

# Layer-1 Informed Internet Topology Measurement

Ram Durairajan\*, Joel Sommers^, Paul Barford\*

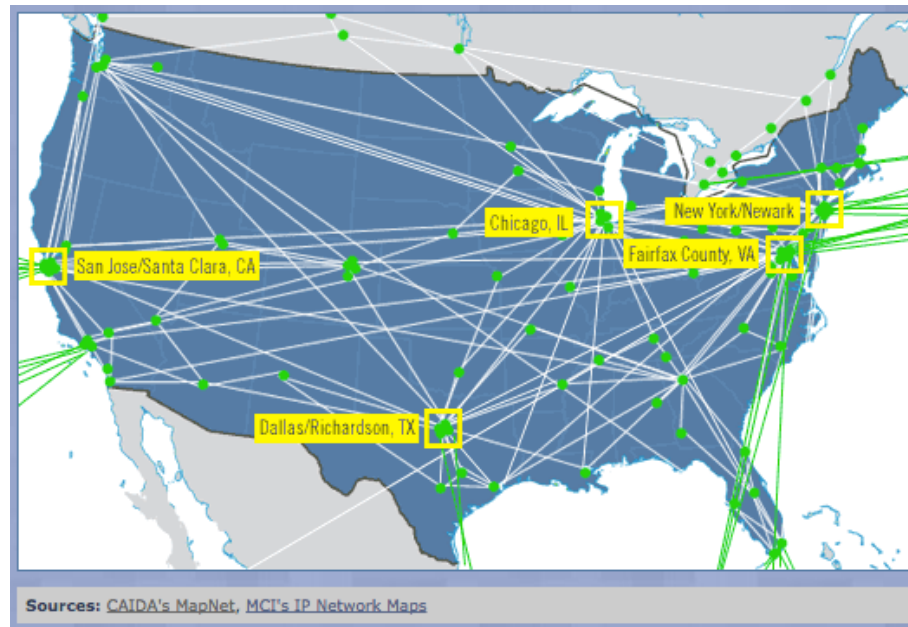
\*University of Wisconsin - Madison      ^Colgate University

# Introduction

- Understanding Internet topology is important
  - Informs performance, security, risk, etc.
- Internet topology mapping is fraught with challenges
  - Huge size and distributed ownership
  - Always in a state of flux

# Existing Approaches

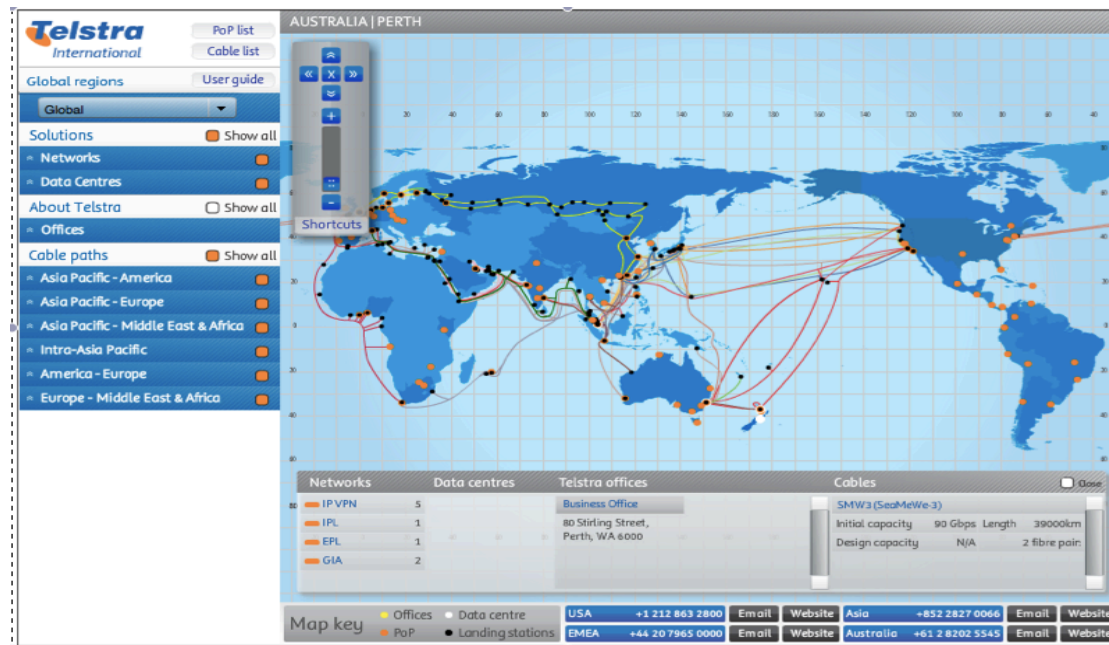
- TTL-limited layer 3 traceroute-like probes
  - Rely on location hints in domain names
  - E.g., CAIDA's Ark, Rocketfuel



Network-layer maps

# Existing Approaches (cont.)

- Search based
  - Maps available at ISP's website
  - E.g., Internet Atlas, Internet Topology Zoo



Physical maps

# Topology measurement challenges

- Problems with TTL-based approaches
  - Management policies/Objectives of providers
  - Lack of visibility of lower layers
- Problems with Search-based approaches
  - ISP acquisition/merge
  - May not be up to date or complete

# Research questions

Can physical maps be used to guide and reinforce the process of collecting network-layer data?

- How do physical maps compare to and contrast with network-layer maps?
  - Atlas vs. Ark comparison study
- How can probe methods be improved to reveal a larger portion of physical infrastructure?
  - POPsicle probing heuristic

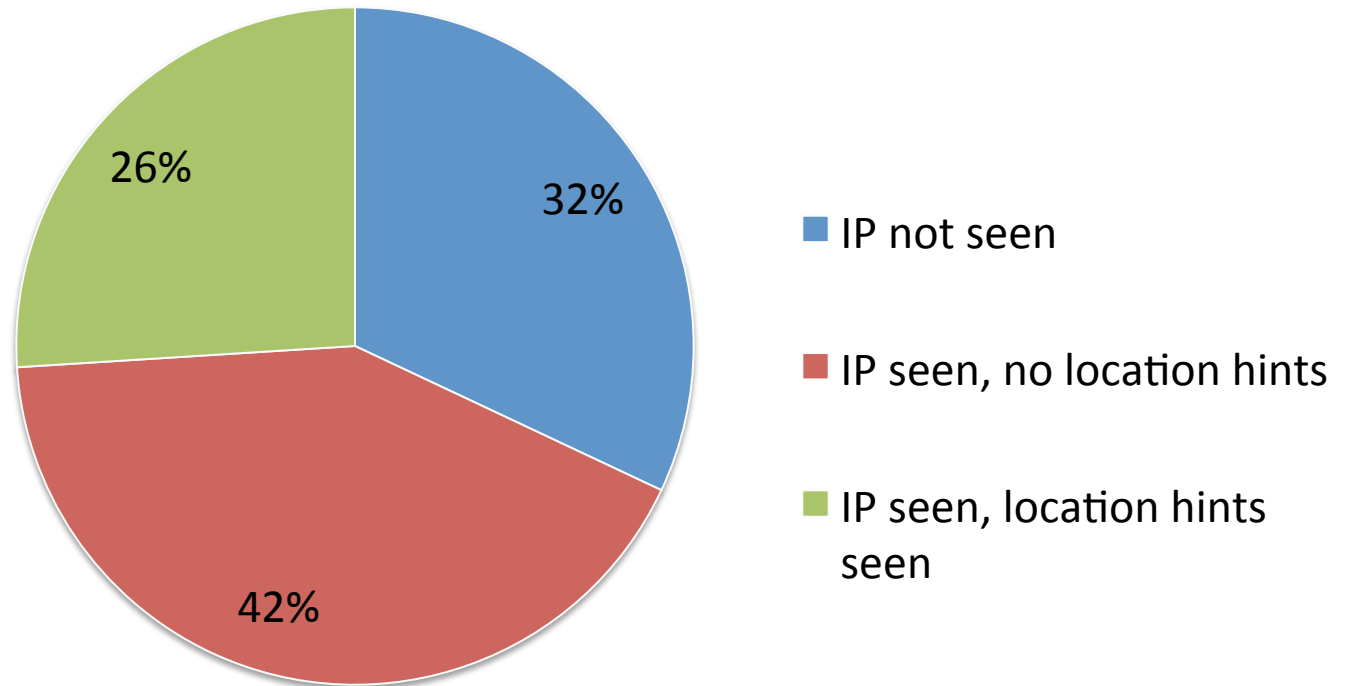
How do physical maps compare to and contrast with network-layer maps?

# Targets for comparison

- We consider 50 networks with footprint in North America
- Atlas
  - 7 Tier-I and 43 regional ISPs
  - 2507 POPs and 3477 links
- Ark
  - Use DNS data and traceroute data
  - PathAudit (Chabarek et al., HotPlanet '13) to decode location hints
    - E.g., for A.B.C.LAX2.D.NET, location code is LAX




# Physical vs. network maps – results I



More nodes and links in physical maps.

# Physical vs. network maps – results 2

Number of probes sent across Internet Service Providers



- Sampling bias in network topology measurements (Shavitt *et. al.*, IEEE Infocom 2009)

# Physical vs. network maps – results 3

- Network map utility
  - 448 distinct networks in North America
    - Greater than physical maps in (worldwide) Atlas repository!
  - Dynamic properties

Results from network-layer maps can be used as guidance for searching physical maps

# Implications

- Differences suggest opportunities for reinforcement
  - Networks in network-layer data
    - Clues for searching new maps
    - Engineering problem
  - Networks in physical data
    - Targets for additional probing
    - Calls for a coordinated topology mapping approach

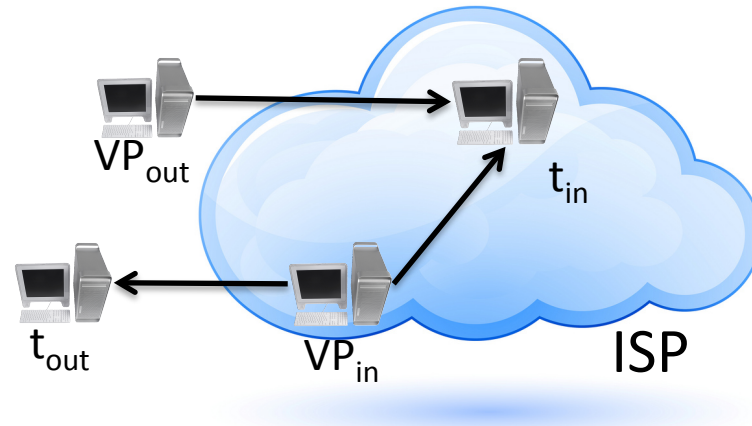
How can layer 3 probe campaigns be designed to reveal a larger portion of physical infrastructure?

# Considerations for targeted probes


- Source-destination selection
  - Vantage point (probing source or VP) and destination selection
    - Internal to an ISP or external to an ISP?
- Scalability
  - Exploit IXPs to aid in node identification
  - Vantage points for multiple networks?
    - Due to layer 2 connectivity

# Source-destination selection

- Leverage publicly available vantage points
  - Planetlab, looking glass and traceroute servers
- Three modalities
  - $VP_{out}$  to  $t_{in}$
  - $VP_{in}$  to  $t_{out}$
  - $VP_{in}$  to  $t_{in}$
- Source-destination selection based on geographical proximity
- 25 ISPs containing 596 target POPs



# Effects of source-destination selection



Sources and destinations within the same AS based on geographic proximity

- Effects of routing
  - $VP_{in}$  to  $t_{in}$ 
    - Greater diversity, more info. on paths, flexible routing
  - $VP_{in}$  to  $t_{out}$  and  $VP_{out}$  to  $t_{in}$ 
    - Interdomain routing



# Scaling perspective with IXPs

IXPs could be the starting point for comprehensive mapping of physical infrastructure

- Enormous amount peering at IXPs
- VPs co-located with IXPs
  - 14 out of 65 have co-located VPs
  - Unique ISPs that peer at 14 IXPs is 625 (from PeeringDB)
  - So, 625 ISPs from these 14 IXPs alone

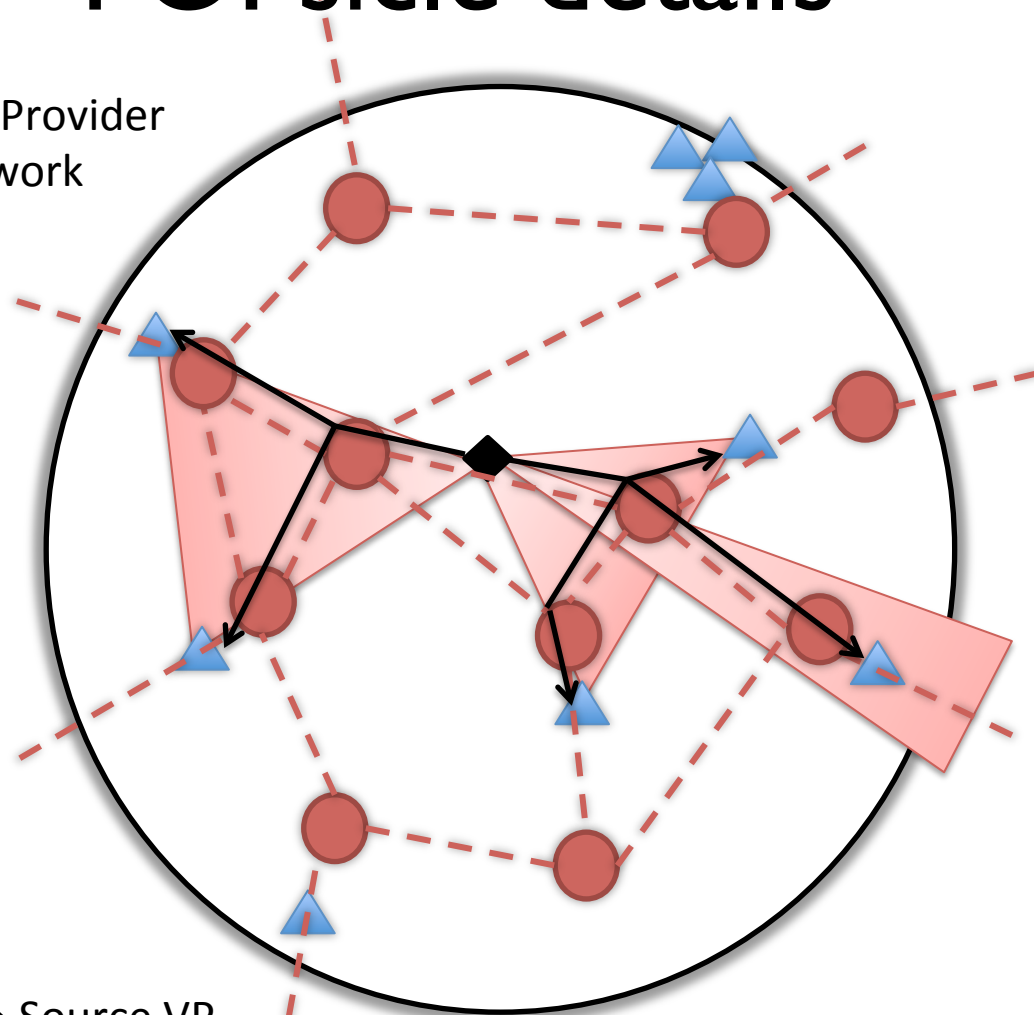
# Pulling it all together

- Goal: use physical maps to enhance network-layer node identification
- Sources:
  - VP located within a target AS
  - VP co-located with IXPs offers broader perspective
- Destinations
  - Send probes toward a target with a known geographic location based on physical map

POPsicle: Probing heuristic based on these insights

# POPsicle details

Service Provider  
Network



- ◆ Source VP
- ▲ Destination VP
- Target POPs

- - - Links
- Probe path

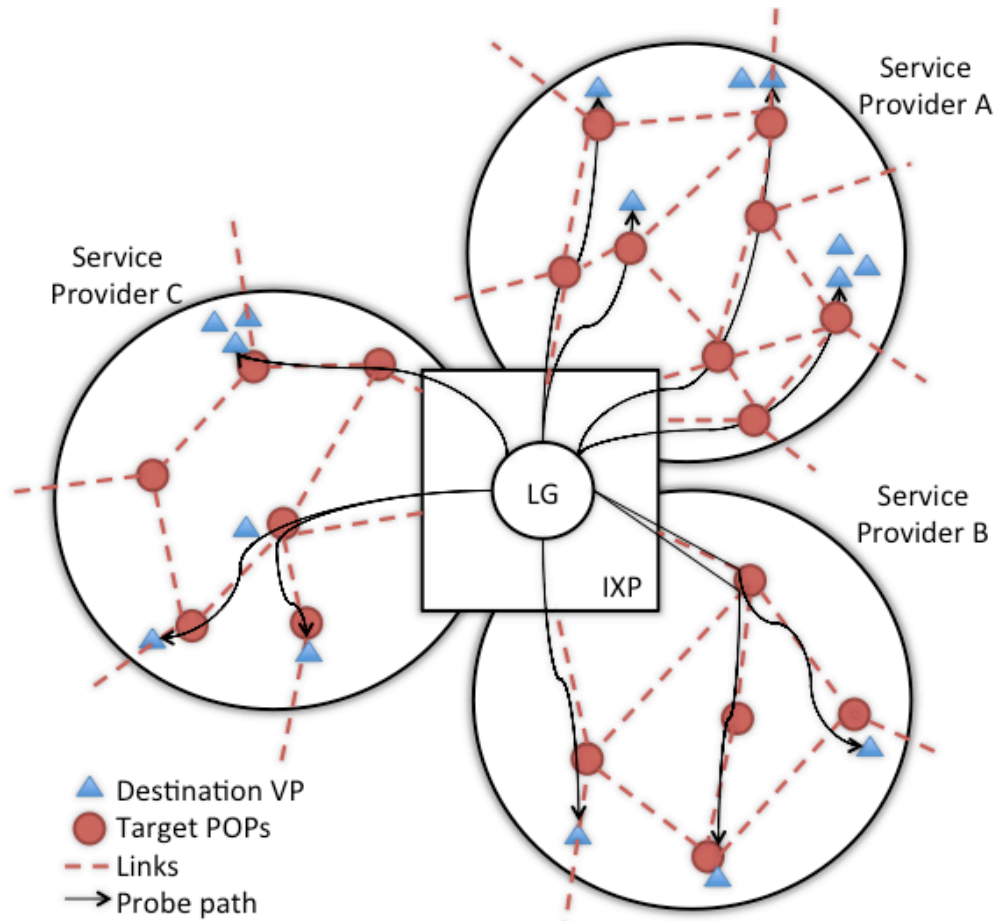
rkrish@cs.wisc.edu

# POPsize - results

- 30 looking glass servers from Atlas
  - server co-located with an IXP
  - ground truth available

	POPs	Datacenters	DNS	NTP	IXPs	Total
POPsize-based probing	149	487	9	627	37	1309
General probing	143	315	1	55	25	539
Ground truth	244	641	13	827	65	1790
Improvement	1.04x	1.54x	9x	11.4x	1.48x	2.42x

# Multiplexing VPs at IXPs



# Multiplexing VPs at IXPs

ISP	POPsize	Ground Truth
BTN	29	29
HE	24	24
Internet2	10	10
Steadfast.net	3	3
Nexicom	9	9
HopOne	3	3
Indiana Gigapop	2	2
MOREnet	4	4
Atlantic Metro	9	12
PaeTec	54	61

# Summary

- First-of-its-kind comparison of physical vs. network-layer maps
- Source-destination pairs within the same AS reveals most physical infrastructure
- POPsicle-based probing identifies 2.4x additional nodes
- IXPs can aid in broadening perspective
- Deployed and demonstrated POPsicle in a real IXP setting

Thank you!  
Questions?

[www.internetatlas.org](http://www.internetatlas.org)