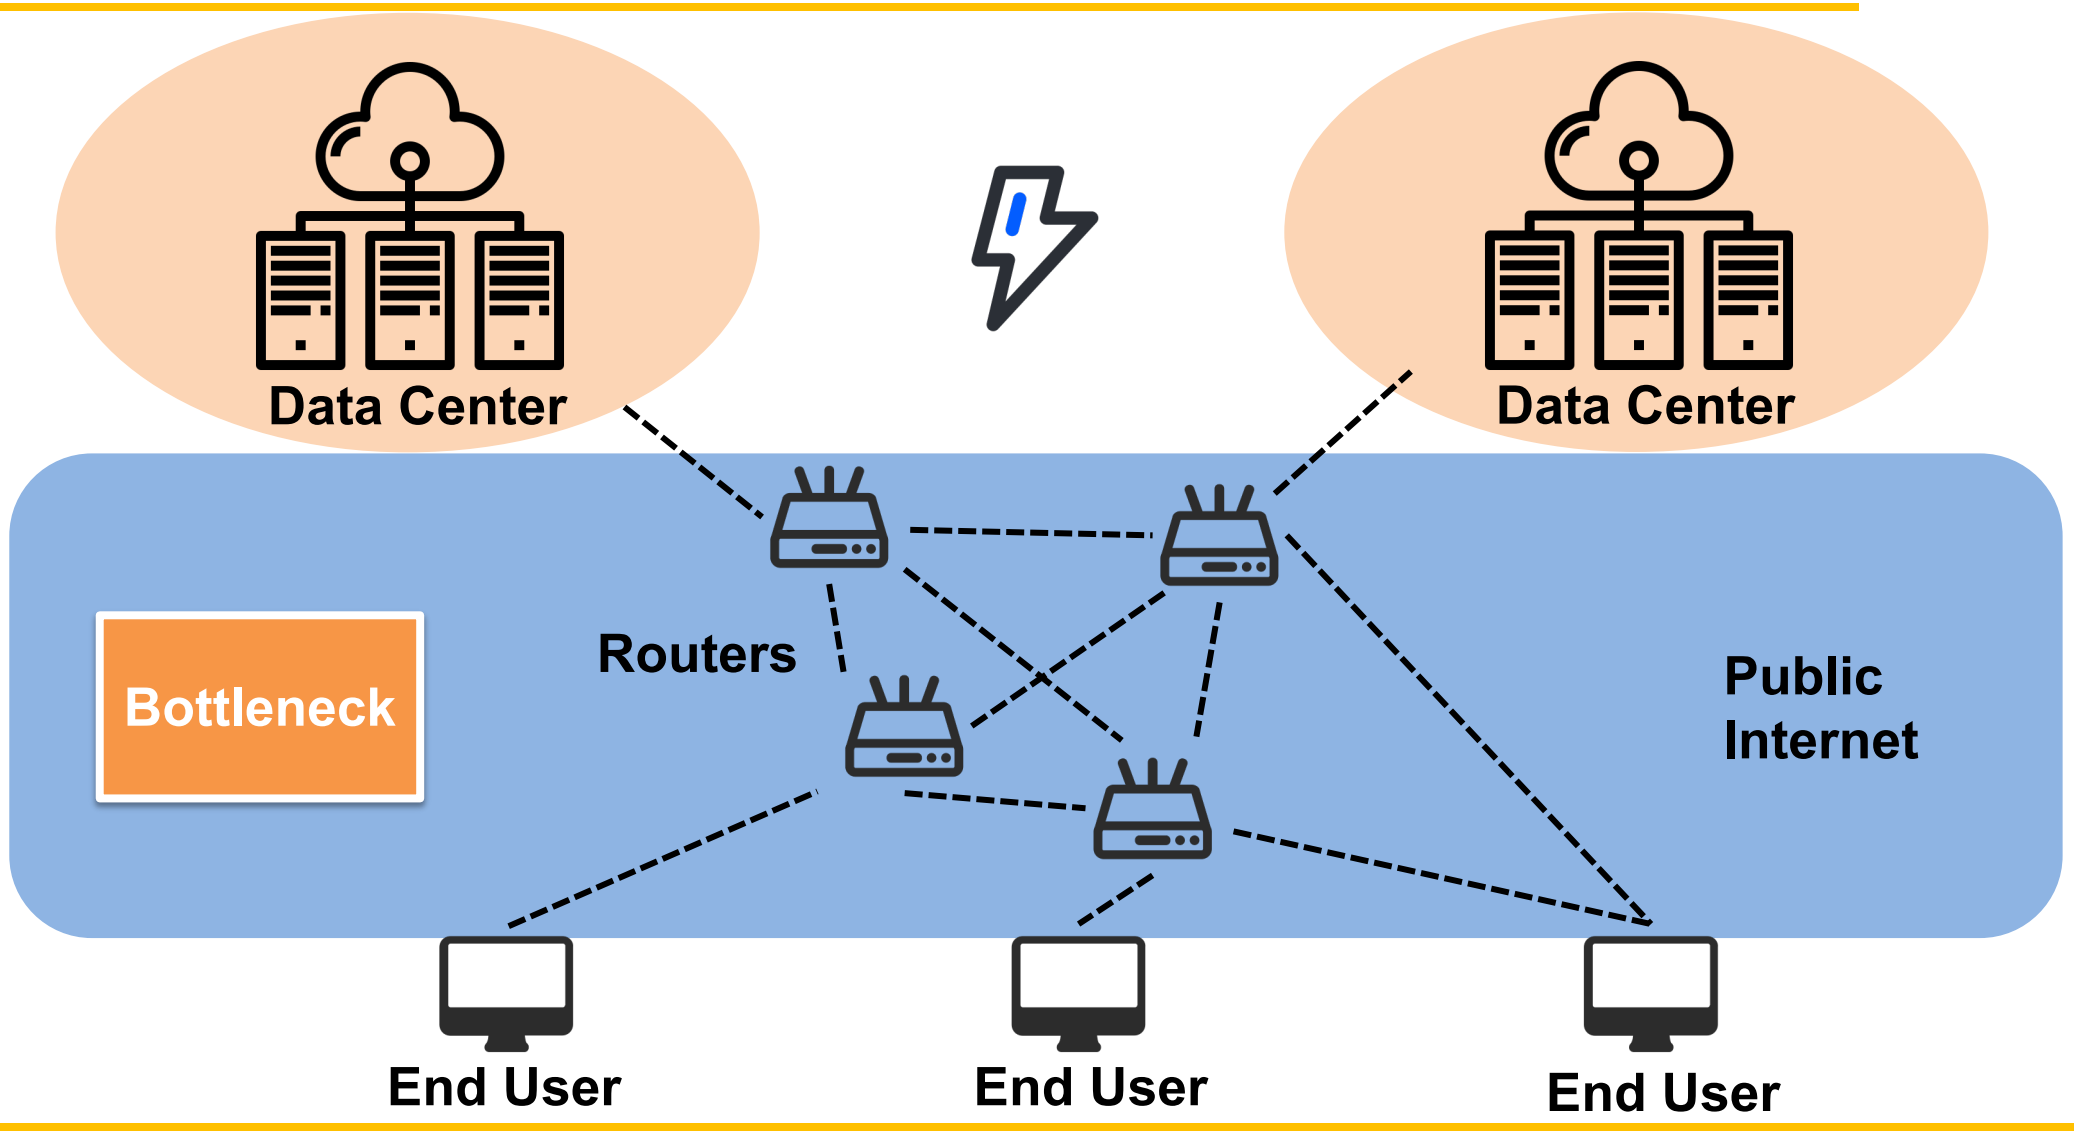

Public Internet: Bottleneck of Cloud Services

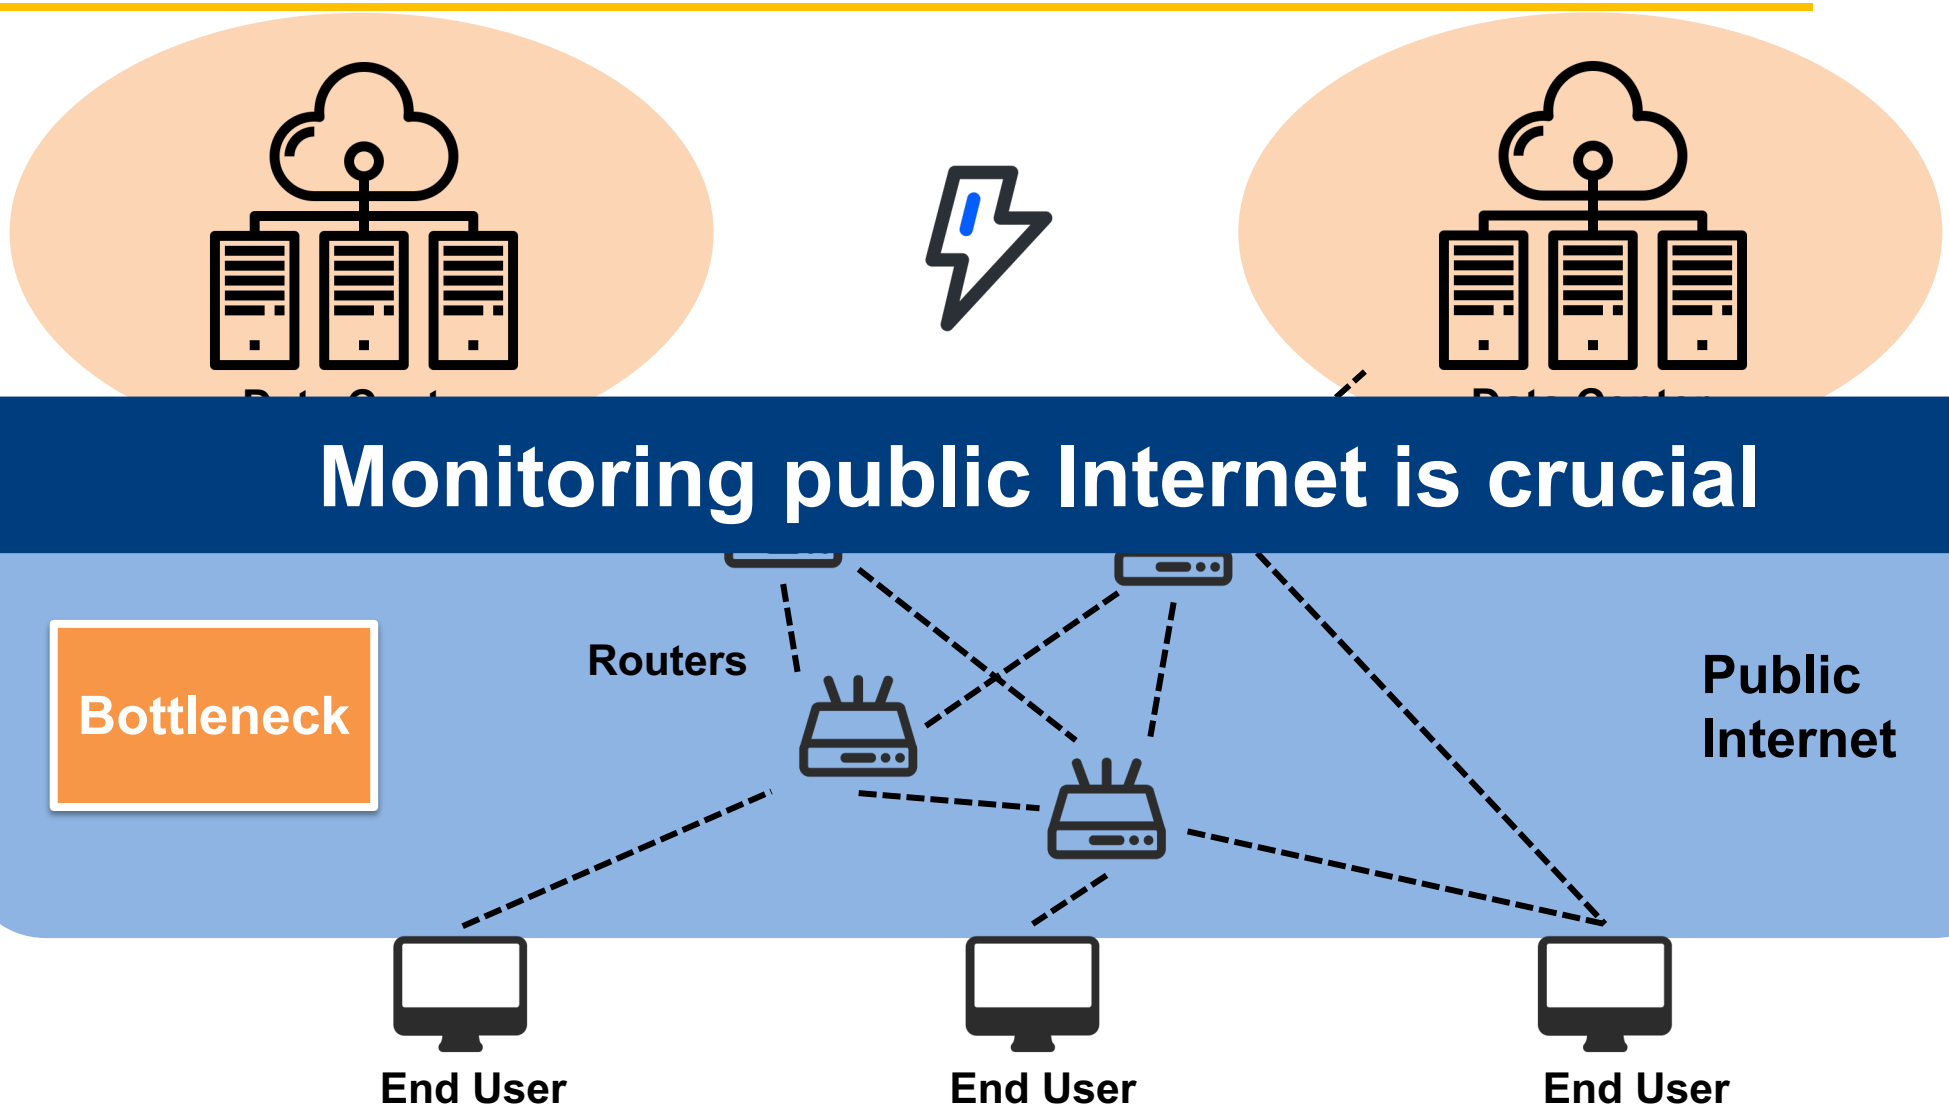

Challenge of Internet Monitoring

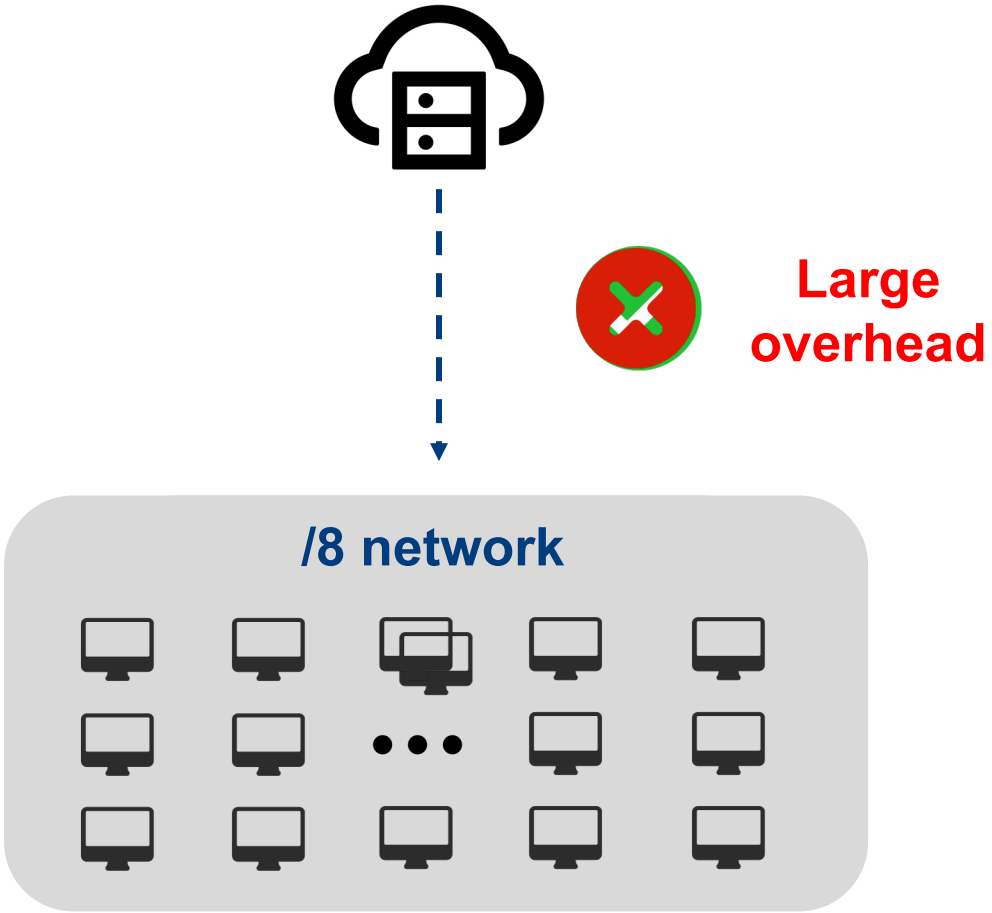

Challenge of Internet Monitoring

Large overhead

How to efficiently achieve large-scale Internet monitoring?

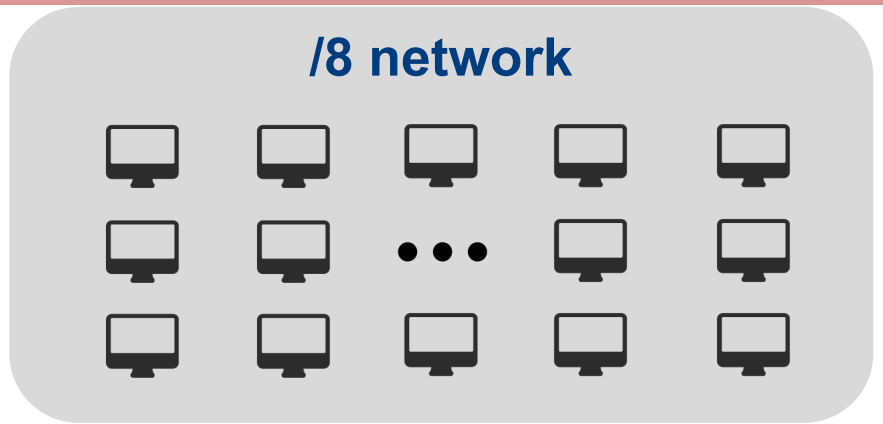

Current Practice of Internet Measurement

- **Two basic assumptions**
  1. **Similarity assumption**: Clients in the same /24 have similar paths
     - Only representative in each /24 needs to probed, reducing probing overhead
  2. **Coverage assumption**: Tracking the performance to each /24 suffices for full-coverage monitoring

Both assumptions are challenged by the increasing prevalence of load balancing

Load Balancing Challenges Similarity Assumption

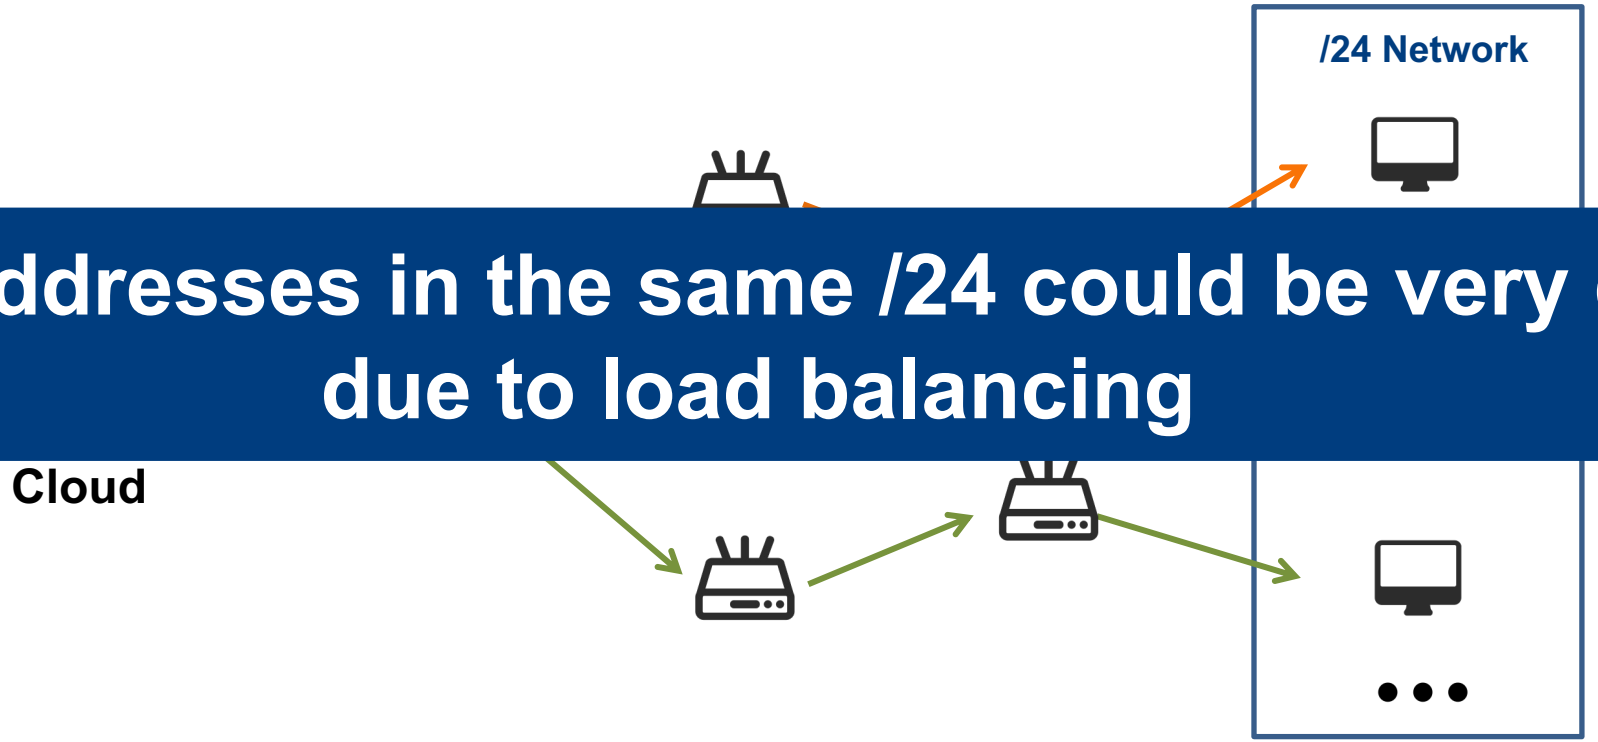

Paths to addresses in the same /24 could be very different due to load balancing

Load Balancing Challenges Coverage Assumption

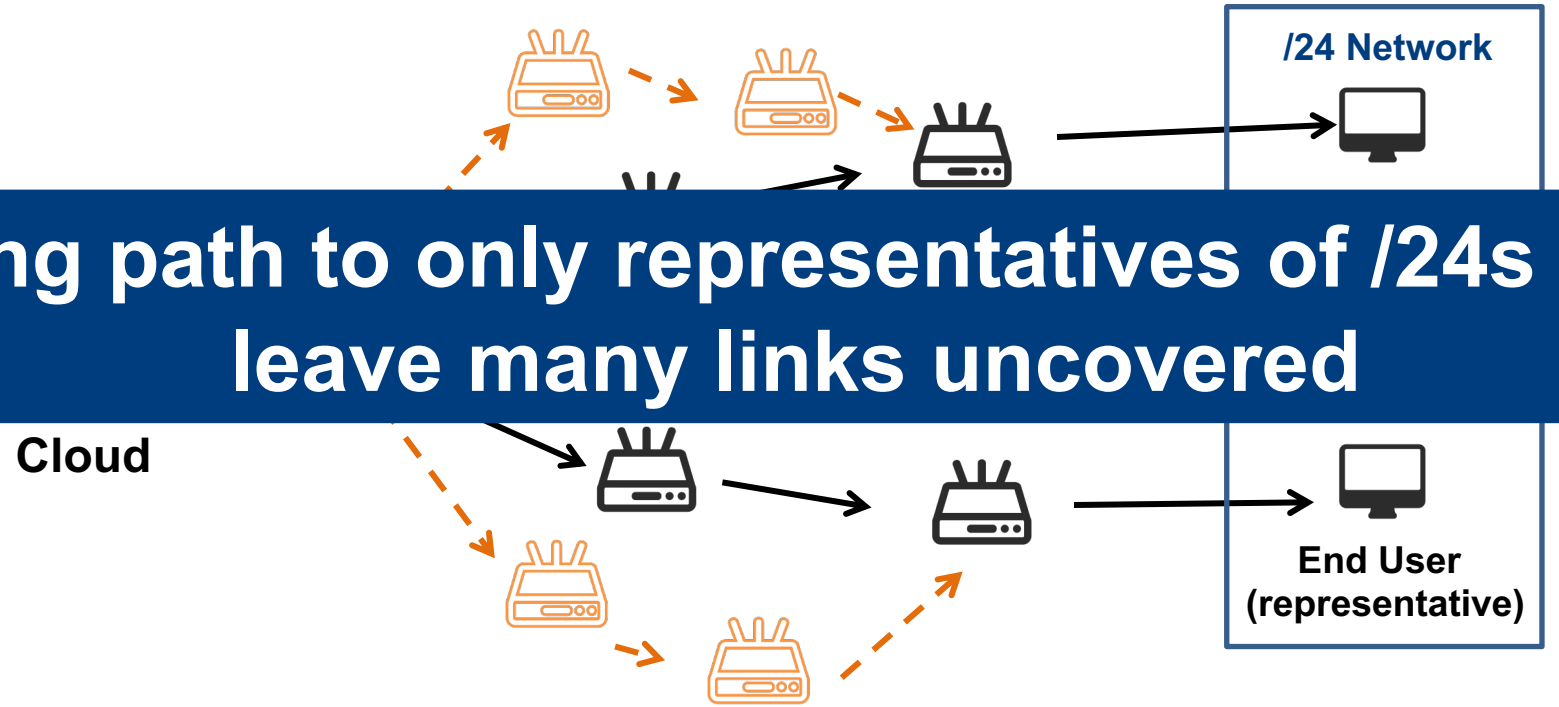

Probing path to only representatives of /24s would leave many links uncovered

Methodology

Ground truth:
all visible links

Use D-Miner [NSDI'20] to find all visible links

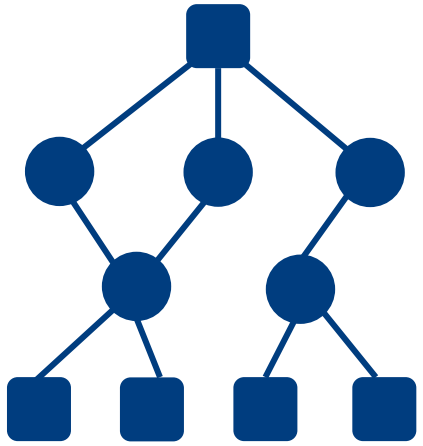

Flow paths

Modify Zmap to find the flow path from vantage points to each end user

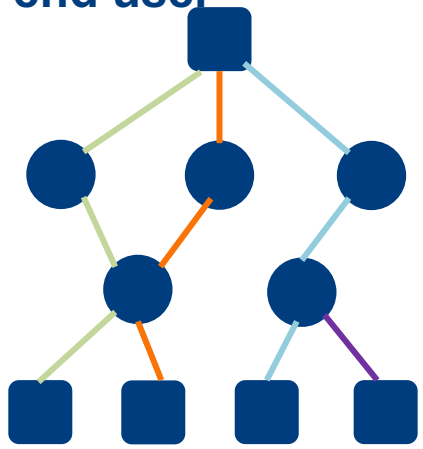

Simulation

- Simulate real-world downstream traffic from DCs to clients
- Evaluate current practice against ground truth

VP and Target Selection

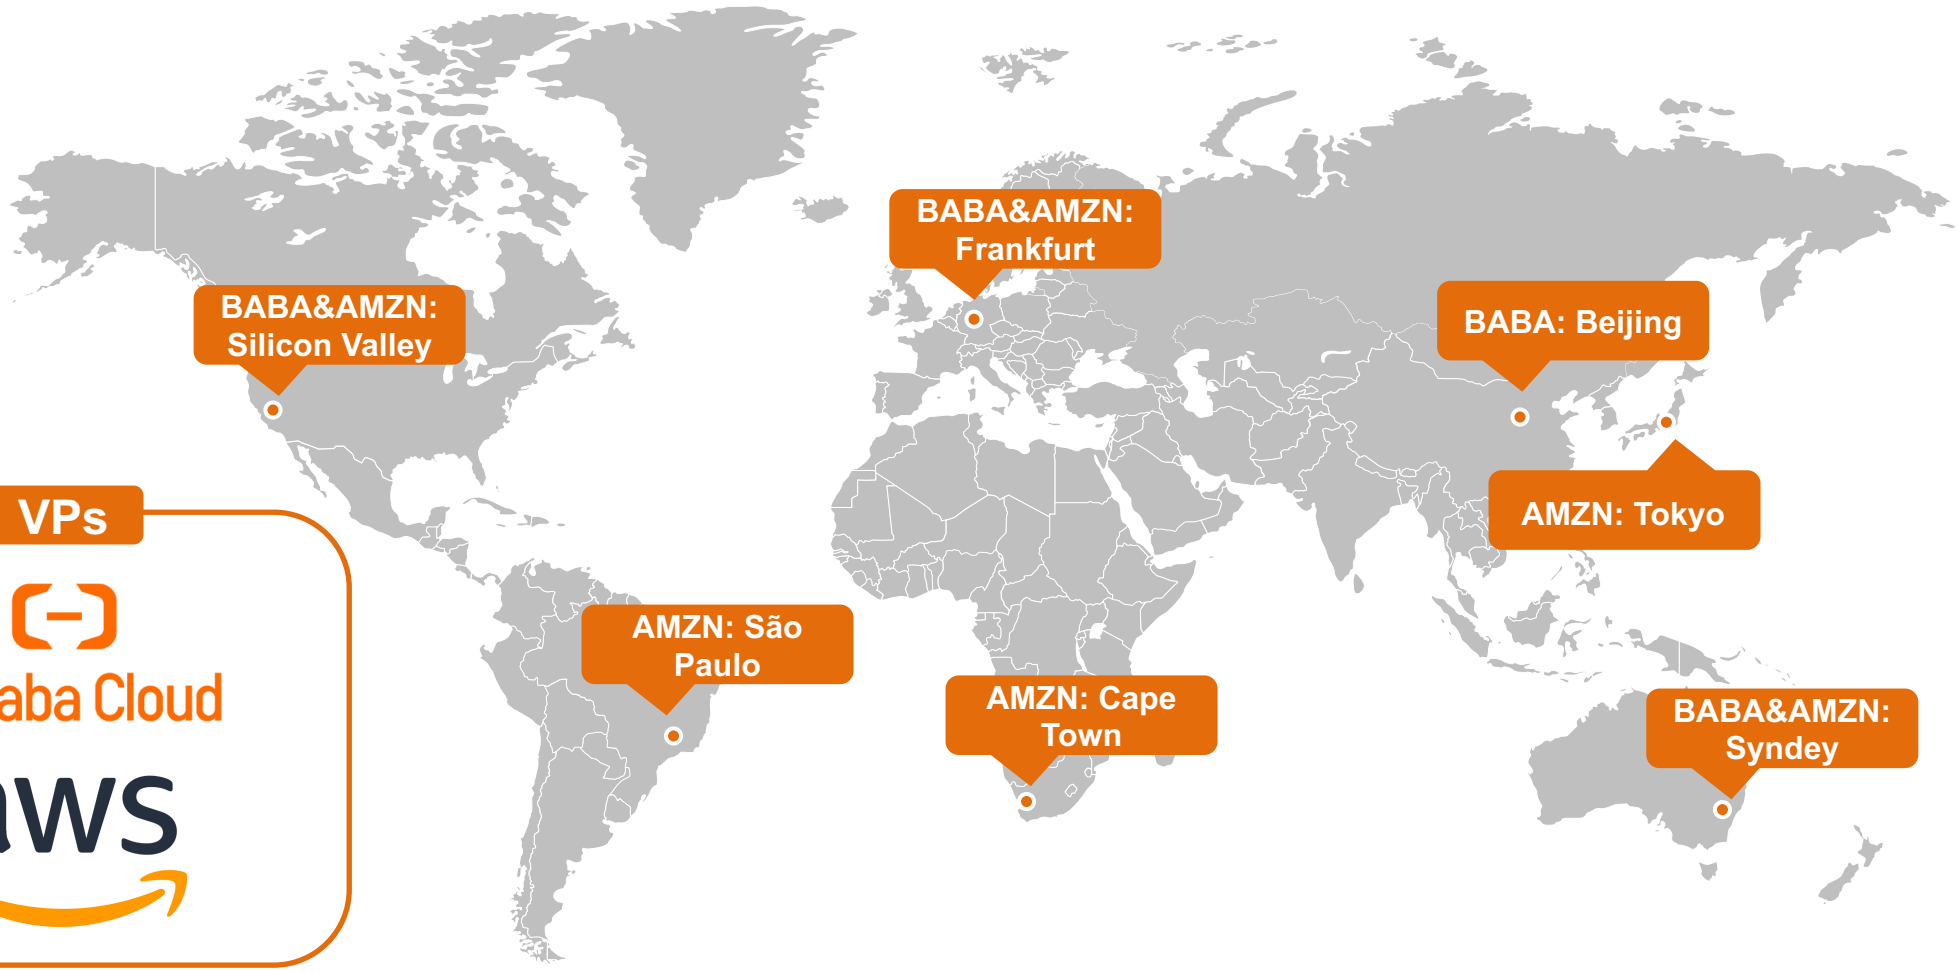

VPs

Alibaba Cloud

aws

VP and Target Selection

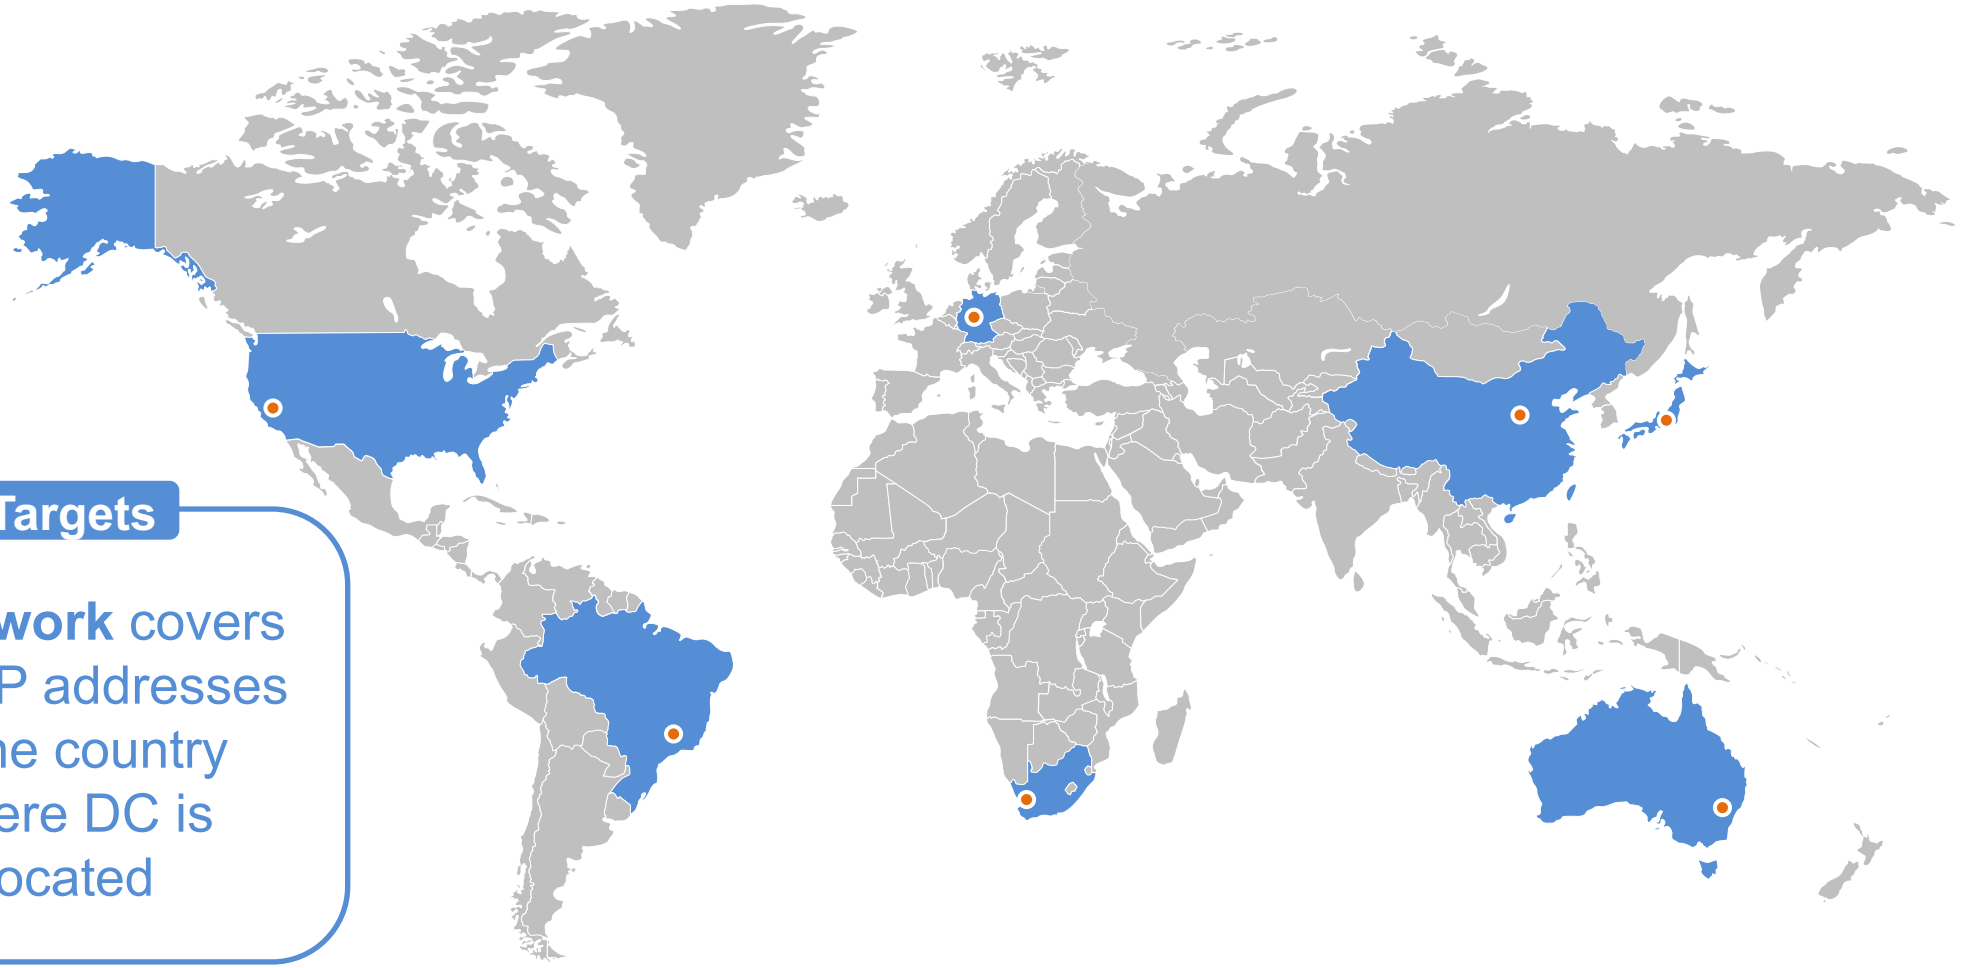

Targets

/8 network covers most IP addresses in the country where DC is located

Evaluating Similarity Assumption of Current Practices

Similarity assumption: Clients in the same /24 have similar paths

Path difference = $(|S_A \cup S_B| - |S_A \cap S_B|)/(|S_A| + |S_B|)$

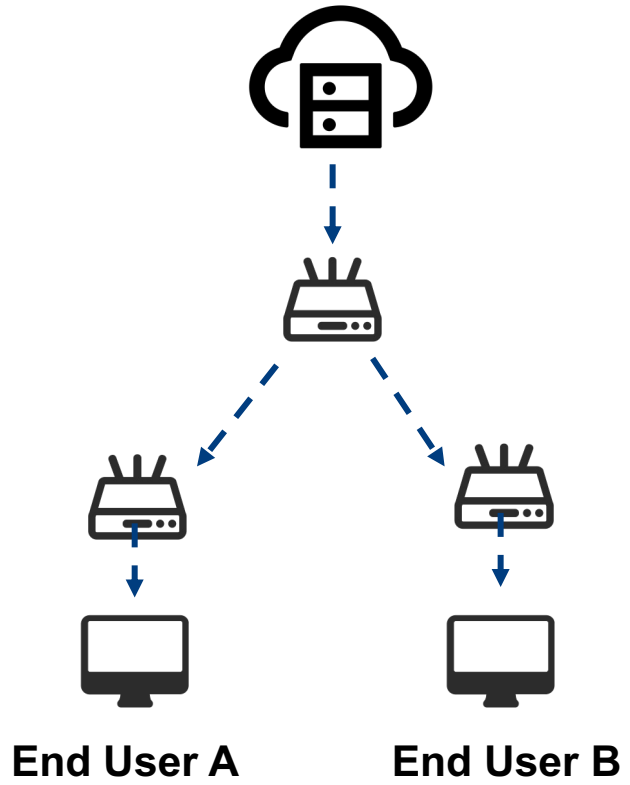

Path difference = (3 – 1) / (2 + 2) = 0.5
(discard last-hop link)

Evaluating Similarity Assumption of Current Practices

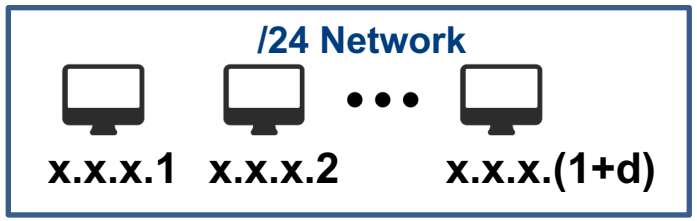

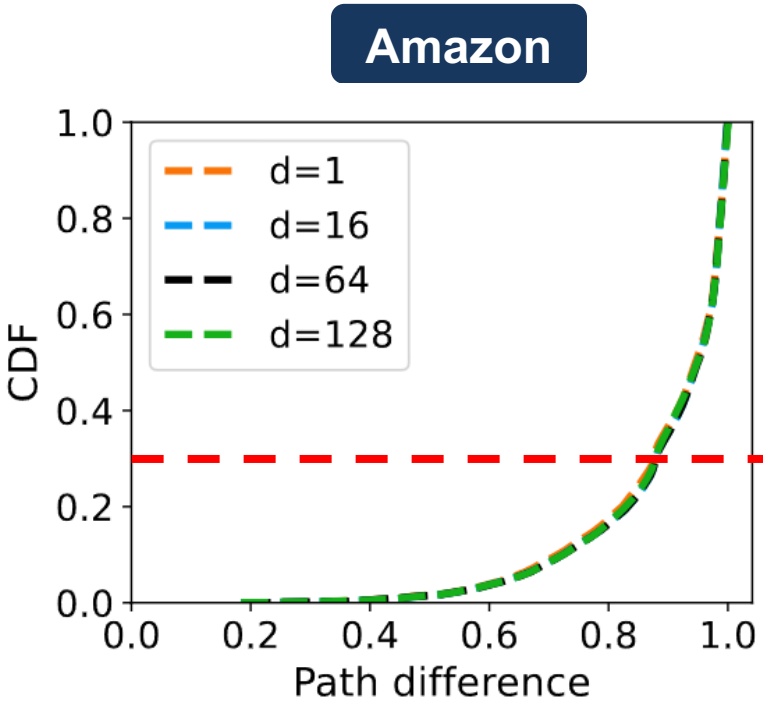

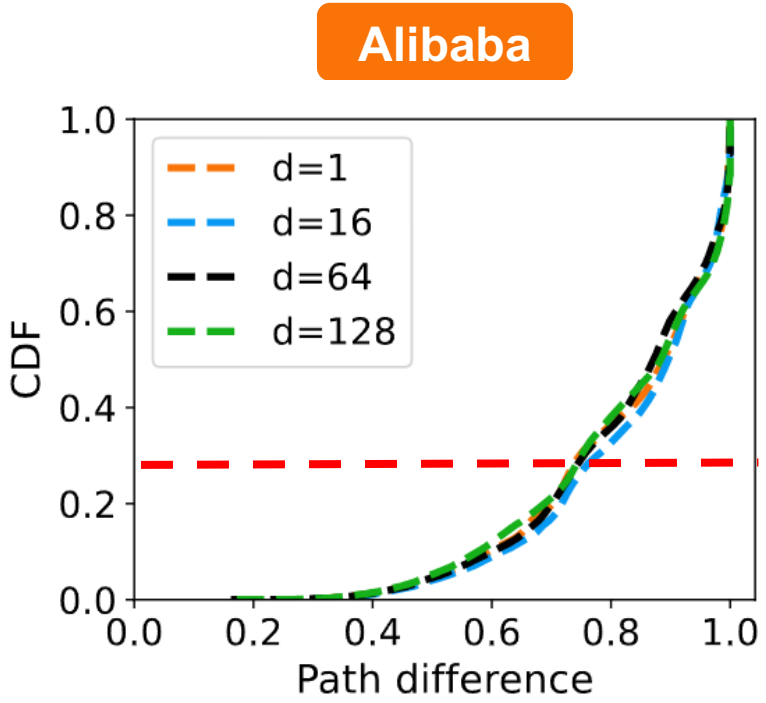

Evaluating Similarity Assumption of Current Practices

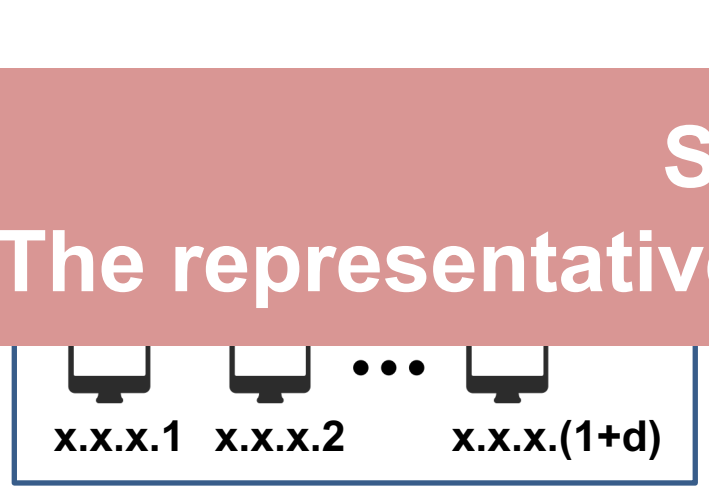

Amazon

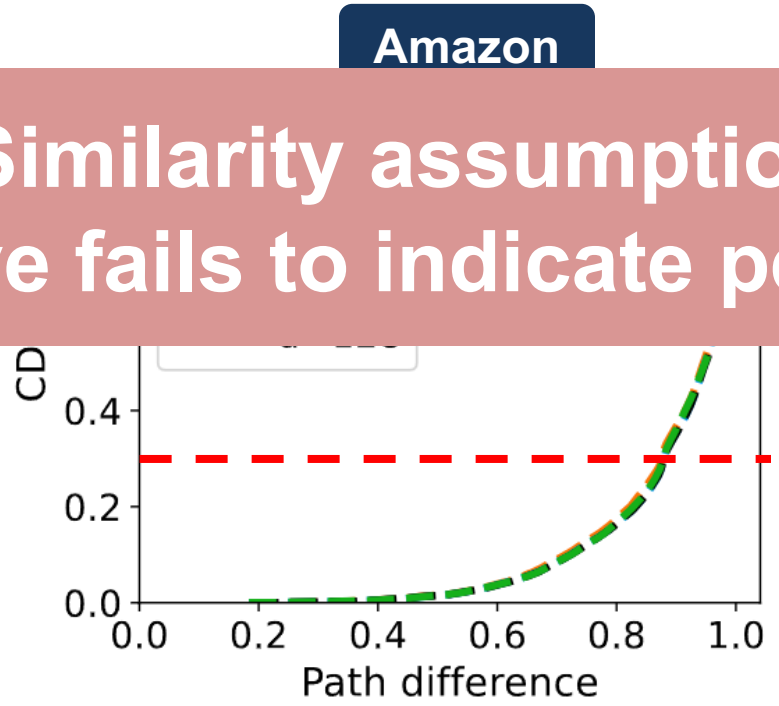

Alibaba

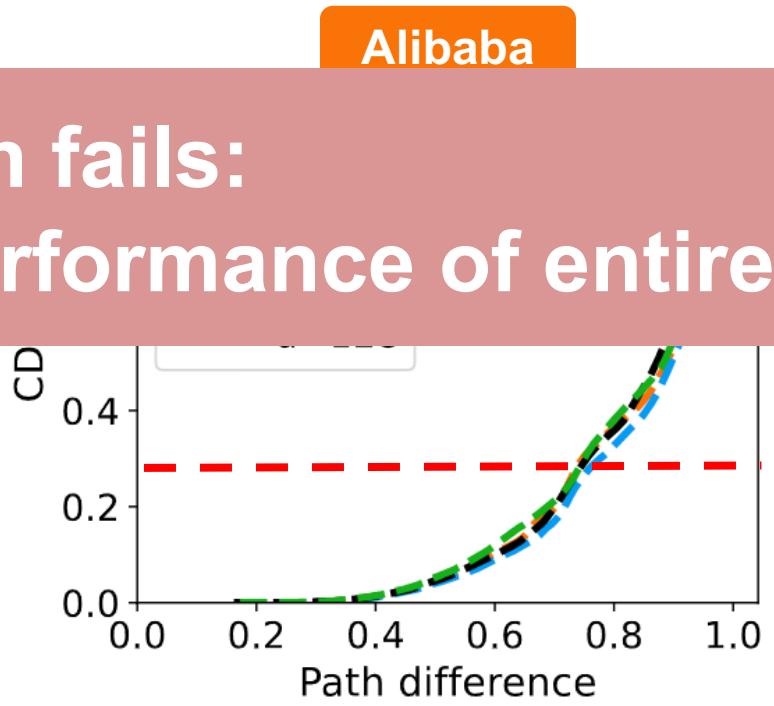

Similarity assumption fails:
The representative fails to indicate performance of entire /24

Evaluating Coverage Assumption of Current Practices

Current Practice: Selecting .1 addr of every /24 as representative

$$\text{Link coverage} = \frac{\#Covered\ links}{\#Total\ links}$$

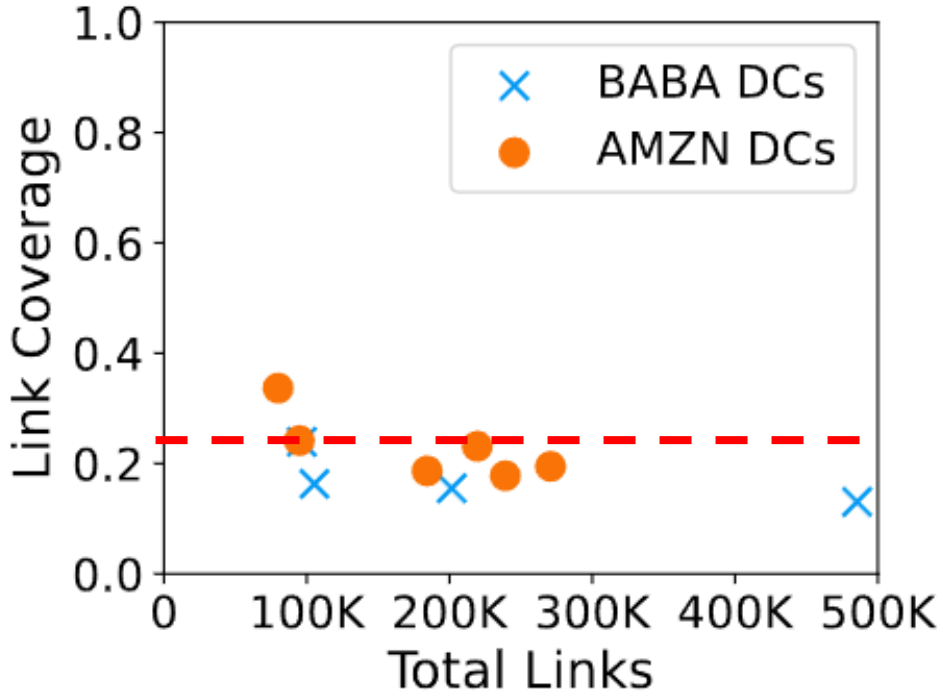

Evaluating Coverage Assumption of Current Practices

Current Practice: Selecting .1 addr of every /24 as representative

Coverage assumption fails:
Current practices leave 70% of links unwatched

$$\text{Link coverage} = \frac{\#\text{Observed links}}{\#\text{Total links}}$$

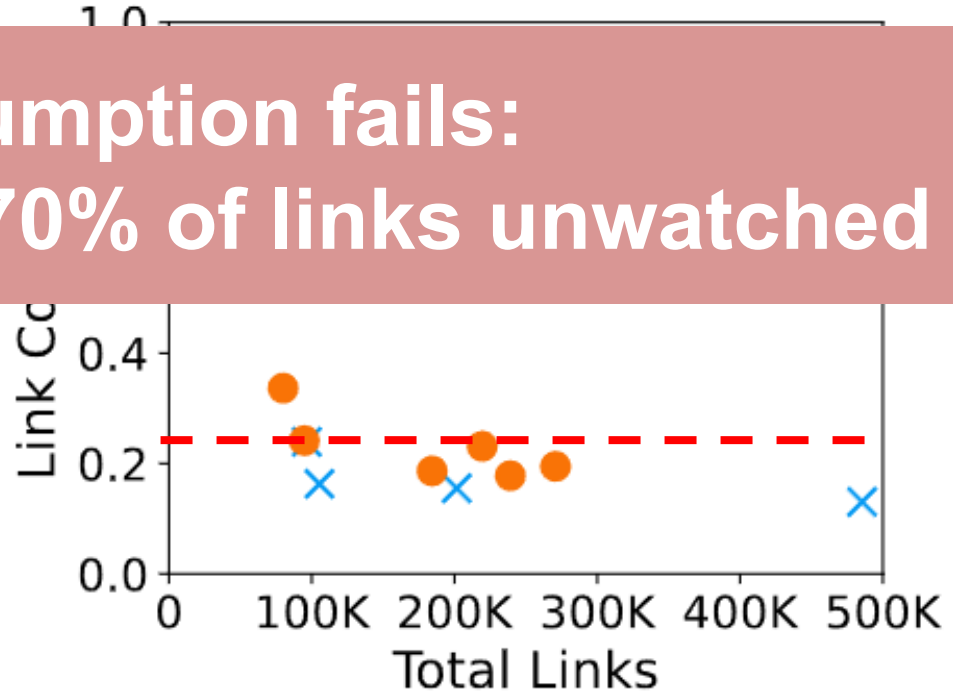

Evaluating Coverage Assumption of Current Practices

Current Practice: Selecting .1 addr of every /24 as representative

Events occurring to 70% of links will not be observed by current practices

$$\text{Link coverage} = \frac{}{\#Total\ links}$$

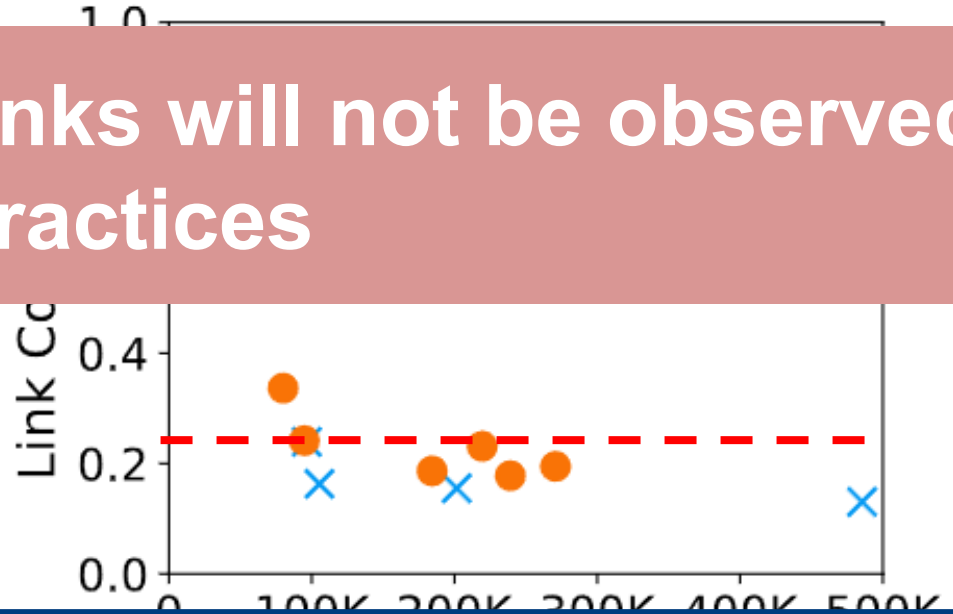

Then, how to improve link coverage?

Traditional wisdom to boost link coverage

Traditional wisdom: Increase the granularity of monitoring.

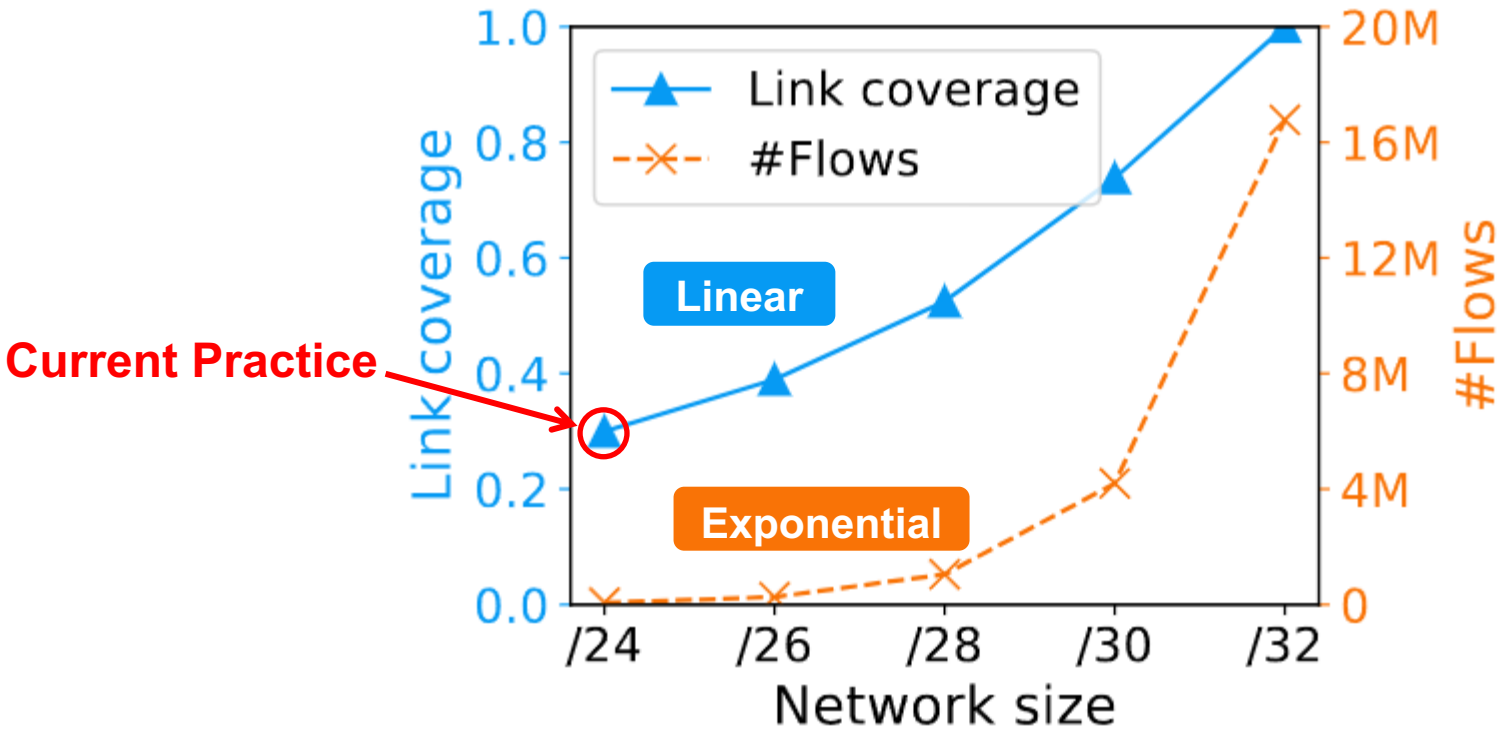

Traditional wisdom to boost link coverage

Traditional wisdom: Increase the granularity of monitoring.

Boosting link coverage by increasing granularity is not scalable

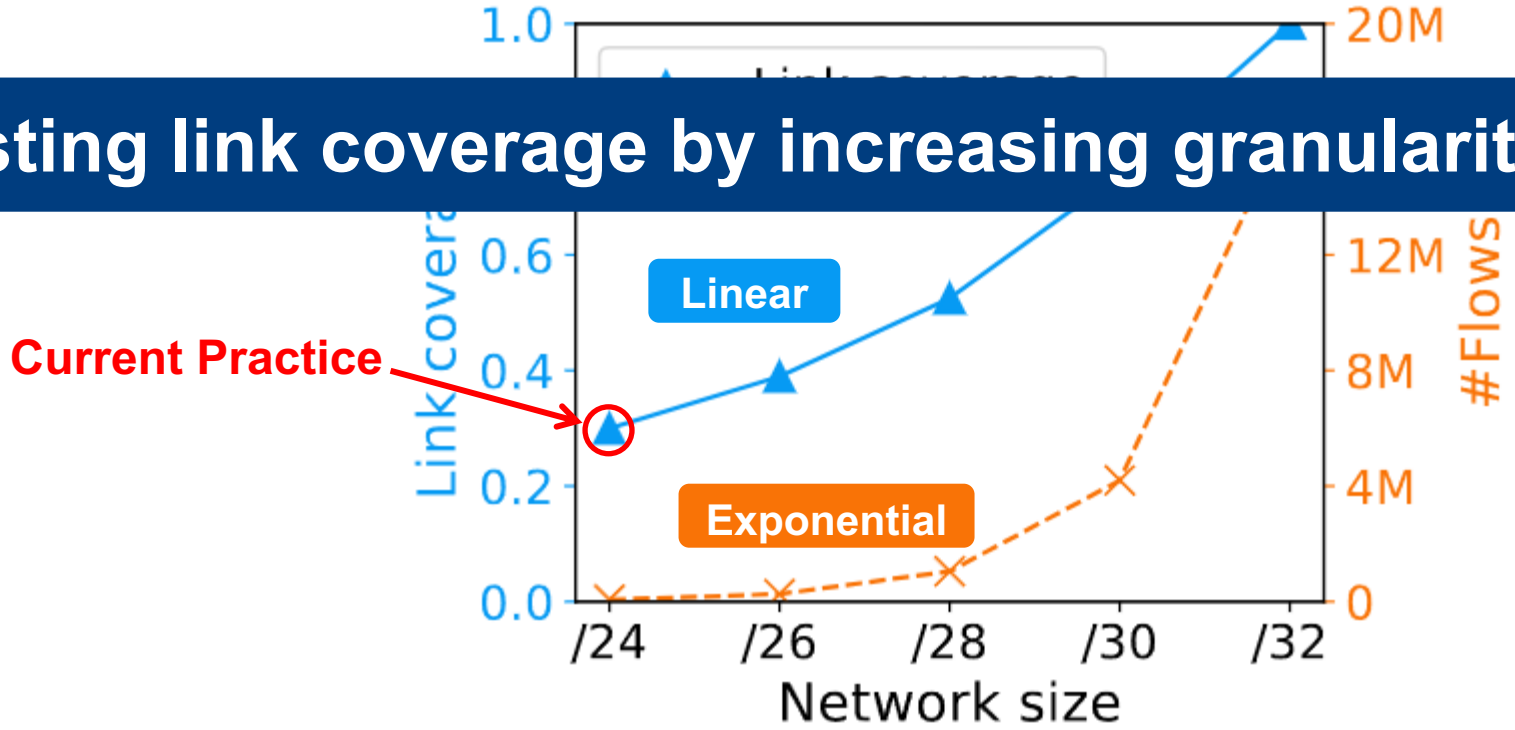

End-to-End Approach To High Link Coverage

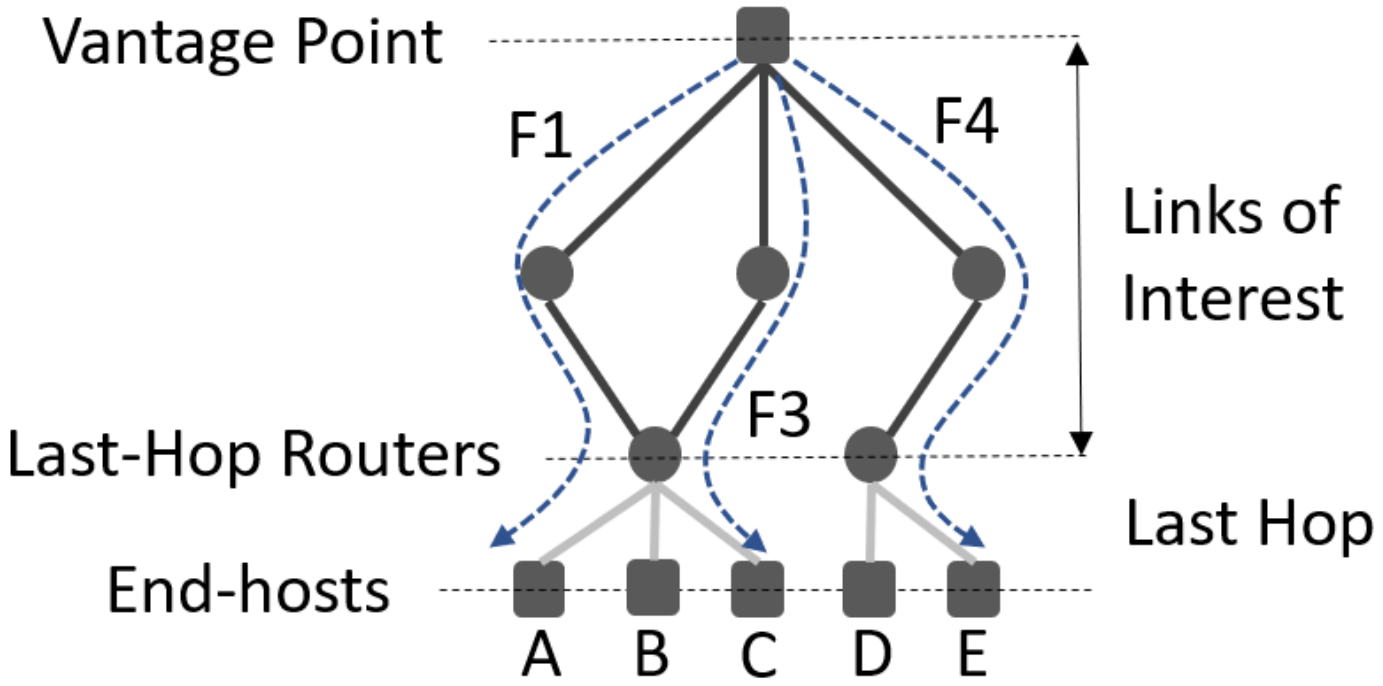

By monitoring end-hosts A,C,E, we can coverage all visible links

Why Our End-to-End Approach is More Scalable?

Normalized link count = $\frac{\text{\# Total links}}{\text{\#/24s}}$

Links scale much slower than network size

Amazon

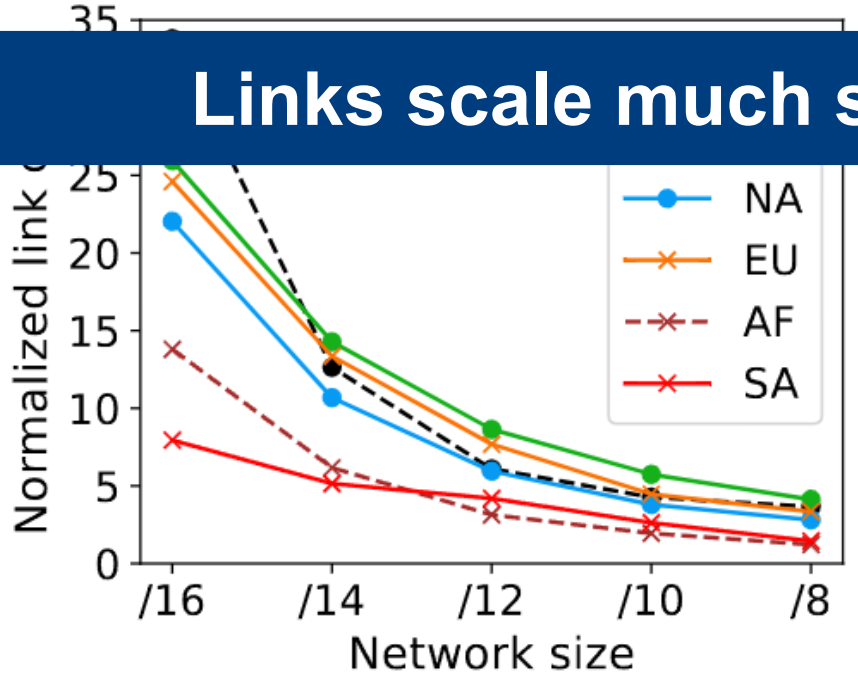

Alibaba

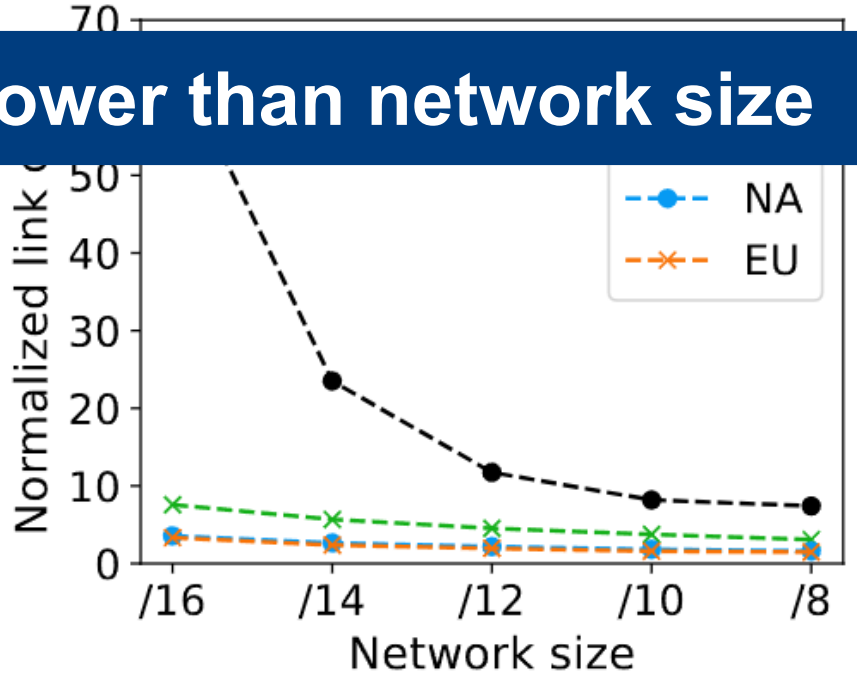

Why Our End-to-End Approach is More Scalable?

$$\text{Normalized link count} = \frac{\#\ \text{Total links}}{\#\ /24\text{s}}$$

Amazon

Alibaba

Can we achieve high link coverage by carefully selecting targets?

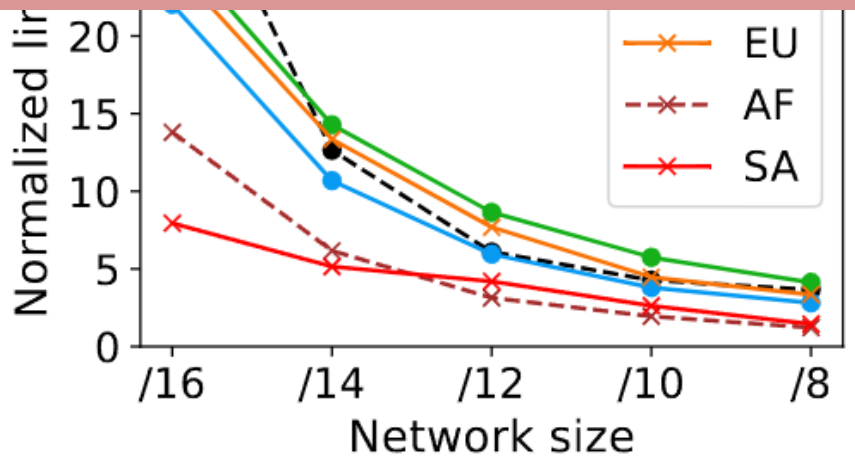

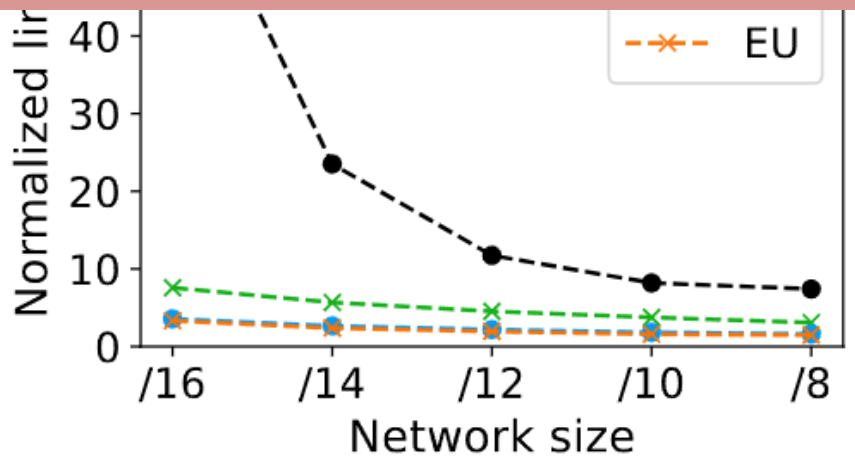

A Greedy End-to-End Approach

Always choose targets contribute most new links

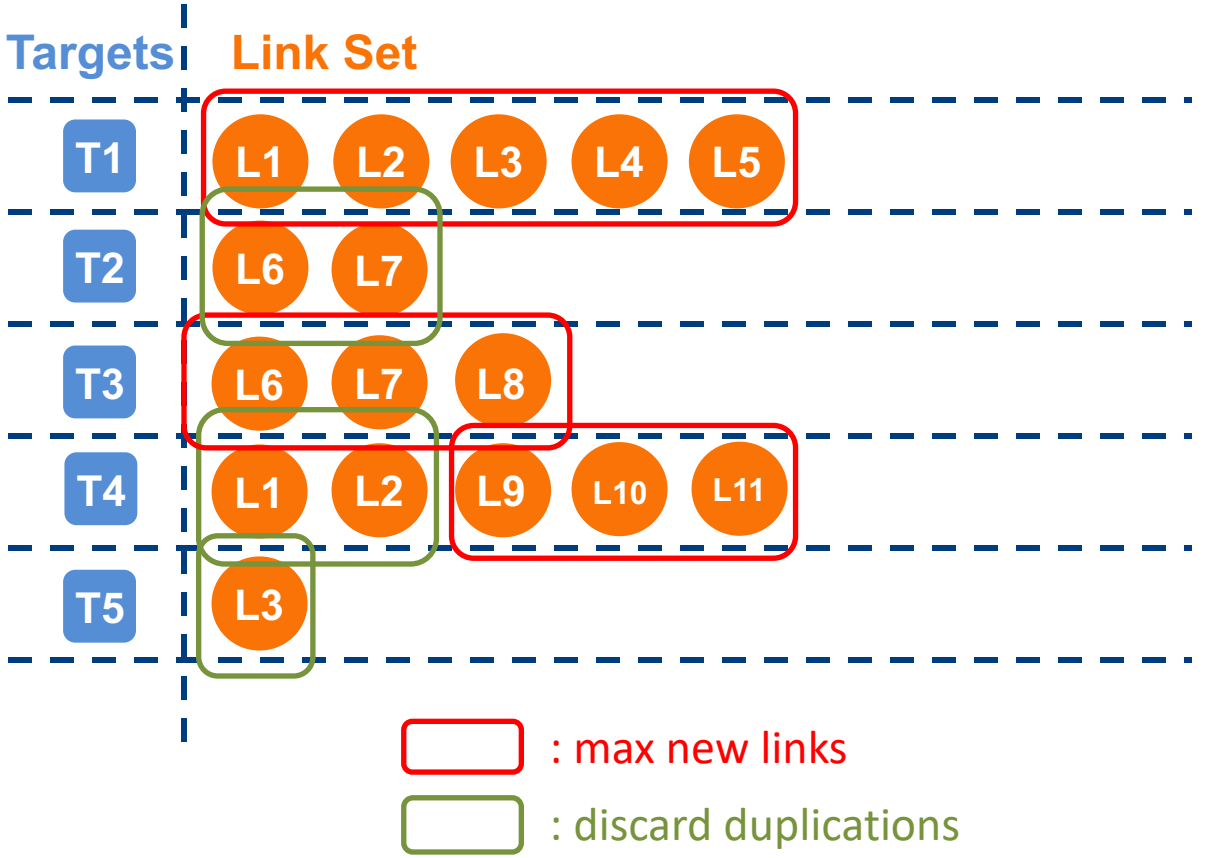

Evaluating the Greedy End-to-End Approach

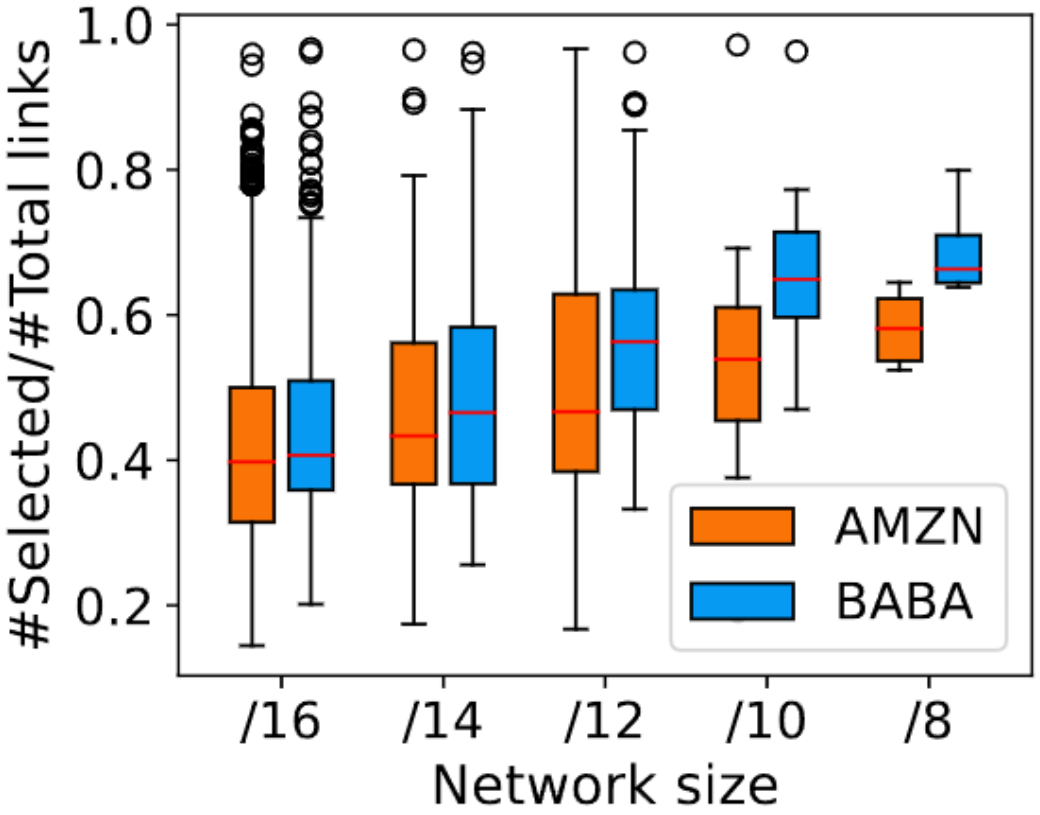

- More difficult to achieve full coverage for larger network
- For /8, monitoring *x* links only requires probing *0.6*x* targets

Trade-off between Coverage & Overhead

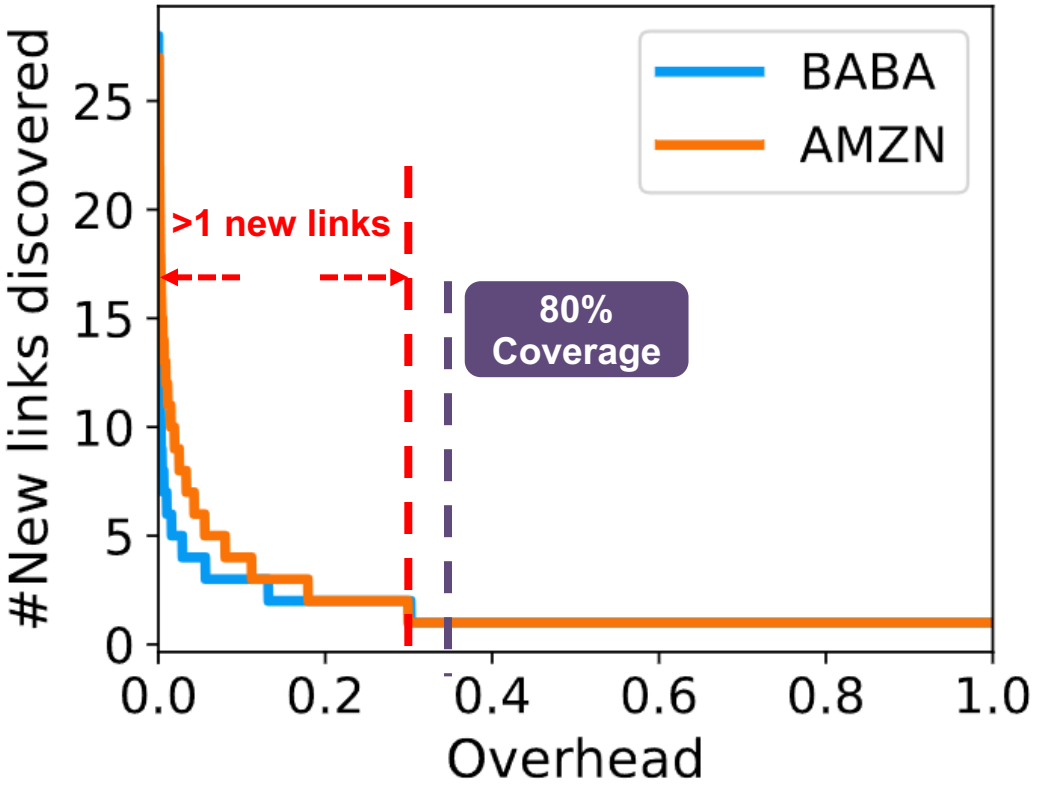

Trade-off

- Only the first 30% of overhead can discover >1 new links

- 1/3 overhead for 80% coverage

Takeaways

- Current practices **fail to monitor the changes of a majority of links** in the Internet, leaving critical links unwatched.

- **High link coverage can be achieved** by carefully selecting probing targets with reasonable overhead.

- Our dataset is published at ***https://github.com/SJTU-NMS-Lab/APNet23***


JOINT INSTITUTE
交大密西根学院

Thank You!

Q & A

Future Directions

- IPv6 exploration
- Fast start without long-time data collection
- Real-time detection on link failure/congestion

Internet Monitoring: Passive vs Active

Passive Measurements

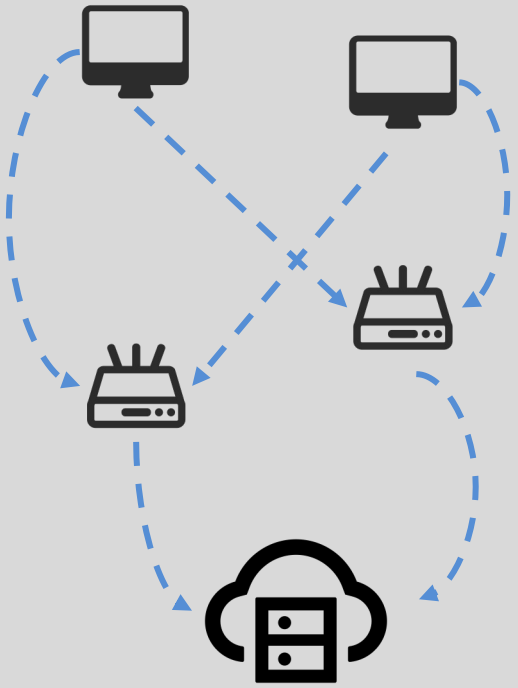

Client Traffic

Active Measurements

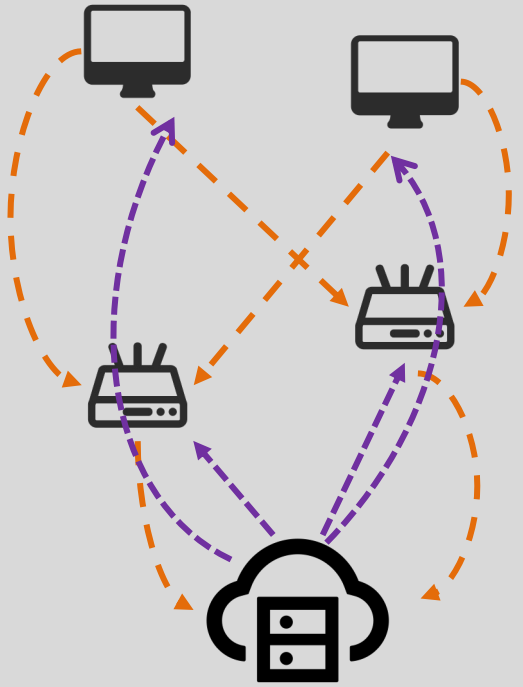

Probes

Responses

Our Work

Our Contributions

- Evaluate the link coverage of two rule-of-thumb practices for scalable Internet measurement from a cloud-centric view
- Evaluate the predictability of performance for client flows to the same /24s
- Propose to achieve high-coverage monitoring with an end-to-end approach
- Estimate the overhead for high-coverage monitoring

Dateset

Ground Truth

- Tool: D-Miner [Vermeulen et al, NSDI' 20]
- Divide /8 into /16s
- Send two back-to-back scans at 100,000pps
- Goal: Find ALL visible links at confidence level of 99%(95% for one scan)

Dateset

Ground Truth

- Tool: D-Miner [Vermeulen et al, NSDI' 20]
- Divide /8 into /16s
- Send two back-to-back scans at 100,000pps
- Goal: Find ALL visible links at confidence level of 99%(95% for one scan)

Random Flow

- Tool: Modified Zmap
- Setup
  - src port: 80
  - dest port: Random
  - TTL: 1~32
- Goal: Simulate the downstream traffic flows from DCs to clients

Dateset

Ground Truth

- Tool: D-Miner [Vermeulen et al, NSDI' 20]
- Divide /8 into /16s
- Send two back-to-back scans at 100,000pps
- Goal: Find ALL visible links at confidence level of 99%(95% for one scan)

Random Flow

- Tool: Modified Zmap
- Setup
  - src port: 80
  - dest port: Random
  - TTL: 1~32
- Goal: Simulate the downstream traffic flows from DCs to clients

Full-coverage Flow

- Tool: Modified Zmap
- Setup
  - src port & dest port: same as D-Miner
- Goal: Cover most of the visible links by carefully selecting client flows

Uncovered links cause poor visibility to the Internet

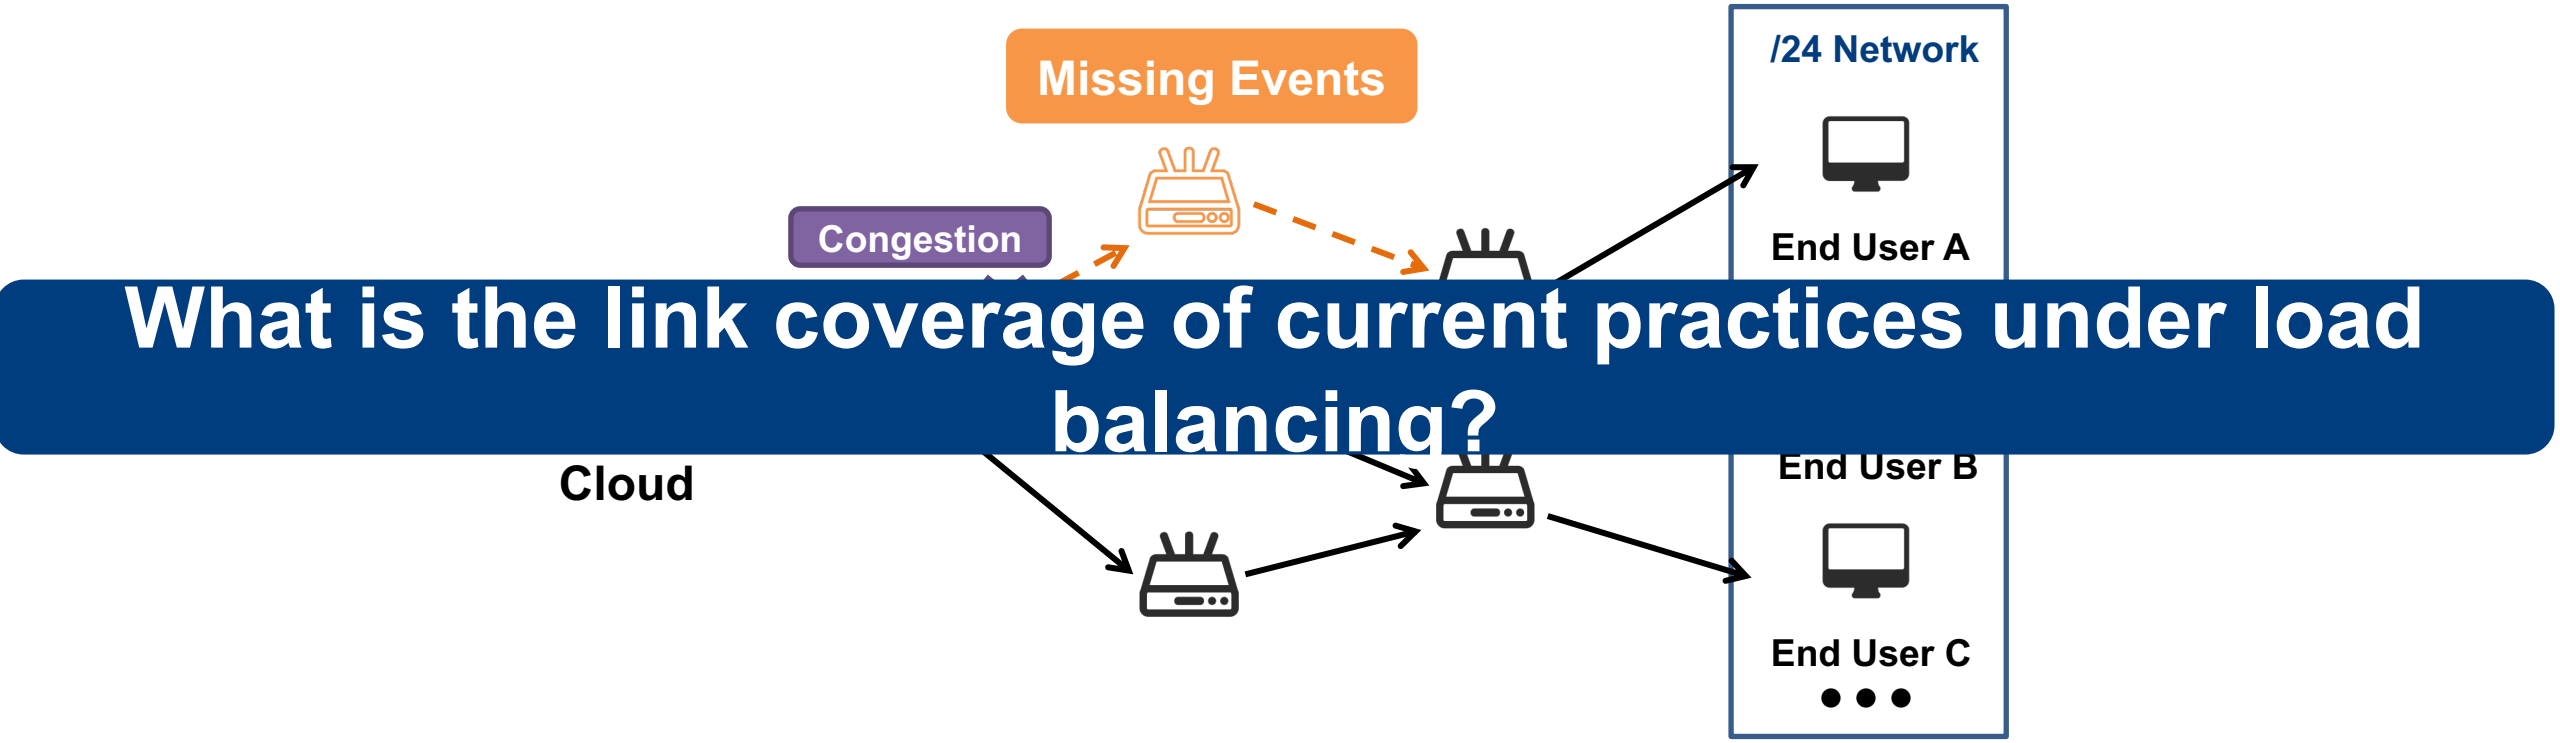

What is the link coverage of current practices under load balancing?

这个图可以用，再加一些说明文字表述出标题的这个意思

Load Balancing Challenges Coverage Assumption