Layer 1-Informed Internet Topology Measurement

Presenter: Markus Ansorge

Technical University of Munich

Munich, 1. June 2017

Autors: Ramakrishnan Durairajan (University of Winsconsin-Madison)

Joel Sommers (Colgate University)

Paul Barford (University of Winsconsin Madison)

Outline

Introduction

State of the art

- Brute-Force IP Search
- Layer 3 TTL-limited probing (= traceroute)

Goal: Improving completeness using layer 1 data

Motivation: Possibility for improved

- Performance
- Security
- Robustness
- Etc.

Datasets

- Setting
  - Time period: September 2011 to March 2013
  - Geo-location: North America

- Dataset: Internet Atlas
  - Map of the physical-layer internet
  - Based on published ISP information

- Dataset: CAIDA's Archipelago (Ark)
  - Map of the network-layer internet
  - Based on large-scale tracerouting

Mapping IP-Addresses to Physical Locations

- Basic Idea: Utilizing location hints in DNS

- Algorithm:
  - Get DNS from IP-address
  - Extract location code using regular expression patterns
  - Retrieve physical location via mapping codes
  - (Classify location into different AS via mapping service)

Mapping IP-Addresses to Physical Locations

- Result:

Total traceroutes processed	2,674,959,041
Number of unique interface IP addresses	14,593,457
Number of unique ASes	31,055
Valid DNS entries found	6,936,146
No associated DNS name found	7,657,311
DNS entries with location hints	704,935
Number of ASes with at least one geographically identifiable interface address	4,135

- Problems:
  - Multiple POPs per city
  - No location hints
  - No AS mapping entry

Comparison between Physical- and Network-Layer Map

- Scale of data: 50 networks
- Findings:

ISP	Physical		Network-layer		Nodes			Links			N_{Index}
	Nodes	Links	Nodes	Links	Intersection	Only in P	Only in N	Intersection	Only in P	Only in N	
AT&T	25	57	39	72	25	0	14	51	6	21	100
Cogent	186	245	122	172	122	64	0	171	74	1	63
NTT	47	216	65	229	47	0	18	189	27	40	57
Tinet	122	132	64	79	57	65	7	79	53	0	37
Sprint	63	102	67	108	63	0	4	98	4	10	54
Level3	240	336	129	237	129	111	0	237	99	0	63
Tata	69	111	0	0	0	69	0	0	111	0	40
Abiline	11	14	8	13	8	3	0	13	1	0	100
Ans	18	25	0	0	0	18	0	0	25	0	94
ATMnet	21	22	0	0	0	21	0	0	22	0	100
Bandcon	22	28	14	22	14	8	0	22	6	0	100
BBNPlanet	27	28	0	0	0	27	0	0	28	0	100
BellCanada	48	65	22	0	22	26	0	0	65	0	56
BellSouth	50	66	0	0	0	50	0	0	66	0	76
BTNorthAmerica	33	76	0	0	0	33	0	0	76	0	85
CompuServe	11	17	0	0	0	11	0	0	17	0	100
DarkStrand	28	31	0	0	0	28	0	0	31	0	96
DataXchange	6	11	0	0	0	6	0	0	11	0	100

Comparison between Physical- and Network-Layer Map

Reason for missing data:

- No location hints
- Blocking traceroute
- Tunneling protocols
- Interface configured with third party IP-addresses

=> Only 13 network comparable

Routing's Source and Destination Selection Effects

Study

- Based on ISP assignment
- Types:
  - $S_{out} \rightarrow D_{in}$
  - $S_{in} \rightarrow D_{out}$
  - $S_{in} \rightarrow D_{in}$

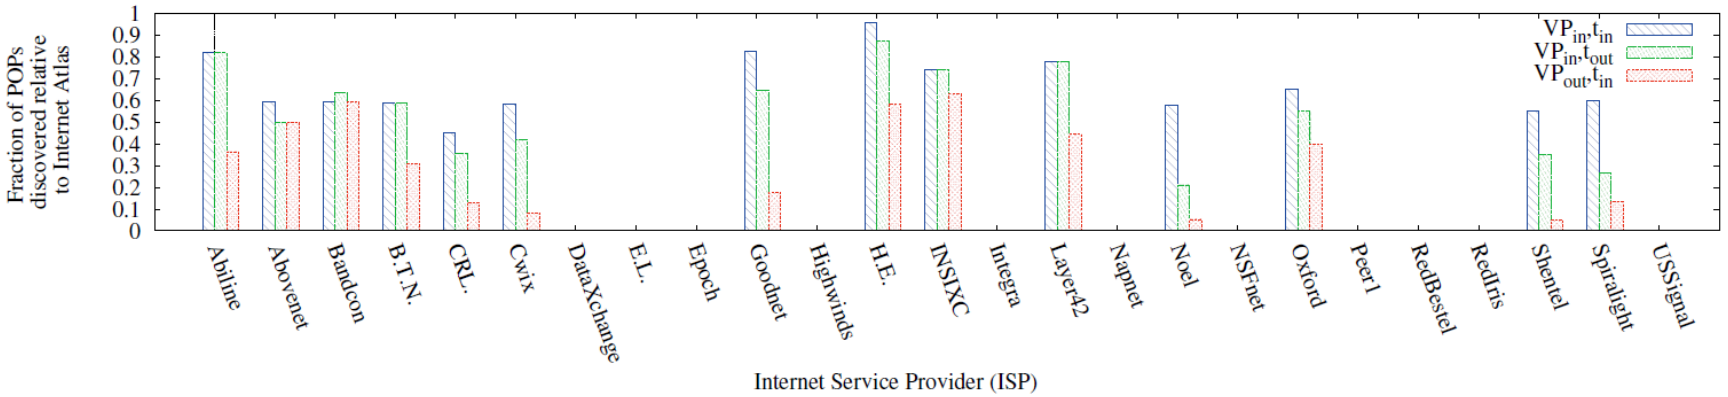

=> Intradomain routing preferable

POPsicle

- Layer 3 probing system
- Purpose-built system
  - Utilizes layer 1 knowledge
  - Deployment: Extension of generalized systems

POPsicle Algorithm

- Input
  - Source VPs
  - Target POPs

1. Traceroute between geographically close VPs
2. Route contains POP ?
   - Finished
   - Go to Step 1

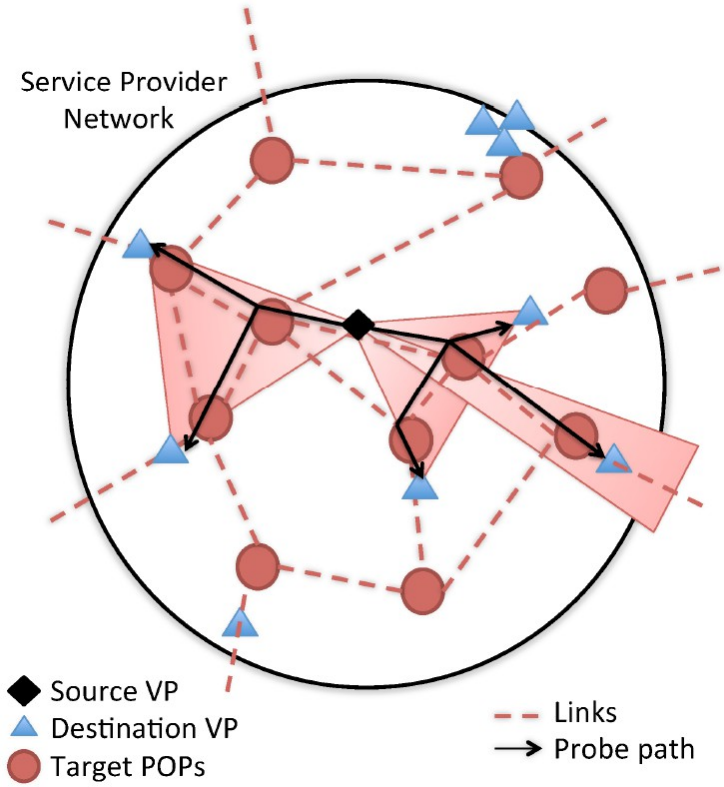

POPsicle Evaluation

- Originally 30 ISP networks planed
- Only 13 suitable

	POPsicle	Atlas	Ark	Rocketfuel
Abovenet	13	22	13	13
BellCanada	34	48	30	29
Centauri	7	14	3	—
Cyberverse	2	2	2	—
Data102	2	2	2	—
HopOne	4	4	4	—
HE	23	24	23	8
Inerail	3	25	3	—
Internet2	10	10	10	10
Interserver.net	2	2	1	—
Steadfast.net	3	3	3	—
Towardex	7	8	6	—
XO	42	80	42	39

POPsicle Evaluation

Results from mapping infrastructural nodes

	POPs (for 13 ISPs)	Datacenters	DNS Servers	NTP Servers	IXPs	Total locations
POPsicle	149	487	9	627	37	1309
Ark	143	315	1	55	25	539
Atlas	244	641	13	827	65	1790
POPsicle compared to Atlas	61.07%	75.98%	69.23%	75.82%	56.92%	73.13%
Ark compared to Atlas	54.60%	49.14%	7.69%	6.65%	38.46%	30.11%
Improvement	1.04x	1.54x	9x	11.40x	1.48x	2.42x

POPsicle Evaluation

Special case: Deployment at Equinix Chicago IXP

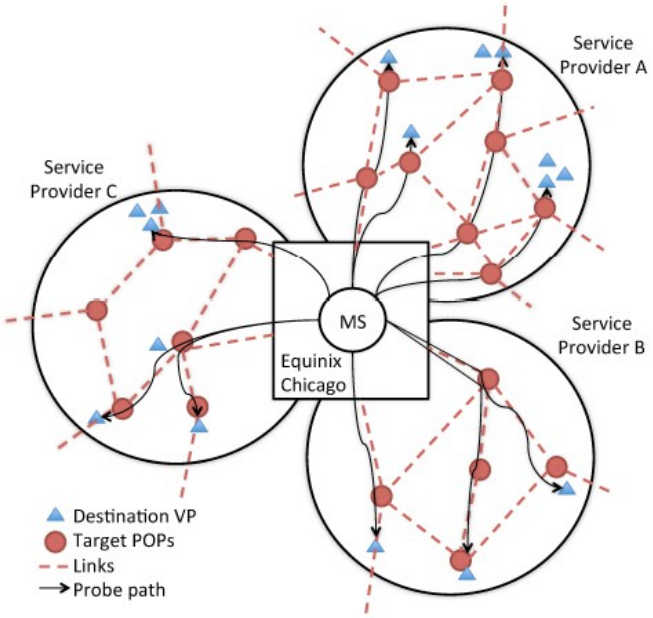

ISP Name	POPsicle	Atlas	Ark
BTN	29	29	28
HE	24	24	23
Internet2	10	10	10
PaeTec	54	61	54
Nexicom	9	9	9
HopOne	3	3	3
Indiana Gigapop	2	2	2
MOREnet	4	4	4
Atlantic Metro	9	12	8
Steadfast.net	3	3	3

Conclusion

- Physical maps typically reveal more nodes/links
- IXPs are great VPs
- POPsicle probing
  - Better results
  - High demands