Topology Inference from BGP Routing Dynamics

David Andersen, Nick Feamster, Steve Bauer, Hari Balakrishnan

MIT Laboratory for Computer Science

October 2002

`http://nms.lcs.mit.edu/ron/`

Current Topologies: AS Topologies

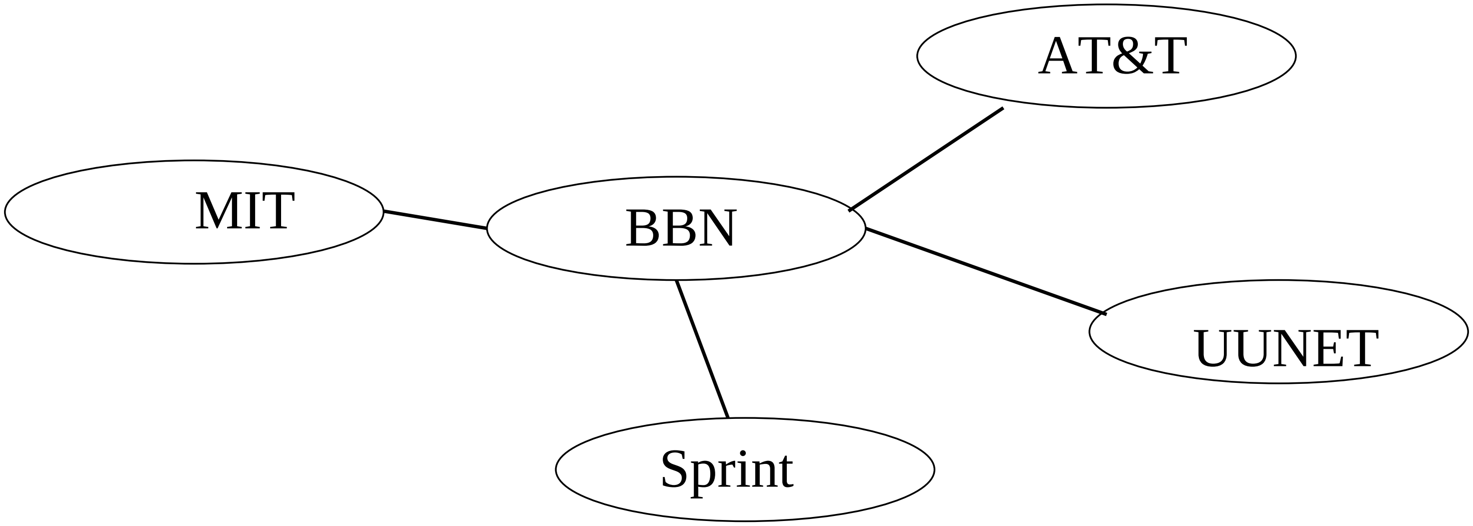

- ✓ Simple to construct
- ✓ Completely passive - BGP snapshot
- ✗ Obnoxiously free of interesting detail

A few paths contain most prefixes

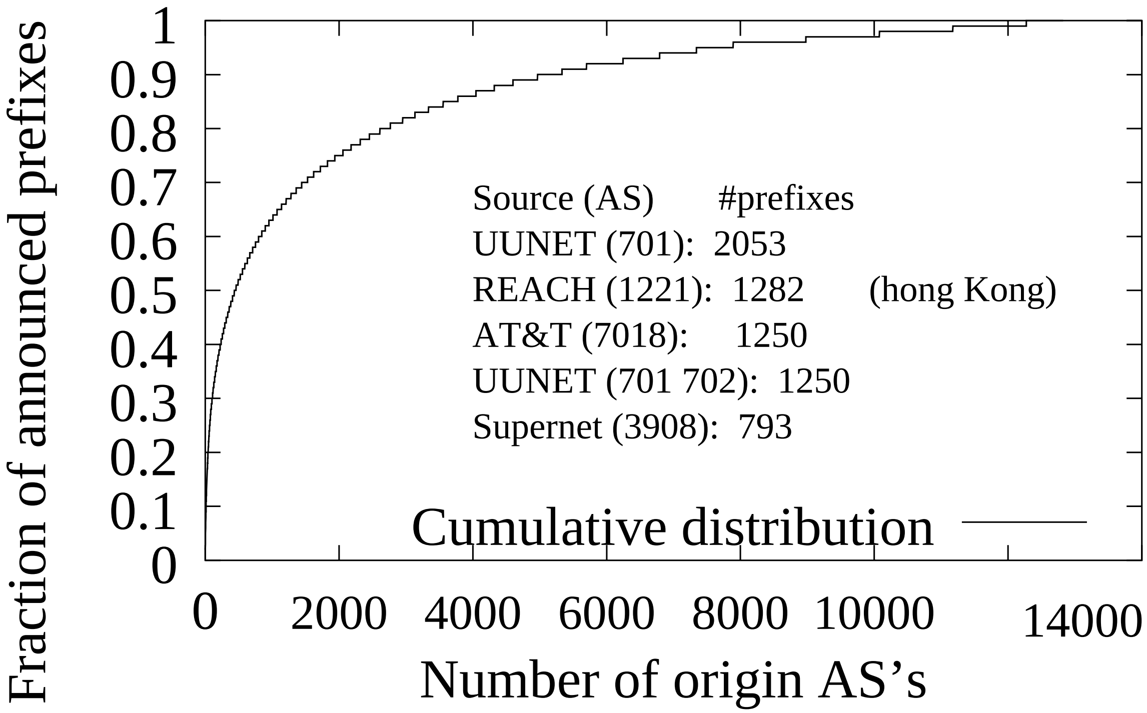

- 13 common paths contain 10% of prefixes
- Binning large ISPs misses critical detail

Current Topologies: Router-Level

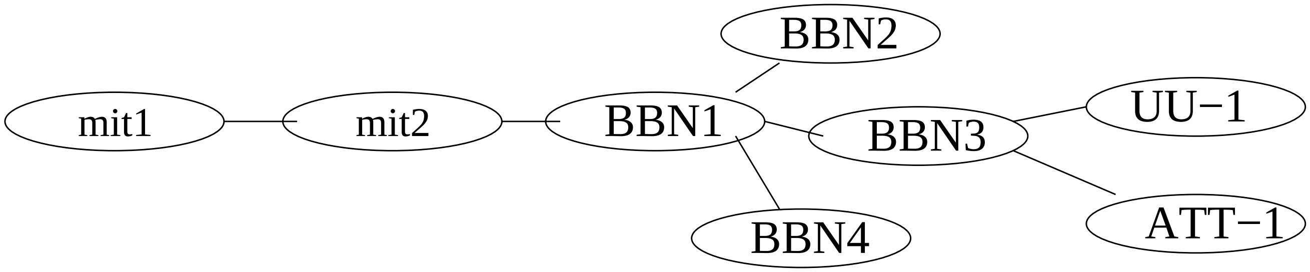

✓ Lots of juicy detail

✗ Requires active probing

- Annoys the paranoid (and can be blocked)

- Consumes time and bandwidth

➜Best of both worlds?

New: Implied *Logical* Topologies

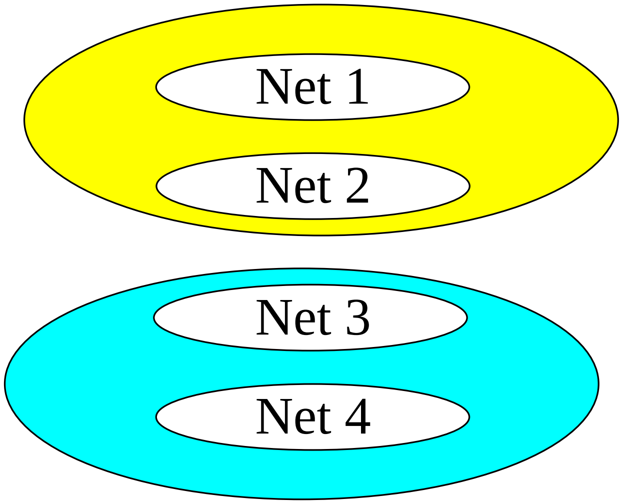

- Group prefixes that “behave similarly”
- What do the resulting clusters mean?

BGP update streams

2002-01-10 23:51:05	198.140.178.0/24
2002-01-10 23:51:05	192.107.237.0/24
2002-01-10 23:55:53	199.230.128.0/23
2002-01-10 23:56:21	216.9.174.0/23
2002-01-10 23:56:21	216.9.172.0/24

- Colored prefixes updated at (nearly) same time

→ Cluster prefixes that often do this

Mechanics

```
2002-01-10 23:51:05    198.140.178.0/24
2002-01-10 23:51:05    192.107.237.0/24


2002-01-10 23:55:53    199.230.128.0/23


2002-01-10 23:56:21    216.9.174.0/23
2002-01-10 23:56:21    216.9.172.0/24
```

- *Group* by 30-second intervals
  (in practice, bin length choice flexible) (BGP min-route-adver time)

Creating BGP update vectors

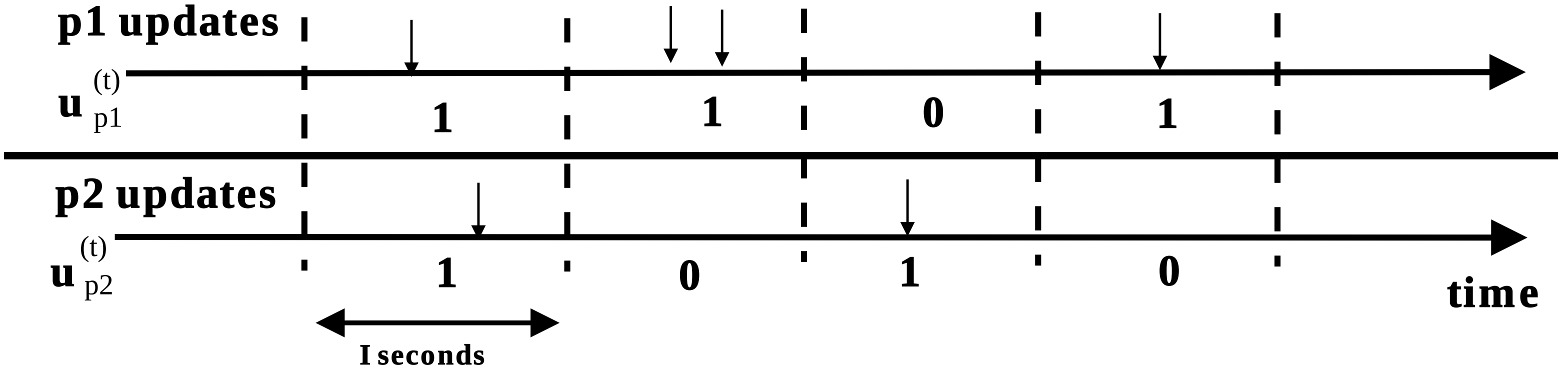

- Update stream is a 0/1 signal

  Did an update happen in time $[t, t + 30s]$*?*

- Now we have a bunch of 0/1 vectors to compare...

BGP update vectors

$time \longrightarrow$

Prefix A	0	0	1	0	1	0	0
Prefix B	1	0	1	0	0	0	1
Prefix C	1	0	1	0	0	0	0

How close are two vectors?

- Correlation coefficient

Correlation Coefficient

A	0	0	1	0	1	0	0
B	1	0	1	0	0	0	1
C	1	0	1	0	0	0	0

$$\mathrm{corr}(p_1, p_2) \quad = \quad \frac{E[(p_1 - \overline{p_1})(p_2 - \overline{p_2})]}{\sigma_{p1}\sigma_{p2}}$$

- Expresses correlation well
- Susceptable to some “coincidental” correlation

How to Group Prefixes?

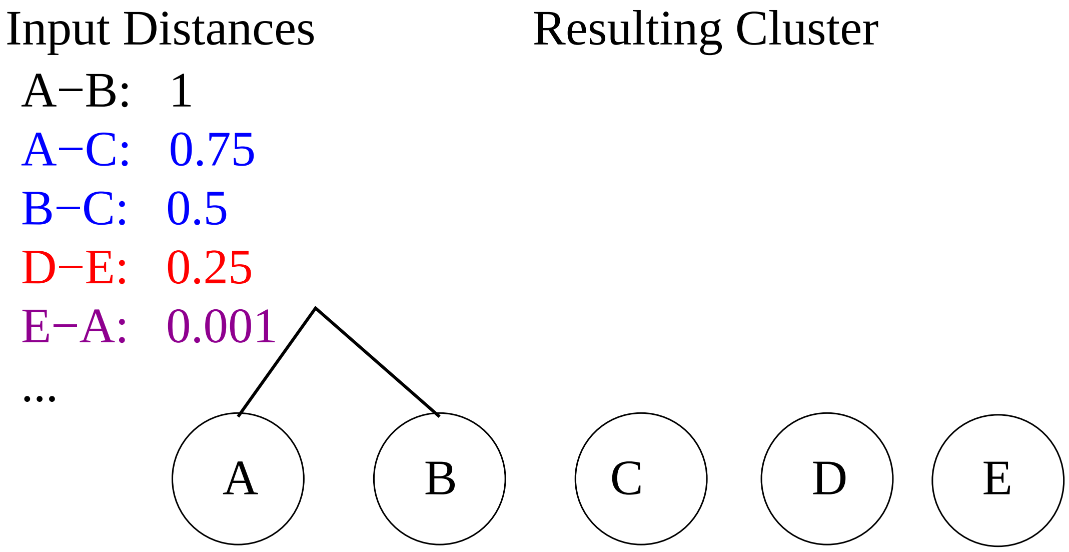
Input Distances

A−B: 1
A−C: 0.75
B−C: 0.5
D−E: 0.25
E−A: 0.001
...

Resulting Cluster

Single-linkage clustering

- Simple and efficient
- Creates a similarty hierarchy: A & B most similar, etc.

How to Group Prefixes?

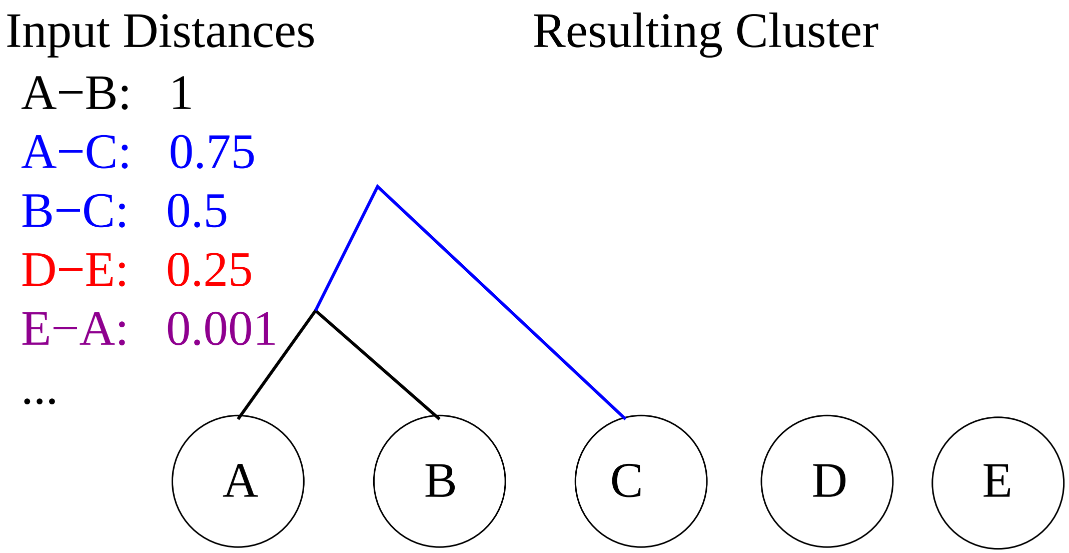

Input Distances

A−B: 1

A−C: 0.75

B−C: 0.5

D−E: 0.25

E−A: 0.001

...

Resulting Cluster

Single-linkage clustering

- Simple and efficient
- Creates a similarty hierarchy: A & B most similar, etc.

How to Group Prefixes?

Input Distances

A−B:	1
A−C:	0.75
B−C:	0.5
D−E:	0.25
E−A:	0.001
...	

Resulting Cluster

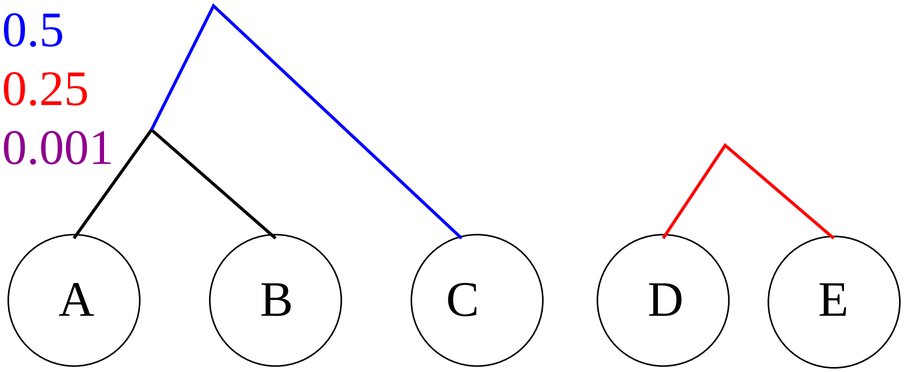

Single-linkage clustering

- Simple and efficient
- Creates a similarty hierarchy: A & B most similar, etc.

How to Group Prefixes?

Input Distances

A−B: 1
A−C: 0.75
B−C: 0.5
D−E: 0.25
E−A: 0.001
...

Resulting Cluster

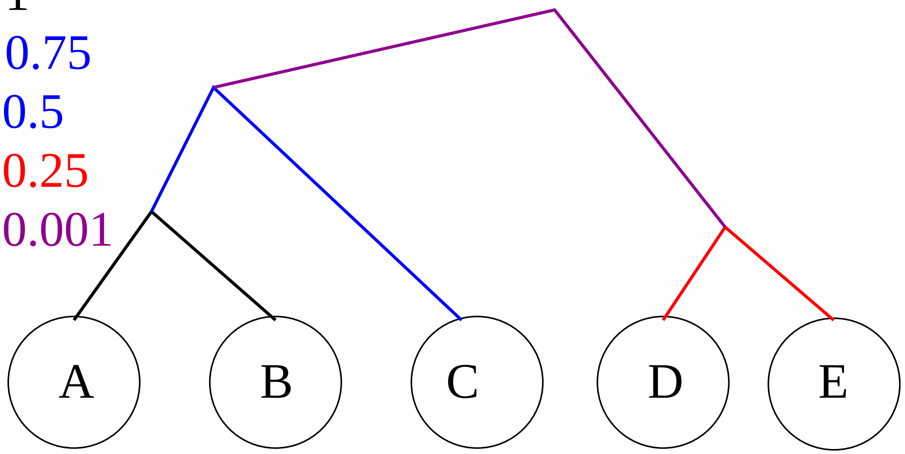

Single-linkage clustering

- Simple and efficient
- Creates a similarty hierarchy: A & B most similar, etc.

Data Capture and Analysis

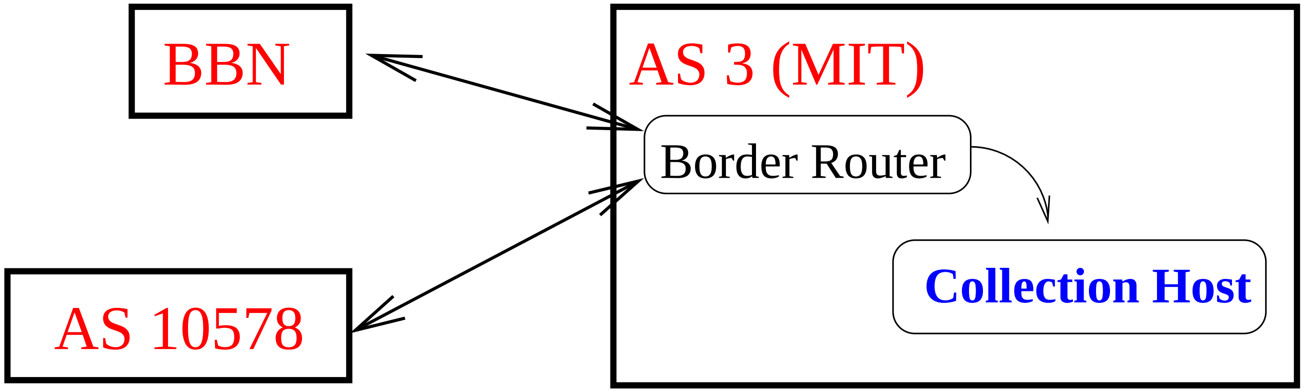

- Studied 90 days of BGP traffic at MIT
- Examined 2 “huge” origin ASes
  - UUNET: 2338 prefixes
  - AT&T: 1310 prefixes
- How do clusters relate to real-word features?

Anecdotes

- Many “expected” results - same city, etc. We’ll get to those in a second.
- 135.36.0.0/16, 135.12.0.0/14. Denver vs. New Jersey. Lucent vs. Agere – a spinoff in 2000, identical network behavior. (... CIA?)
- 6 Sandia labs prefixes - internet2 routes, but flapped to backup UUNET route.
- Many transient discoveries: backups, etc.

Topological similarities

Measureable quantities: path, location

- Compute pairwise similarity for metric (shared path length, or shared pop)
- *Average* similarity as clustering proceeds
- If match with logical clustering,
  similarity strongest for leaf clustering,
  weakest at end.

➔Logical topology: integration of topological, organizational, and administrative factors.

Leaves share more hops in traceroute

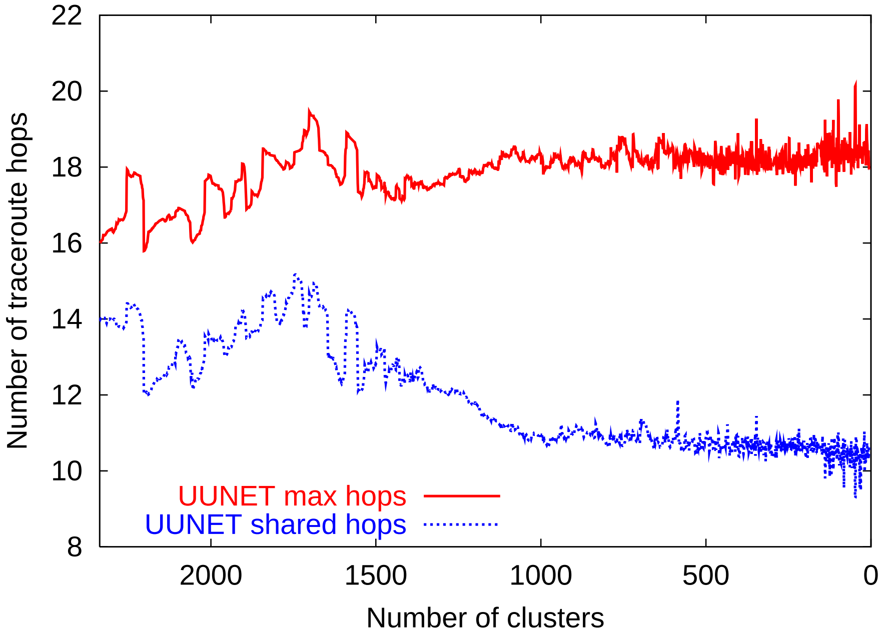

- Path length varies less with clustering
- More shared hops in earlier clustering
- Data noisy: loops, etc., but still works

Leaves often share the ISP POP

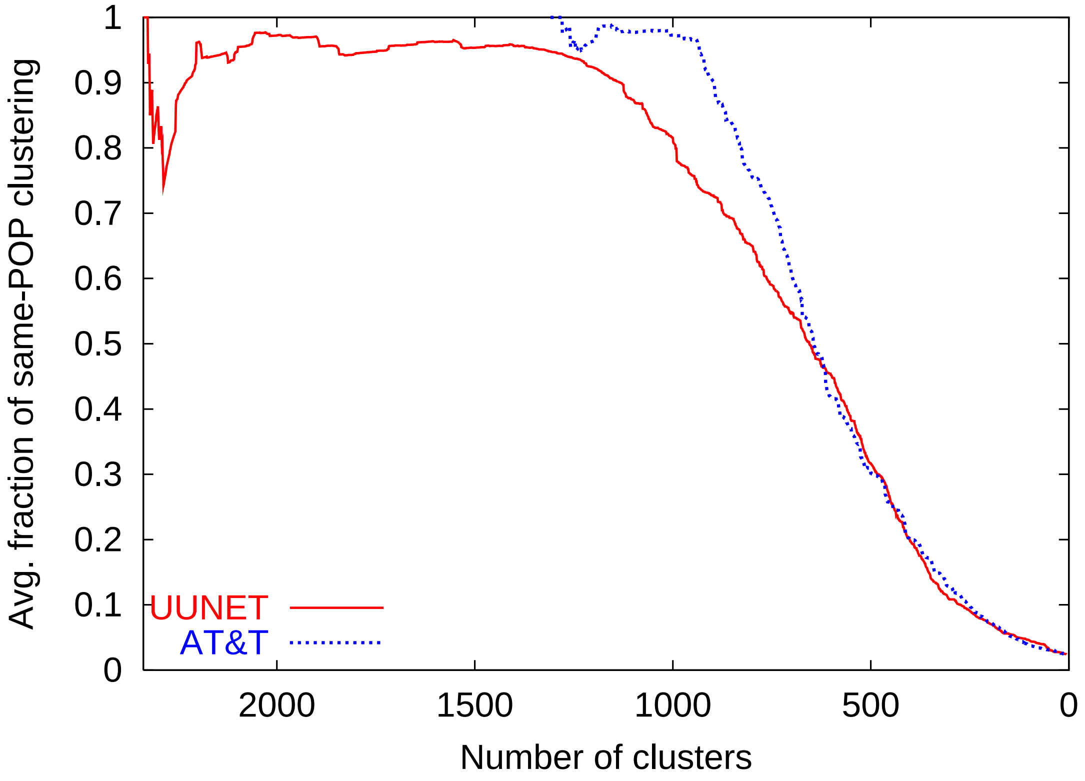

- UUNET: 50% clustered at 95% accuracy
- AT&T: 30% clustered at 97% accuracy

What does it all mean?

- Update clusters reflect reality:
  - Topology
  - Prefix assignment
  - Fate sharing
- Passive window into remote networks
- Facilitate network mapping and data collection
- What else can be extracted from this signal? Similar signals?