Topology Inference from BGP Routing Dynamics

David Andersen, Nick Feamster, Steve Bauer, Hari Balakrishnan
M.I.T. Laboratory for Computer Science
{dga,feamster,bauer,hari}@lcs.mit.edu

Internet Topology Estimation

- Current techniques:
  - Passive: BGP routing table-based AS-level topologies
  - Active: Traceroute-based intra-AS topologies
  - Combinations of the above (e.g., Rocketfuel)
- Our approach:
  - Passively observe prefix relationships inside an AS.
  - Treat BGP updates as a signal that contains logical topology information

Table-Based Autonomous System Topologies

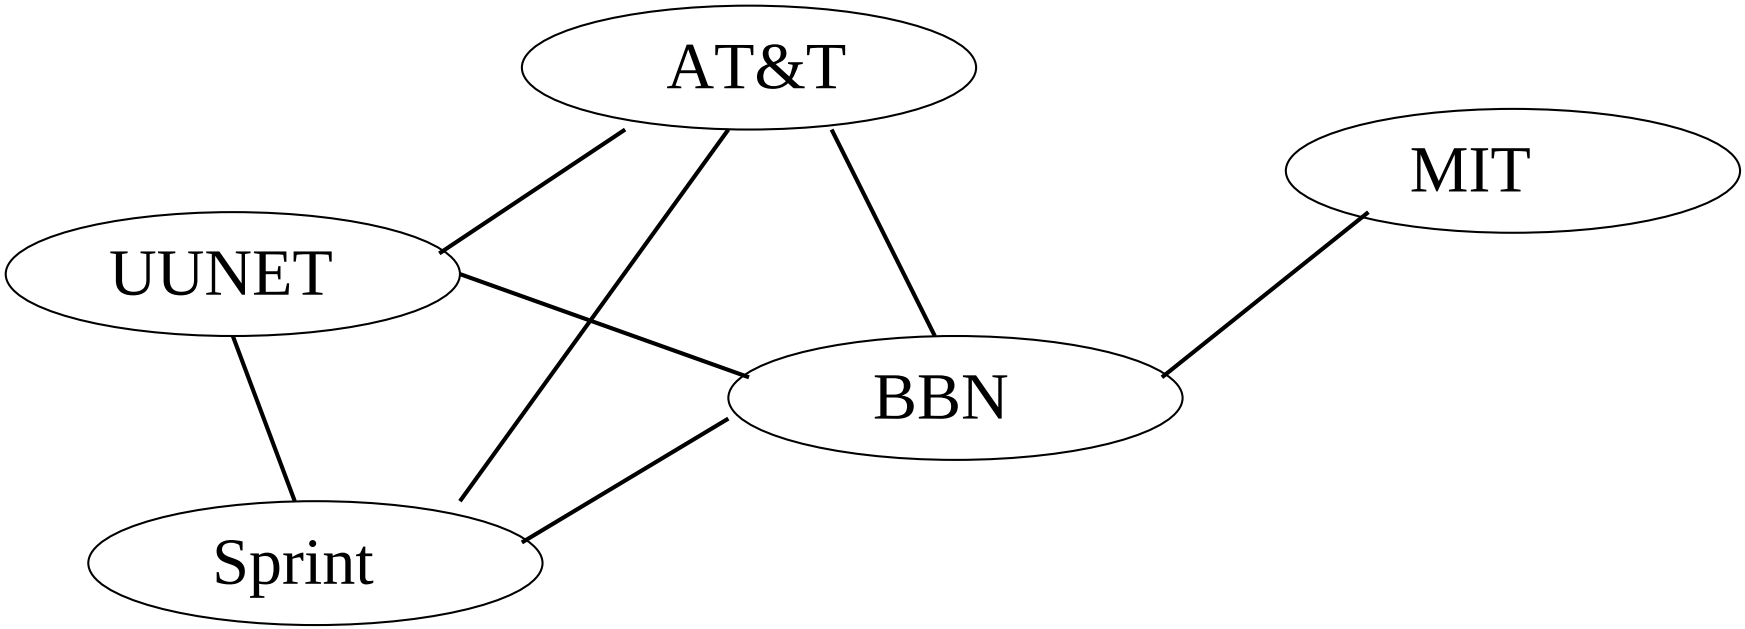

- Simple and passive
- Lacks detailed information...why?

Most prefixes are advertised by common paths

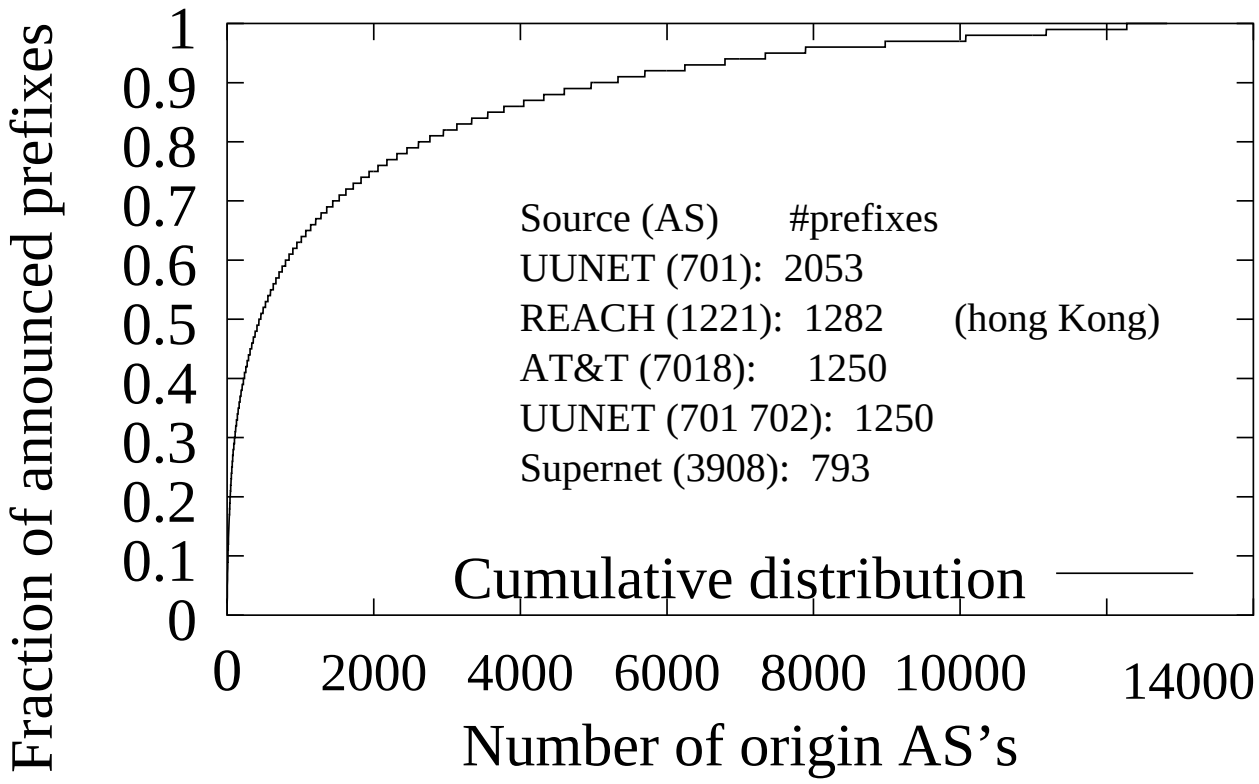

- 13 common paths contain 10% of prefixes
- Not all prefixes within an AS are like.

Router-level Topologies

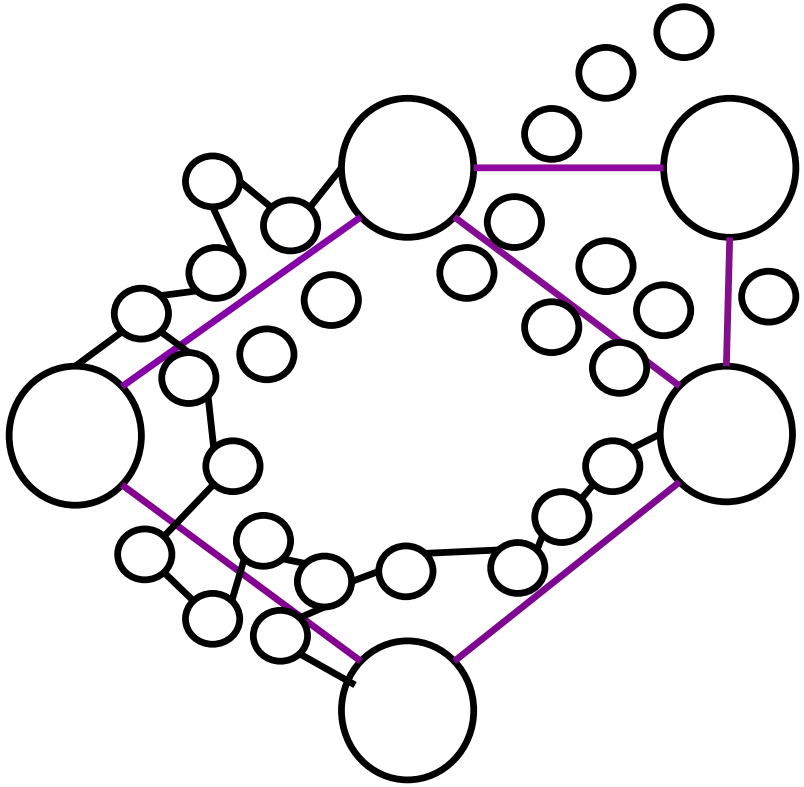

- Contains considerable sub-AS detail
- Requires active probing
  - Traceroute can be blocked, and can generate complaints
  - Potentially bandwidth intensive

Logical Topologies from Routing Dynamics

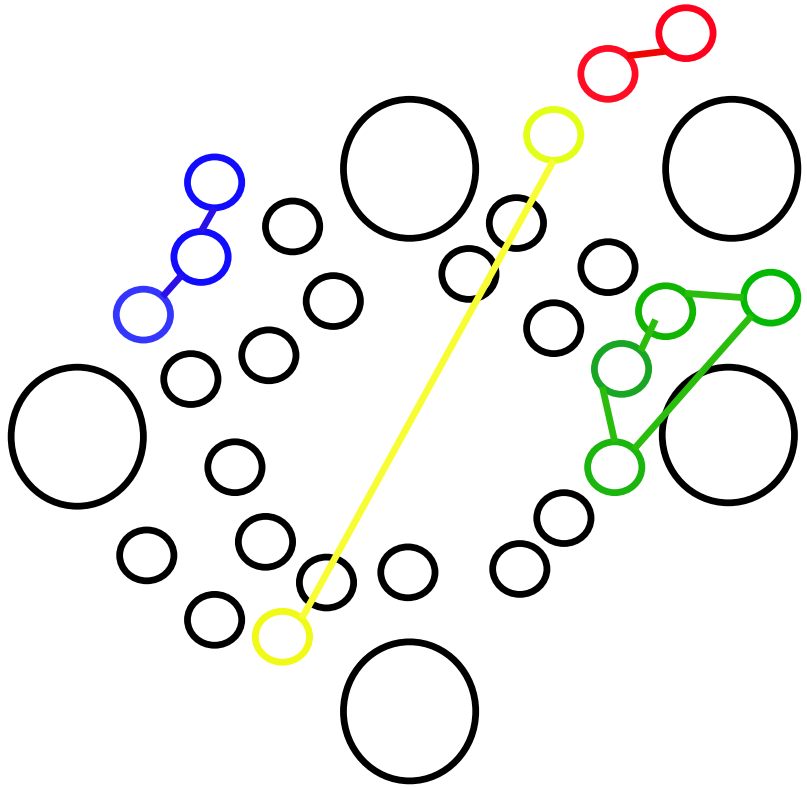

- Group prefixes that have similar update patterns.
- Do these groups share common characteristics?

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 experienced a routing change at nearly the same time.
- Might they have something in common?
- Processing:
  - Filter out noise (e.g., session resets).
  - Divide timeseries into discrete 30-seconds windows.

Timeseries Details

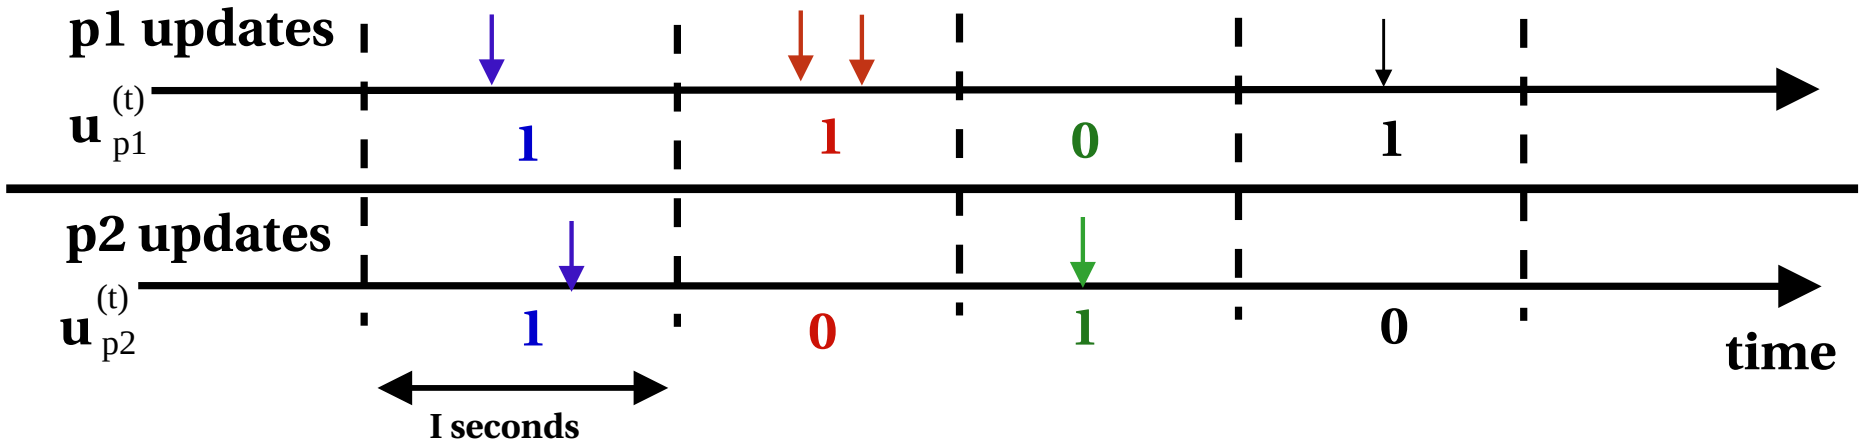

- Update stream is 0/1 signal.
  - If prefix has at least one update w/i window: 1
  - Otherwise: 0
- Now, can use techniques to compare these two prefixes

How to Measure "Closeness"?

	t_1	t_2	t_3	t_4	t_5	t_6	t_7
A	0	0	1	0	1	0	0
B	1	0	1	0	0	0	1
C	1	0	1	0	0	0	0

$$\sigma_{xy} = \frac{E[(x(t) - \mu_x)(y(t) - \mu_y)]}{\sigma_x \sigma_y}$$

- Correlation Coefficient
  - Expresses correlation well
  - One coincidence may falsely indicate perfect correlation (if update traffic is low)

How to Group Prefixes?

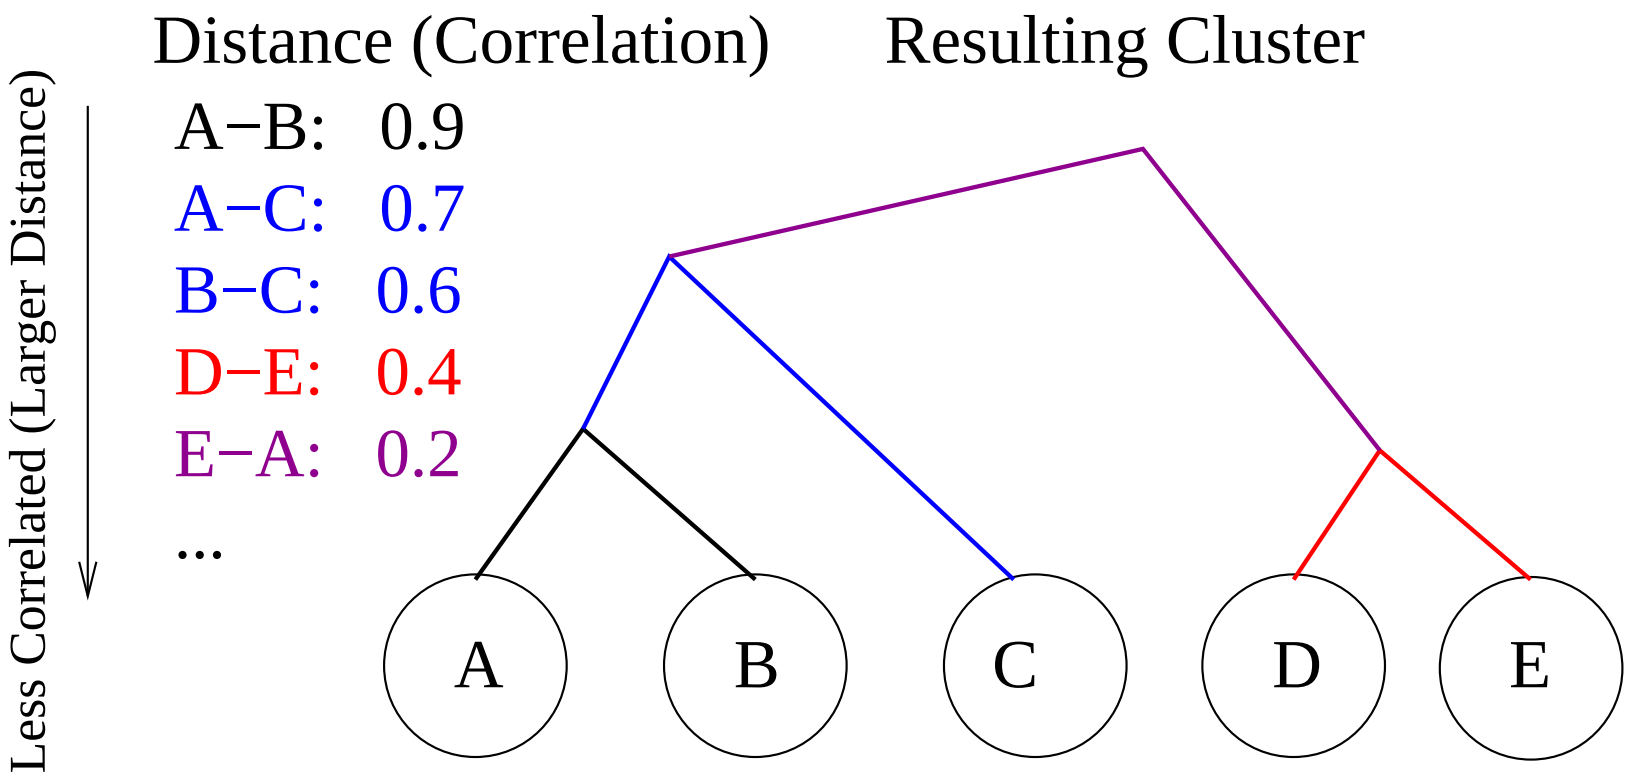

- Single-linkage clustering
- Simple and efficient

Data Collection

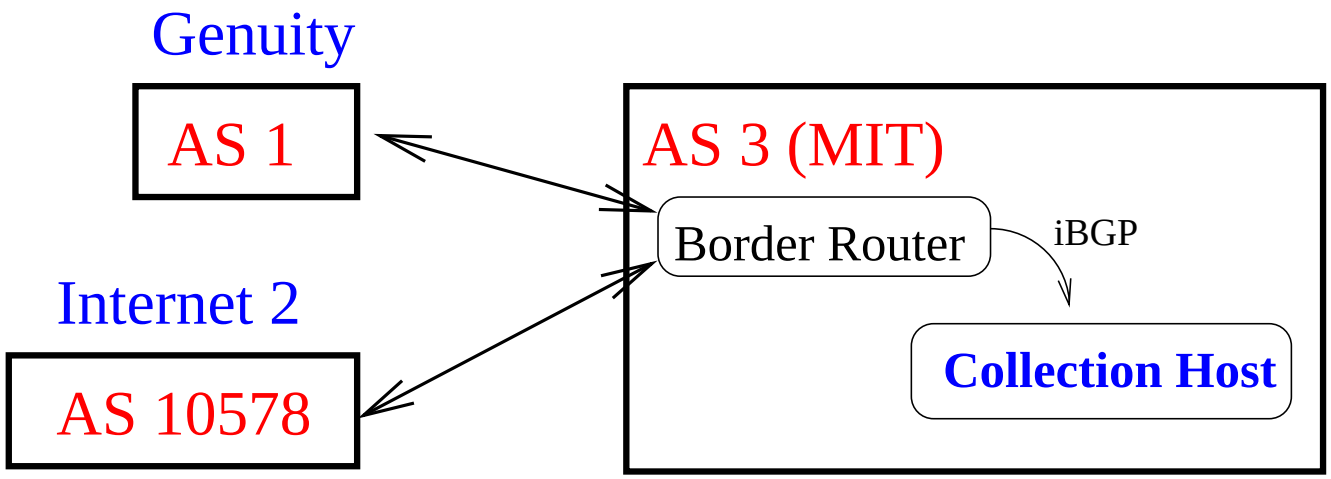

- iBGP Session: Changes to best route only.
- 3 months of BGP updates originating from
  - UUNet (2338 prefixes)
  - AT&T (1310 prefixes)
- Didn’t consider when prefixes change origin

Fun Anecdotes

- 135.36.0.0/16, 135.12.0.0/14. Denver vs. New Jersey. Lucent vs. Agere -- a spinoff in 2000, identical network behavior.
- 6 Sandia labs prefixes - internet2 routes, but flapped to backup UUNET route.
- Many transient discoveries: backups, etc.

Prefix Groups Share Traceroute Hops

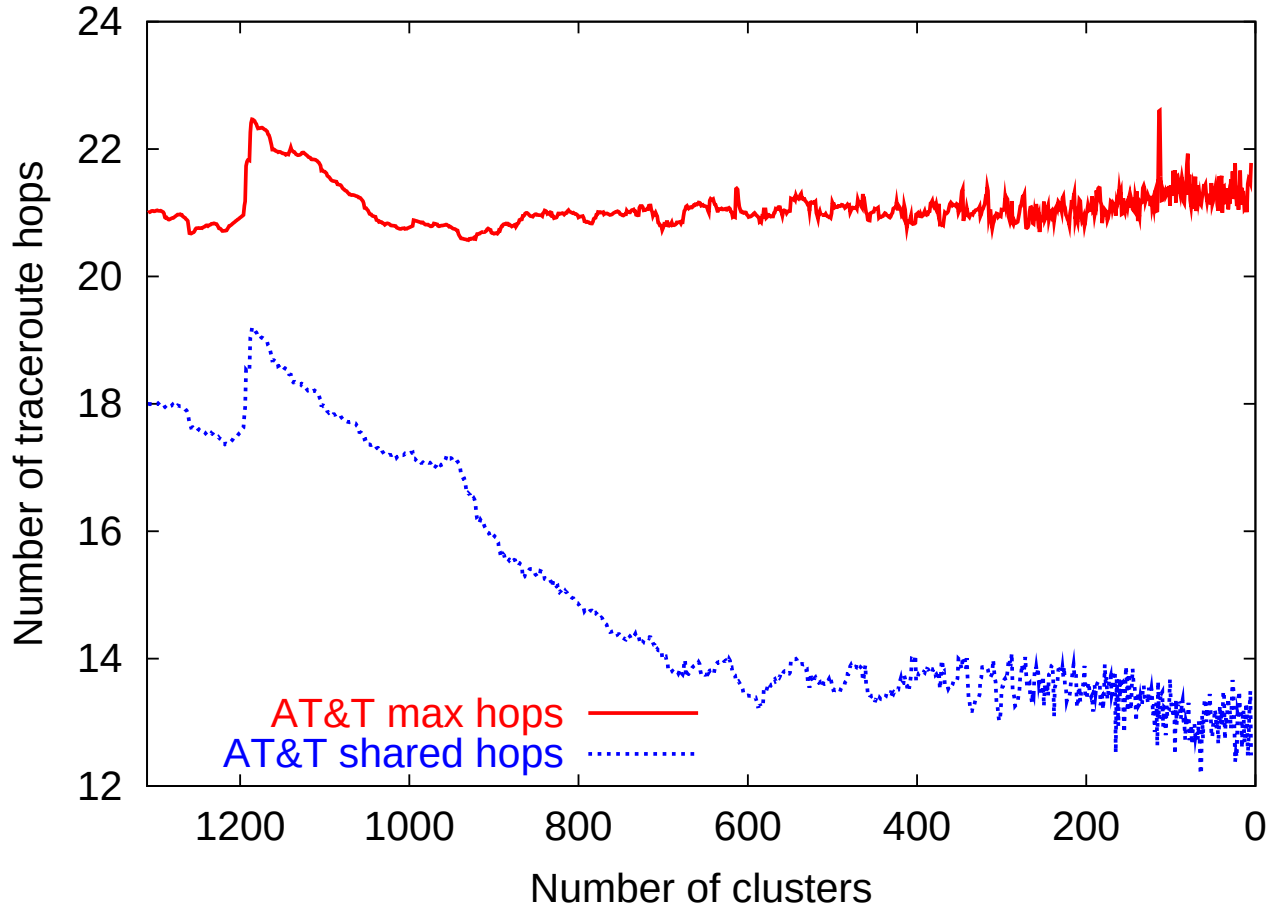

- Tough to figure out the meaning here
  - Many shared hops on customer routes
  - Standard traceroute "noise"

Prefix Clusters Often Share the Same PoP

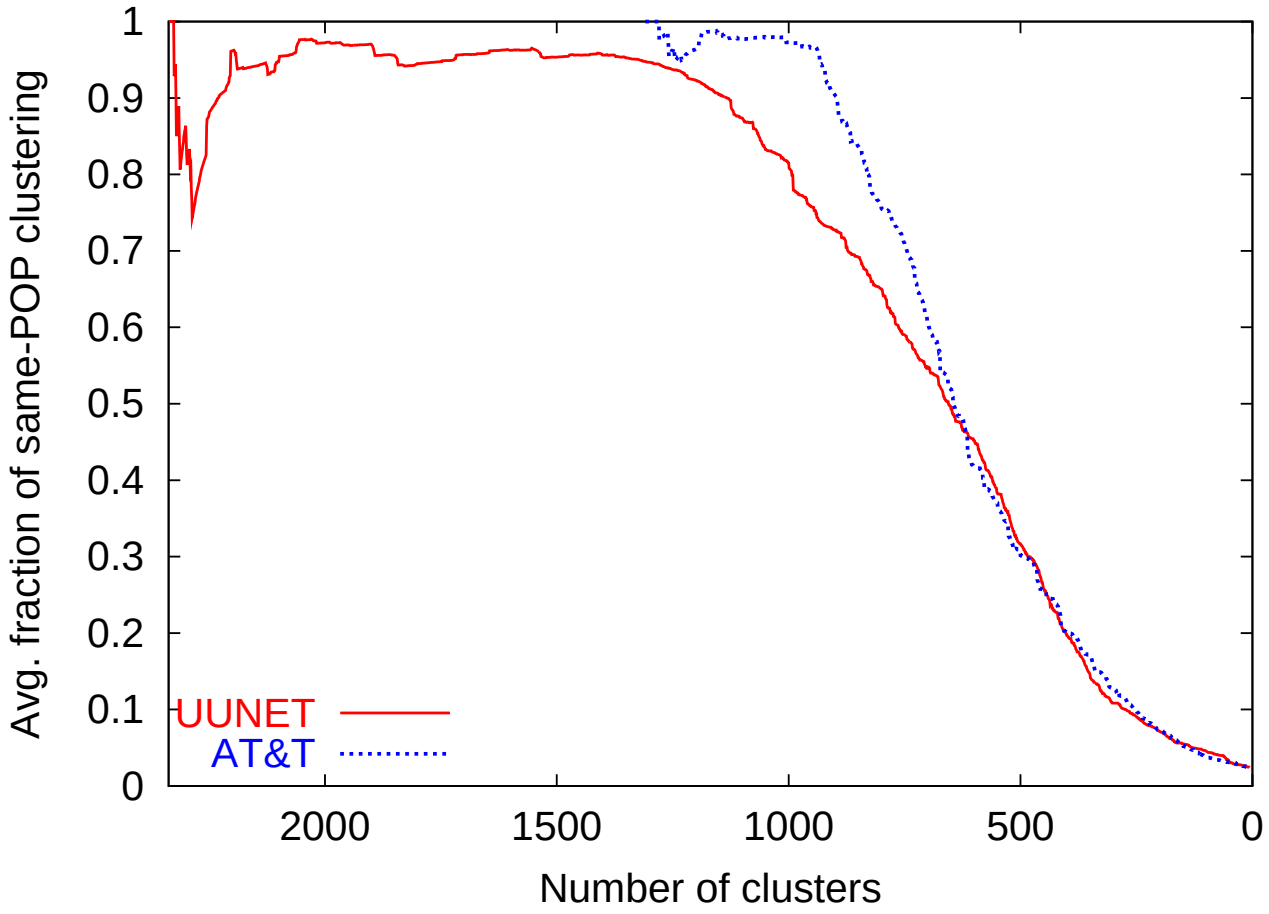

- UUNet: 50% Clustered, 97% Accuracy
- AT&T: 30% Clustered, 95% Accuracy

Prefix Clusters are Close in Address Space

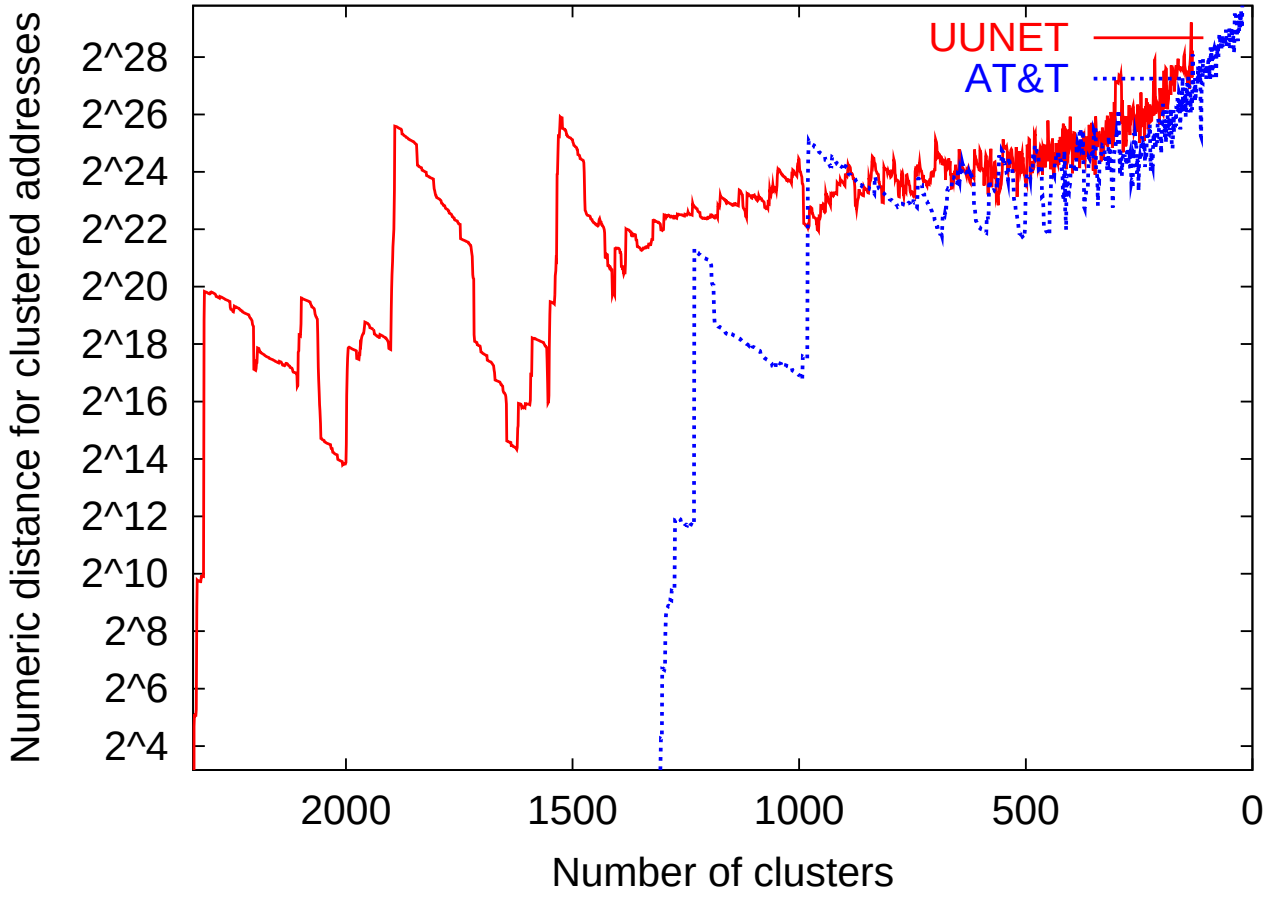

- Clusters that form first are closer in address space.

Other Thoughts

- What about other correlations?
  - Changes in the AS path to guess about failure points?
- Other signal weirdness (cf. BGP misconfiguration). Can we see:
  - Load balancing?
  - Route hijacking?
  - Failover/backup relationships?
  - Policy slips due to NTAF?

Conclusion

- More to passive topology than the snapshot!
- Routing dynamics reveal interesting details about logica topology
  - Fate sharing
  - Prefix assignment
  - History :)
- Uses only passive measurements.
- Can reduce the amount of data needed for topology mapping, etc.