

TDTS21: Advanced Networking


Lecture 7: Internet topology

Based on slides from P. Gill and D. Choffnes
Revised 2015, 2021 by N. Carlsson

Measuring the Internet's topology

2

- What do we mean by topology?
 - ▣ Internet as graph
 - ▣ Edges? Nodes?
 - ▣ Node = Autonomous System (AS); edge = connection.
- Edges labeled with business relationship
- Customer → Provider
- Peer -- Peer




The outputs

3

CAIDA's IPv4 AS Core AS-level Internet Graph

Archipelago January 2014



15412	12041	p2c
15412	12486	p2c
15412	12880	p2c
15412	13810	p2c
15412	15802	p2c
15412	17408	p2c
15412	17554	p2c
15412	17709	p2c
15412	18101	p2c
15412	19806	p2c
15412	19809	p2c
15413...		

So how do we measure this graph?


4

- ❑ Passive approach: BGP route monitors
 - ❑ Coverage of the topology
 - ❑ Amount of visibility provided by each neighbor
- ❑ Active approach: Traceroute
 - ❑ From where?
 - ❑ Traceroute gives series of IP addresses not ASes
- ❑ Active approach: TransitPortal
 - ❑ Much more control over what we see
 - ❑ ...scalability/coverage?

Passive approach: BGP Route Monitors

5

- Receive BGP announcements from participating ASes at multiple vantage points



www.routeviews.org

“originally motivated by interest on the part of operators in determining how the global routing system viewed their prefixes and/or AS space”

www.routeviews.org

Going from BGP Updates to a Topology

6

- Example update:
- TIME: 03/22/11 12:10:45
- FROM: 12.0.1.63 AS7018
- TO: 128.223.51.102 AS6447
- ASPATH: 7018 4134 9318 32934 32934 32934
- 69.171.224.0/20

AT&T (AS7018) is telling
Routeviews (AS 6447) about this route.

This /20 prefix can be reached via
the above path

Going from BGP Updates to a Topology

7

- Key idea
 - ▣ The business relationships determine the routing policies
 - ▣ The routing policies determine the paths that are chosen
 - ▣ So, look at the chosen paths and infer the policies
- Example: AS path “7018 4134 9318” implies
 - ▣ AS 4134 allows AS 7018 to reach AS 9318
 - ▣ China Telecom allows AT&T to reach Hanaro Telecom
 - ▣ Each “triple” tells something about transit service

Why are peering links hard to see?

□ The challenge:

- BGP announcements **do not reflect complete connectivity** information


Regional ISP won't see the peering edge!

- They are an agreement to transit traffic

Neither will Routeviews

Local ISP will only tell his customers about the peering link.

(ASes only transit traffic if it generates revenue!)



Combination of **no valley routing policy** and a **lack of monitors in stub ASes** mean missing **up to 90%** of peering links of content providers! (Oliveria et al. 2008)

Active approach: Traceroute

9

□ Issue: Need control over end hosts to run traceroute

□ H

□ http

□ C

□ H

□ RIPE

□ D

□ C

□ Das




Universities)



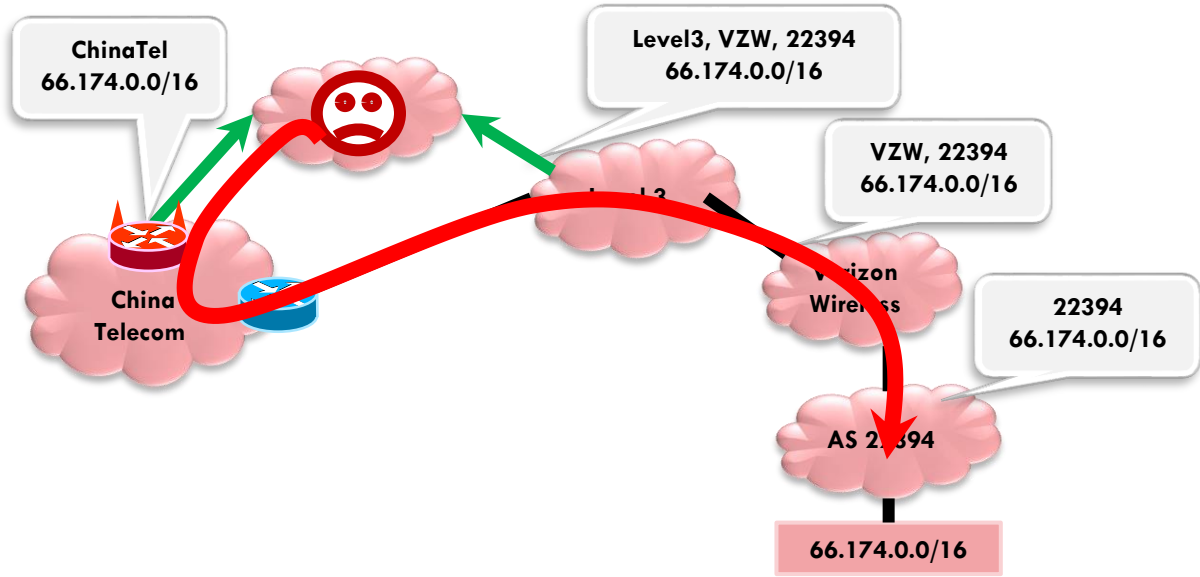
□ Bittorrent plug in that does measurements

□ O(200) ASes with Dasu clients

Traceroute vs Announced Path



Traceroute vs Announced Path

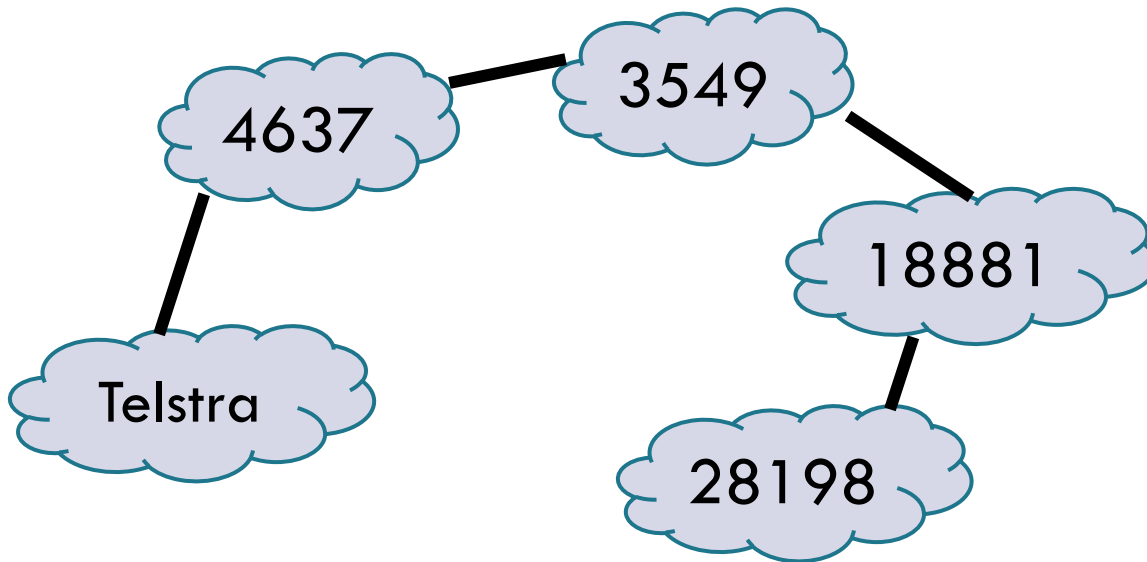


Interception typically results in differences between

- Announced AS-PATH
- Data path (traffic)

Policy checks if legit reason(s)

Traceroute vs Announced Path

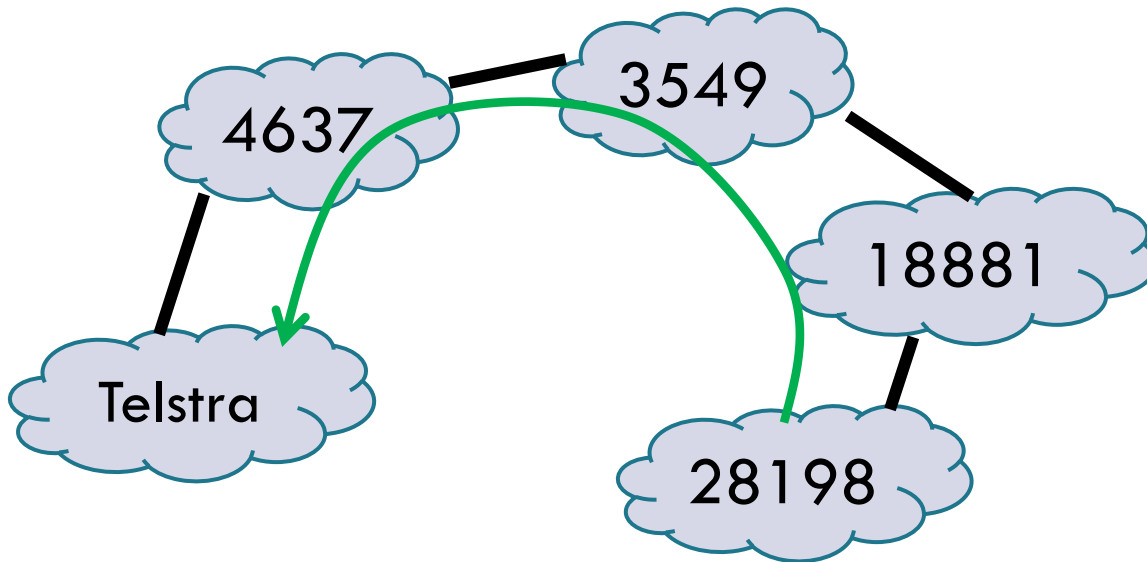


Sometimes differences

- Announced AS-PATH
- Data path (traffic)

Many legit reason(s)

Traceroute vs Announced Path




Sometimes differences

- Announced AS-PATH
- Data path (traffic)

Many legit reason(s)

AS-PATH: 177.52.48.0/21 | 1221 4637 3549 18881 28198

Traceroute vs Announced Path



Sometimes differences


— Announced AS-PATH

— Data path (traffic)

Many legit reason(s)

AS-PATH: 177.52.48.0/21 | 1221 4637 3549 18881 28198

Traceroute vs Announced Path



Sometimes differences

— Announced AS-PATH

— Data path (traffic)

Many legit reason(s)

AS-PATH: 177.52.48.0/21 | 1221 4637 3549 18881 28198

Traceroute:

... (initial hops)

9. telstraglobal.net (134.159.63.202) 164.905 ms


10. impsat.net.br (189.125.6.194) 337.434 ms

11. spo.gvt.net.br (187.115.214.217) 332.926 ms

12. spo.gvt.net.br (189.59.248.109) 373.021 ms

13. host.gvt.net.br (189.59.249.245) 343.685 ms

Traceroute vs Announced Path



Sometimes differences

— Announced AS-PATH

— Data path (traffic)


Many legit reason(s)

AS-PATH: 177.52.48.0/21 | 1221 4637 3549 18881 28198

AS HOPS in traceroute: 1221 1221 1221 1221 4637 4637 4637

4637 4637 3549 3549 3549 18881 18881 18881 18881 28198

Traceroute vs Announced Path



Sometimes differences

— Announced AS-PATH

— Data path (traffic)

Many legit reason(s)


AS-PATH: 177.52.48.0/21 | 1221 4637 3549 18881 28198

AS HOPS in traceroute: 1221 1221 1221 1221 4637 4637 4637

4637 4637 3549 3549 3549 18881 18881 18881 18881 28198

Traceroute-PATH: 1221 4637 3549 18881 28198

Traceroute vs Announced Path



Sometimes differences

— Announced AS-PATH

— Data path (traffic)

Many legit reason(s)

AS-PATH: 177.52.48.0/21 | 1221 4637 3549 18881 28198

AS HOPS in traceroute: 1221 1221 1221 1221 4637 4637 4637


4637 4637 3549 3549 3549 18881 18881 18881 18881 28198

Traceroute-PATH: 1221 4637 3549 18881 28198

Distribution properties of AS graph

Zipf popularity...



... and long tails



$$v_r \propto r^{-\alpha}$$

Zipf popularity...


... and long tails




$$v_r \propto r^{-\alpha}$$

Zipf popularity...

... and long tails




$$v_r \propto r^{-\alpha}$$




$$\log v_r = \log v_1 - \alpha \log r$$

Zipf popularity...

... and long tails




$$v_r \propto r^{-\alpha}$$




$$\log v_r = \log v_1 - \alpha \log r$$

Zipf popularity...

... and long tails



$$v_r \propto r^{-\alpha}$$




$$\log v_r = \log v_1 - \alpha \log r$$


$y(x) = x_0 - \alpha x$

Zipf popularity...

... and long tails



$$v_r \propto r^{-\alpha}$$




$$\log v_r = \log v_1 - \alpha \log r$$


$y(x) = x_0 - \alpha x$

Zipf popularity...

... and long tails



$$v_r \propto r^{-\alpha}$$



$$\log v_r = \log v_1 - \alpha \log r$$

Power law, Pareto, and Zipf

- Power-law, Pareto, Zipf (in some sense the same)
 - ▣ Power-law: $f(x) \sim x^{-\eta}$ (probability of value x)
 - ▣ Pareto: $F(x) = P[X > x] = \int f(x) dx \propto x^{-\kappa}$ (cumulative prob.)
 - ▣ Zipf: $v_r \propto r^{-\alpha}$ (discrete representation; frequency v_r of rank r)
 - ▣ Parameters related as: $\kappa = \eta - 1 = 1/\alpha$
 - E.g., paper and references therein: “A Tale of the Tails: Power-laws in Internet Measurements”, IEEE Network, Mahanti et al., 2013

Heavy-tail distributions ...

- “A probability distribution is said to have a heavy tail if the tail is not exponentially bounded”
- ... and then there are many many other “heavy tail” distributions, variations and generalizations, including distributions such as log-normal, various generalized Zipf/Pareto distributions, etc.

