Internet Topology Research

Matthew Luckie

WAND Network Research Group
Department of Computer Science
University of Waikato

Internet Topology

Why should we care?

- Impacts on the design and operation of routing protocols
- Understand choices in network design
- Some indication as to the efficiency of a route between two hosts

This talk

- Overview of topology measurement techniques
- Motivation to annotate captured data with reverse DNS entries
- Some results

Topology Measurement

- Playing games with traceroute(8)
  - IP topology
  - Discover peerings not visible with Routeviews
  - Observe structure of individual ASes
    - Rocketfuel: Neil Spring et al.
- Specialist projects for Internet-scale topology measurement
  - Skitter (CAIDA)
  - Archipelago + Scamper (CAIDA + WAND)
  - DIMES
  - Others; PlanetLab + ScriptRoute

Traceroute basics

- Series of TTL-limited probes
  - solicit ICMP Time Exceeded messages from routers on the path
  - ICMP message has original probe embedded, need to identify probe to know which TTL it is for
- UDP probes to high-numbered ports
  - each probe is identified by a unique destination port
- ICMP echo request probes
  - each probe is identified by ICMP sequence
- TCP SYN probes
  - each probe is identified by IP-ID

scamper

- Parallelised Internet measurement
  - takes a list of IP addresses
  - traceroute in parallel as required to fill a specified packets-per-second rate
- http://www.wand.net.nz/scamper/
  - code freely available

CAIDA Archipelago

- Measurement infrastructure distributed across the globe
- Probe with scamper for Internet topology
- Results are centrally collected and made available for further analysis

Challenges

- Traffic can look like scanning
- Redundant probing
- Load-balancing routers may break validity of traceroute output
- Translating IP topology to router topology

traceroute and abuse

- UDP probes to a series of high-numbered ports may appear as port scanning
  - Particularly when a destination or middlebox silently discards probes
  - UDP is the default traceroute method
- ICMP PING ATTACK!
- TCP probes to routers may be monitored by operators

Redundant probing

- The first few hops from a source are likely to be the same to any destination
- The last few hops to a destination are likely to be the same from most sources

- Solution: technique known as Doubletree
  - B. Donnet, T. Friedman, et al.
  - distributed measurement systems build a shared topology
  - systems begin probing somewhere in the middle where they are more likely to discover new links

Load Balancing

- Traditional UDP traceroute uses a different destination port to identify each probe
  - Routers may load balance based on 5-Tuple (src,dst IP / src,dst port / IP protocol)
  - May result in false IP links being reported
- Solution: paris traceroute
  - Augustin et al.
  - Identify probes using different UDP checksum value
  - Their recent work presents techniques to find all paths in a load balanced path

Router vs. IP topology

- Traceroute discovers interface IP addresses
- Routers have multiple interfaces
- Goal: IP topology to router topology
  - Resolve *router aliases*

Router Alias Resolution

- UDP probes to high-numbered ports
  - watch which ones get port unreachable from the same source IP address
  - Source: Pansiot and Grad
- Solicit responses from candidate address pairs, look for sequential IP-ID values
  - Implemented in Ally (Spring et al.)

Router Alias Resolution

- DNS, similar names for each interface may indicate they belong to the same router
  - Rocketfuel (Spring et al.)
- Analytical Alias Resolution
  - Mehmet Gunes, Kamil Sarac
  - Point-to-point links tend to be allocated out of /30 or /31
  - Two different /30 or /31 pairs observed at adjacent hops are likely to be aliases for two routers: e.g.:

192.107.171.49	130.217.2.2
130.217.2.1	192.107.171.51

DNS and Topology Discovery

- Reverse DNS entries give some indication as to the role or location of each interface on a path

DNS and Topology Discovery

```
traceroute to cider.caida.org (192.172.226.123)
 1 lo2.akl-grafton-bba2.ihug.net (203.109.128.167) 46.041 ms
 2 gi1-1.akl-grafton-bdr2.ihug.net (203.109.130.110) 48.862 ms
 3 gi2-10.akl-grafton-bdr1.ihug.net (203.109.130.50) 178.426 ms
 4 Gi15-2.gw1.akl1asianetcom.net (203.192.166.41) 48.645 ms
 5 po2-0.gw1.lax1.asianetcom.net (202.147.61.189) 185.788 ms
 6 lax-cenic-equinix-exch.cenic.org (206.223.123.7) 185.832 ms
 7 calren3-cust.lsanca01.transitrail.net (137.164.131.242) 183 ms
 8 dc-lax-dc2--lax-dc1-ge--2.cenic.net (137.164.22.5) 184 ms
 9 dc-tus-dc1--lax-dc2-pos.cenic.net (137.164.22.43) 185 ms
10 dc-sdsc-sdsc2--tus-dc1-ge.cenic.net (137.164.24.174) 190 ms
11 pinot.sdsc.edu (198.17.46.56) 202 ms
12 cider.caida.org (192.172.226.123) 187 ms
```

Auckland (Grafton bridge), NZL
Los Angeles, CA
Equinix Exchange, LA, CA
Tustin, CA
San Diego, CA

DNS and Topology Discovery

- **undns** (part of Neil Spring's Scriptroute) contains a database of DNS to location

```
4648 \.global-gateway\.net\.nz {
    \.([a-z]{2})[bcs][rw][0-9]\.global-gateway\.net\.nz$  loc=1 {
        tk "Tokyo, Japan"
        ak "Auckland, NewZealand"
        sy "Sydney, Australia"
        sj "SanJose, CA"
        la "LosAngeles, CA"
    };
}
```

Goal

- Extend scamper to resolve IP to hostname mappings while probing, store in collected data files
- Use data to guide router alias resolution, guide location inference.

An aside: Comparing traceroute methods

- Multiple traceroute probing techniques exist
- UDP traceroute (traditional)
  - Variation: UDP-paris
- ICMP echo traceroute
  - Variation: ICMP-echo-paris
- TCP SYN traceroute
  - Variation: parasitic traceroute (paratrace)

Comparing traceroute methods

- Single source, 3 different destination sets
  - Alexa 500: top 500 websites ranked by Alexa
  - Router 500: 500 random routers on path to these websites
  - Random 1703: 1703 random IP addresses in unique routeviews IP prefixes

Dataset #1: 428 webservers

	completed	unreach	loop	gaplimit
udp	178 (41.6%)	26	15	209
udp-paris	180 (42.1%)	27	9	212
icmp	322 (75.2%)	9	16	81
icmp-paris	327 (76.4%)	10	10	81
tcp (p 80)	405 (94.6%)	0	11	12

15 targets (3.5%) observed the same sequence of IP hops

Dataset #2: 500 random routers

	completed	unreach	loop	gaplimit
udp	288 (57.6%)	72	1	139
udp-paris	286 (57.2%)	73	1	140
icmp	392 (78.4%)	66	3	39
icmp-paris	394 (78.8%)	67	1	38
tcp (p 80)	273 (54.8%)	75	1	151

33 targets (6.6%) observed the same sequence of IP hops

Dataset #3: 1703 random IP addresses

	completed	unreach	loop	gaplimit
udp	108 (6.3%)	180	181	1234
udp-paris	111 (6.5%)	178	139	1275
icmp	174 (10.2%)	204	172	1156
icmp-paris	174 (10.2%)	206	135	1188
tcp (p 80)	152 (8.9%)	188	137	1226

609 targets (40.6%) observed the same sequence of IP hops

Comparing traceroute methods

- Initial observations:
  - UDP traceroute gives relatively poor results
  - ICMP-echo traceroute tends to give best