Issues with Inferring Internet Topological Attributes

Lisa Amini[ab], Anees Shaikh[a], Henning Schulzrinne[b]

[a]*IBM T.J. Watson Research Center*

[b]*Columbia University*

Internet Mapping

- Goal
  - Create mathematical and graph theoretic models of the Internet
    - detect pathologies
    - improve existing protocols
    - validate proposals for new protocols
    - predict the future evolution of the Internet
- Issues
  - Rapid growth
  - No single entity has complete representation
  - Conflicting empirical data (routing tables, traceroute)
- Can the impact of ambiguities in empirical data be quantified?

BGP Routing Table

- IP network prefix advertisements include AS_PATH

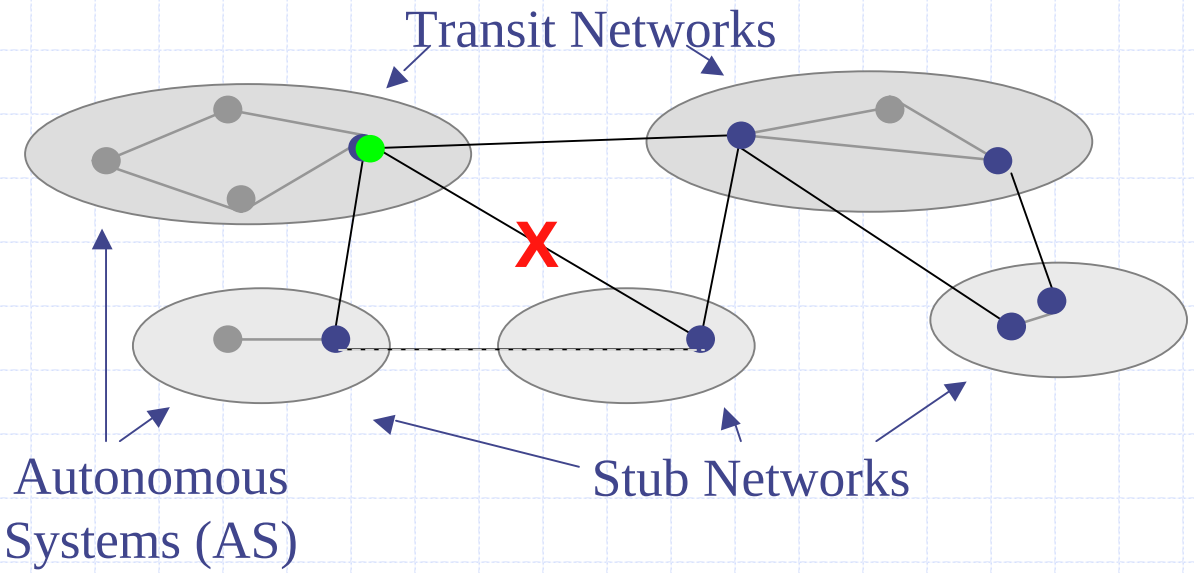

- Partial Information
  - Single Viewpoint, Route Selection, Route Filtering
- Intra- vs Inter- domain
  - Static Routes, Source Routing, Multi-hop BGP Sessions, Route Stuffing

traceroute

- Set TTL to elicit ICMP response from intermediate routers
  - Associate router IP address with AS number

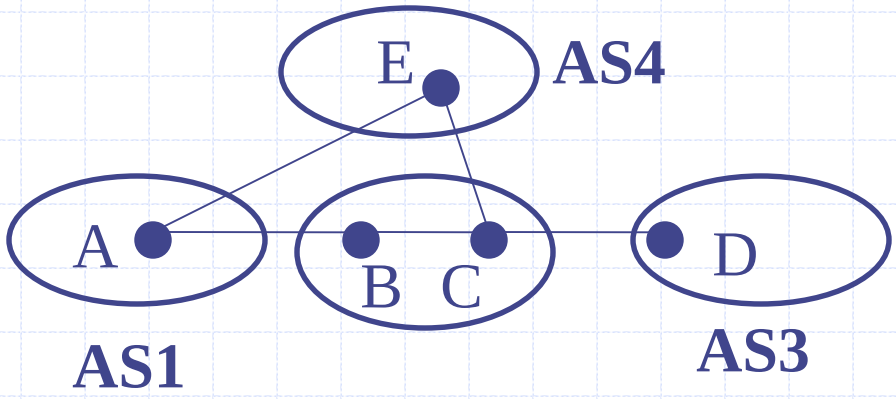

- AS Number Resolution
  - Registry Data, Multiple AS Numbers
- ICMP Message Generation
  - RFC1812 – use transmitting interface address for source in response

Methodology

- Two-way measurements
  - Randomly select a accessible source and destination
  - Request BGP AS_PATH to source/destination from local router at destination/source
  - Request traceroute to source/destination from in each direction
- One-way measurements
  - Randomly select accessible source
  - Randomly select IP address prefix from BGP table convert to destination IP address in subnet
  - Request BGP AS_PATH to destination from local router at source
  - Request traceroute to destination from source
- Events centrally scheduled at Poisson intervals
  - Issue: failure tracking

Data Sources

- Looking Glass
  - HTTP interface
    - Traceroute to target address
    - BGP route entry (AS_PATH)
  - Geographically and topologically diverse
    - 92 with BGP and traceroute enabled
    - 25 countries, ~28% in US
  - www.traceroute.org
- Oregon Route Views
  - BGP routing tables
  - Peers with 57 other routers
  - Most major ISP's represented
  - www.antc.uoregon.edu/route-views

Discard Criteria

- Looking Glass server not responding
- Incomplete traceroute output
- Node address in 10.x.x.x, 172.16.x.x-172.32.x.x, or 192.168.x.x range.
- Route did not terminate in target AS
- Intermediate node did not respond to ICMP echo
- No matching BGP/traceroute data for same period
- For two-way measurements:
  - No matching reverse probe for same time period

Dataset Summary

- D1: Randomly pair LG's to collect forward and reverse paths of each type
  - 116302 attempted measurements
  - 8464 unique routes (2840 fully paired)
- D2: Oregon Route Views tables for 18 days corresponding with D1.
- D3: Randomly pair LG with advertised network address at Poisson intervals.
  - 62645 attempted measurements
  - 27185 unique routes
- D4: Oregon Route Views tables for 11 days corresponding with D3.

Cursory Dataset Comparison

Dataset	D1	D2	D3	D4
Date Collected	3/02	3/02	4/02	4/02
Collection Duration (days)	18	18	11	11
Data Source	LG	ORE	LG	ORE
Number of nodes/AS's	337	13054	7640	13226
Number of edges	1937	53816	25812	55410

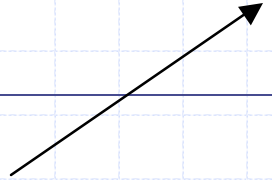

100% of top 20 AS's in D2, ranked by degree

AS Path Asymmetry

- Path(A,B) ⊙Path(B,A)

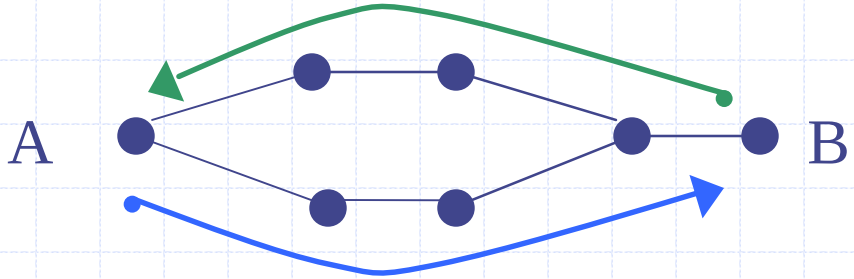

- 1995 (Paxson) study: AS path asymmetry at ~30%
  - Total path asymmetry at ~50%
- Current (1Q2002) level according to traceroute data: 69.8%
  - Artificially inflated?
- BGP-assisted resolution: 61.4%
- BGP AS_PATH asymmetry: 56.3%.

Distribution of Hop Differences

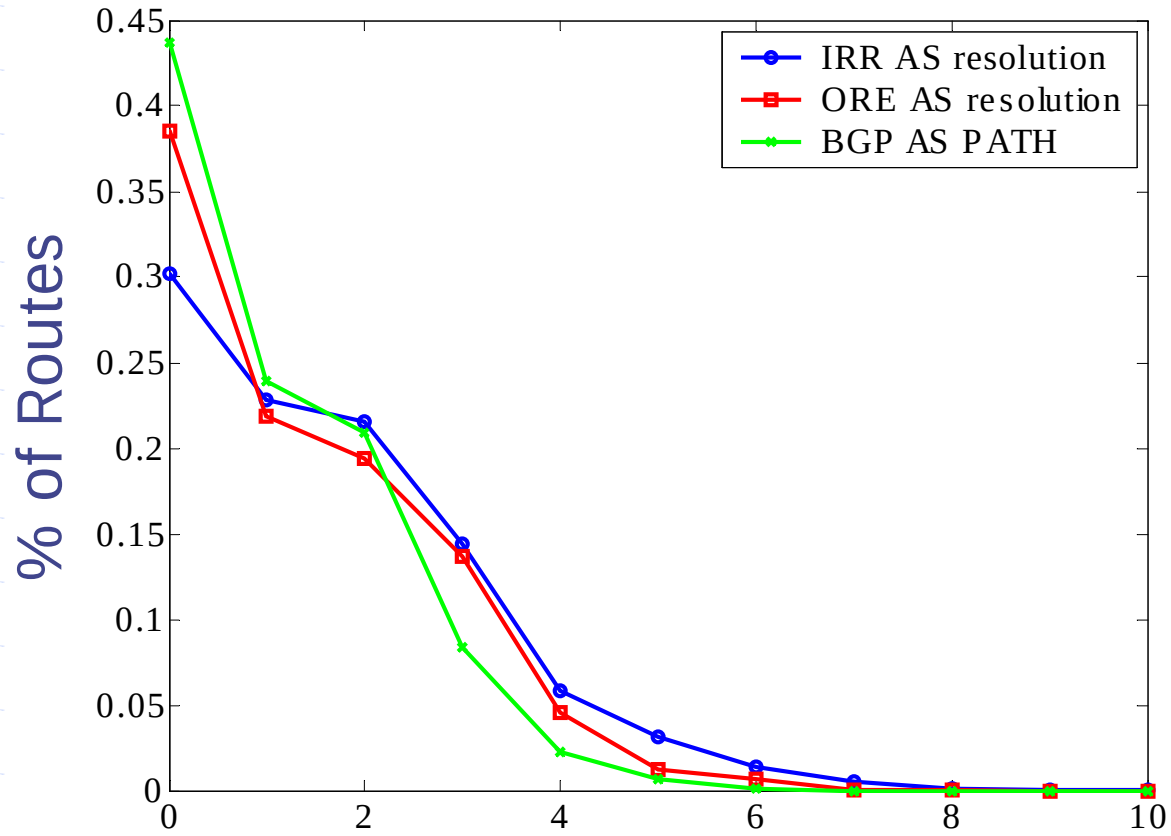

Nearly 15% difference between traceroute and AS_Path results

Of the routes that were asymmetric, nearly 60-80% differed by only one or two hops.

Path Length

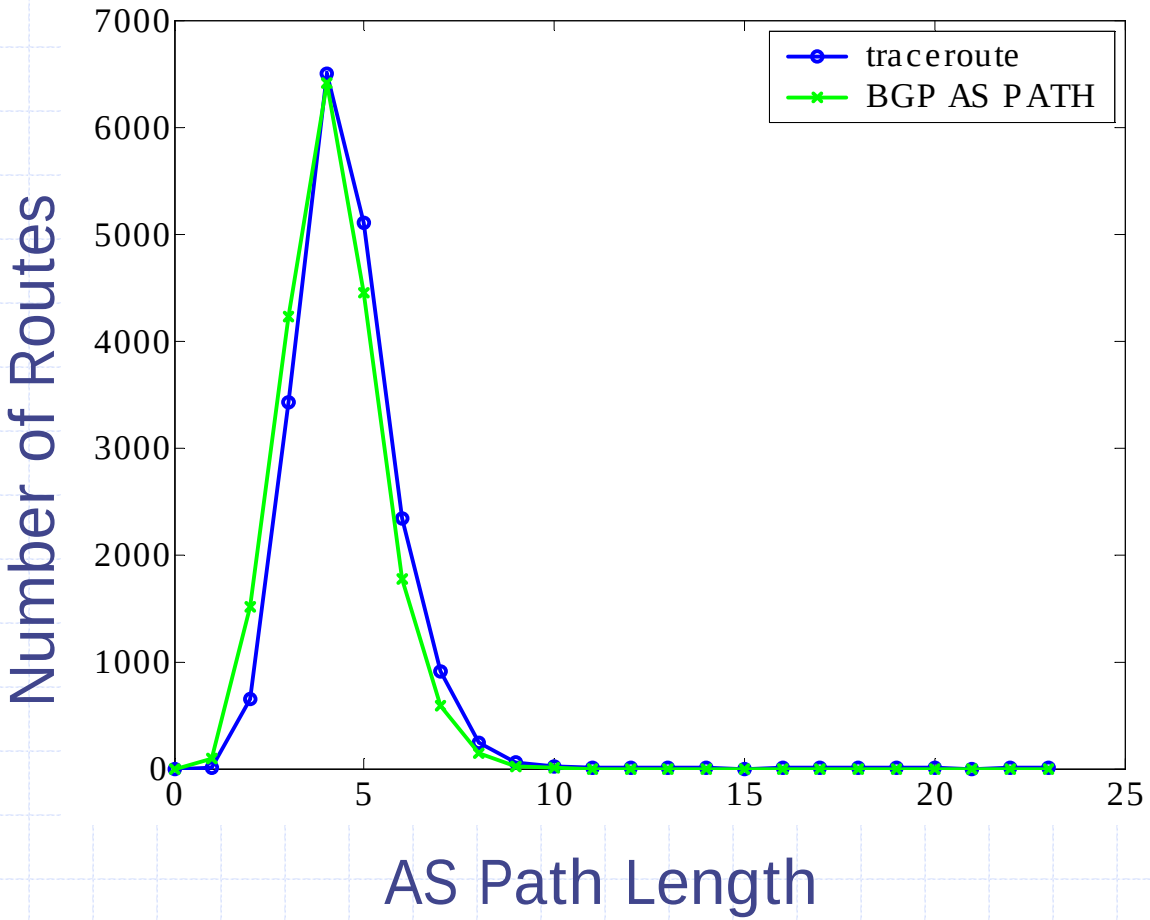

Mean Path Length
traceroute: 4.49 hops
BGP AS_PATH: 4.15 hops

32.7% of routes differed in path length

Path Component Differences

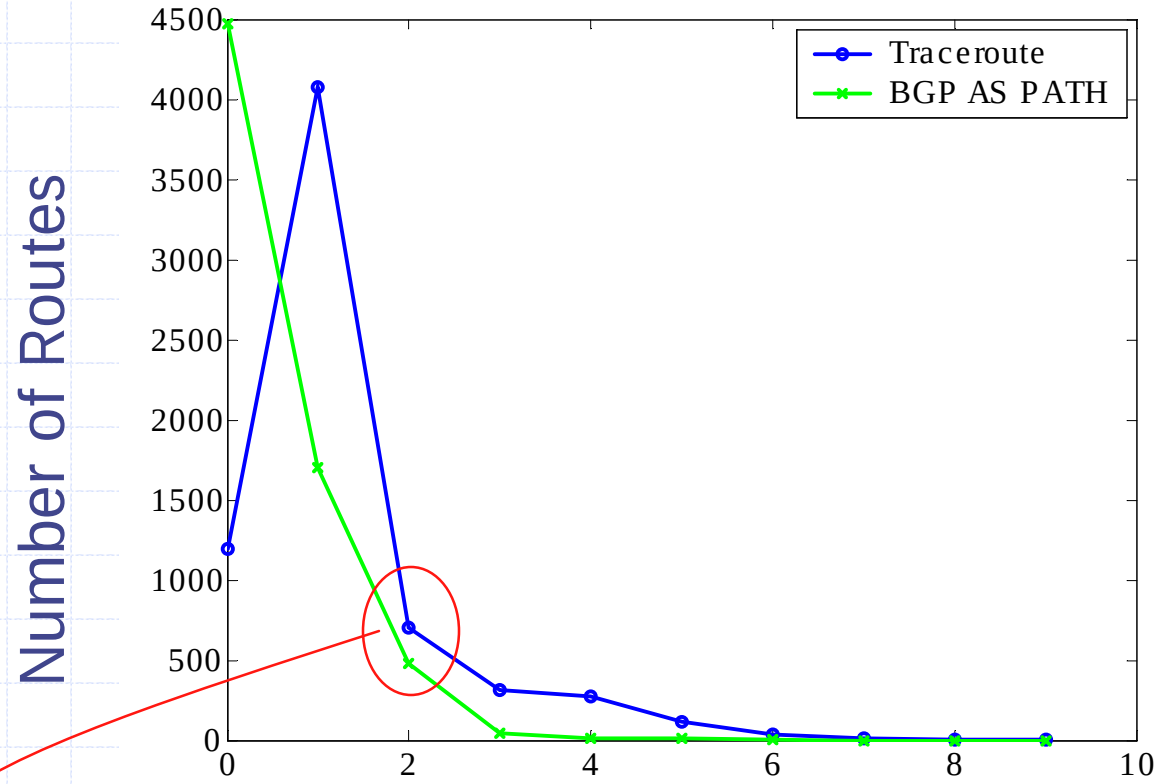

for ~500 Routes, AS_Path had 2 additional AS
for ~700 Routes, traceroute had 2 additional AS

Comparison of AS hops not represented in corresponding traceroute/BGP path (if different)

Neither BGP nor traceroute paths strictly longer/shorter

74% of differing length routes: traceroute path had a single additional node not in BGP ***AS_PATH***

AS Degree

- Internet's AS topology be represented with purely mathematical formulation?
  - Hierarchical connectivity and routing policies must be represented
  - Hierarchical representation (AS degree) conforms to power laws
- Analysis based on D3, D4 datasets
  - Nodes/edges discovered in traceroute paths
  - Nodes/edges discovered in BGP AS_PATHs
  - Nodes/edges discovered in Oregon Route View BGP table

traceroute

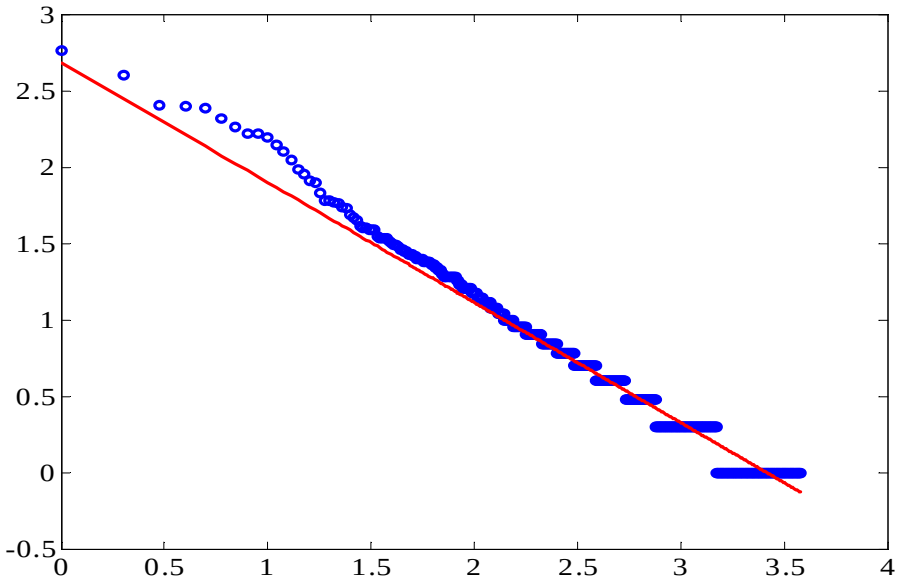

BGP AS PATH

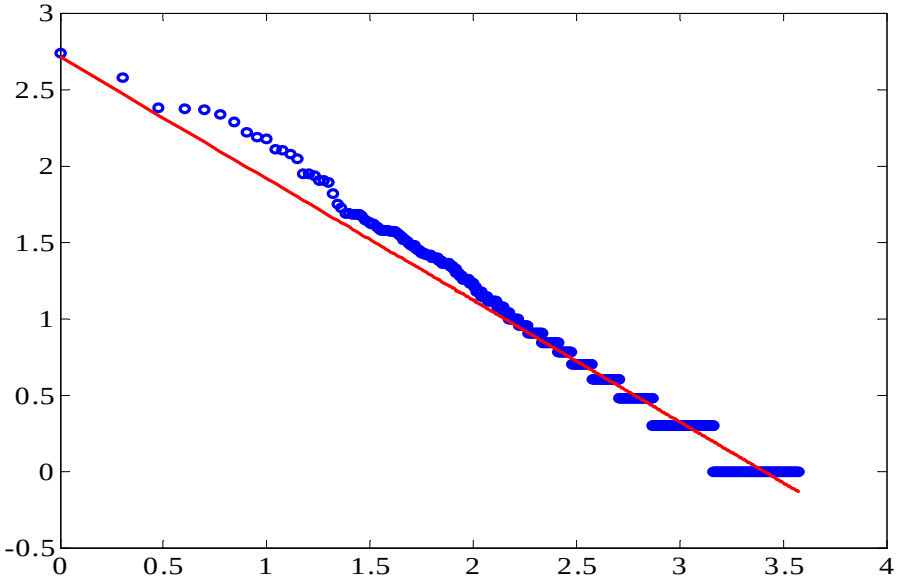

AS Degree (in $\log_{10}$)

Rank ($\log_{10}$)

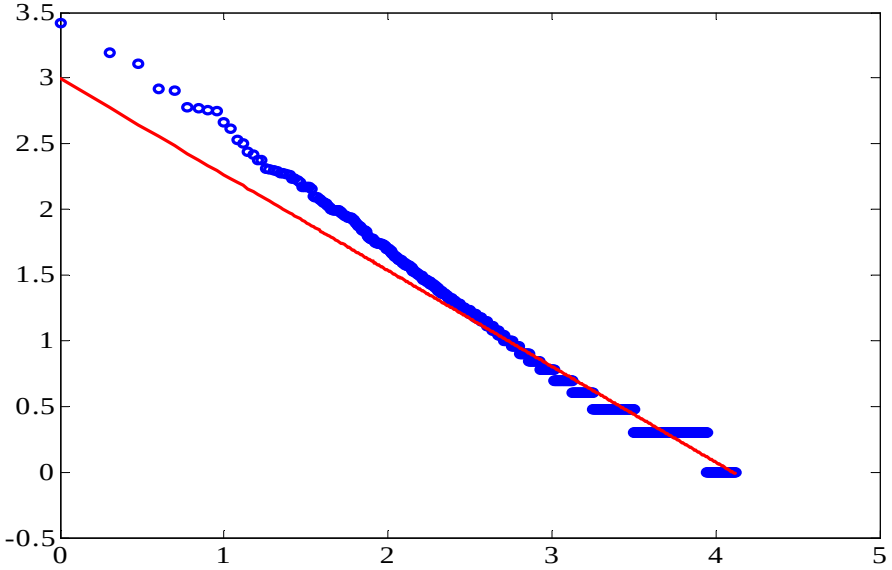

Oregon router

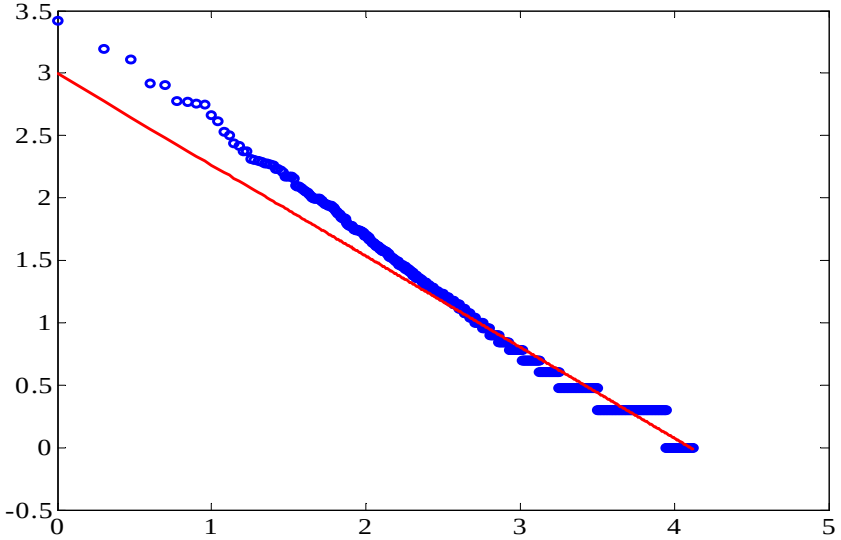

All

AS Degree Comparison

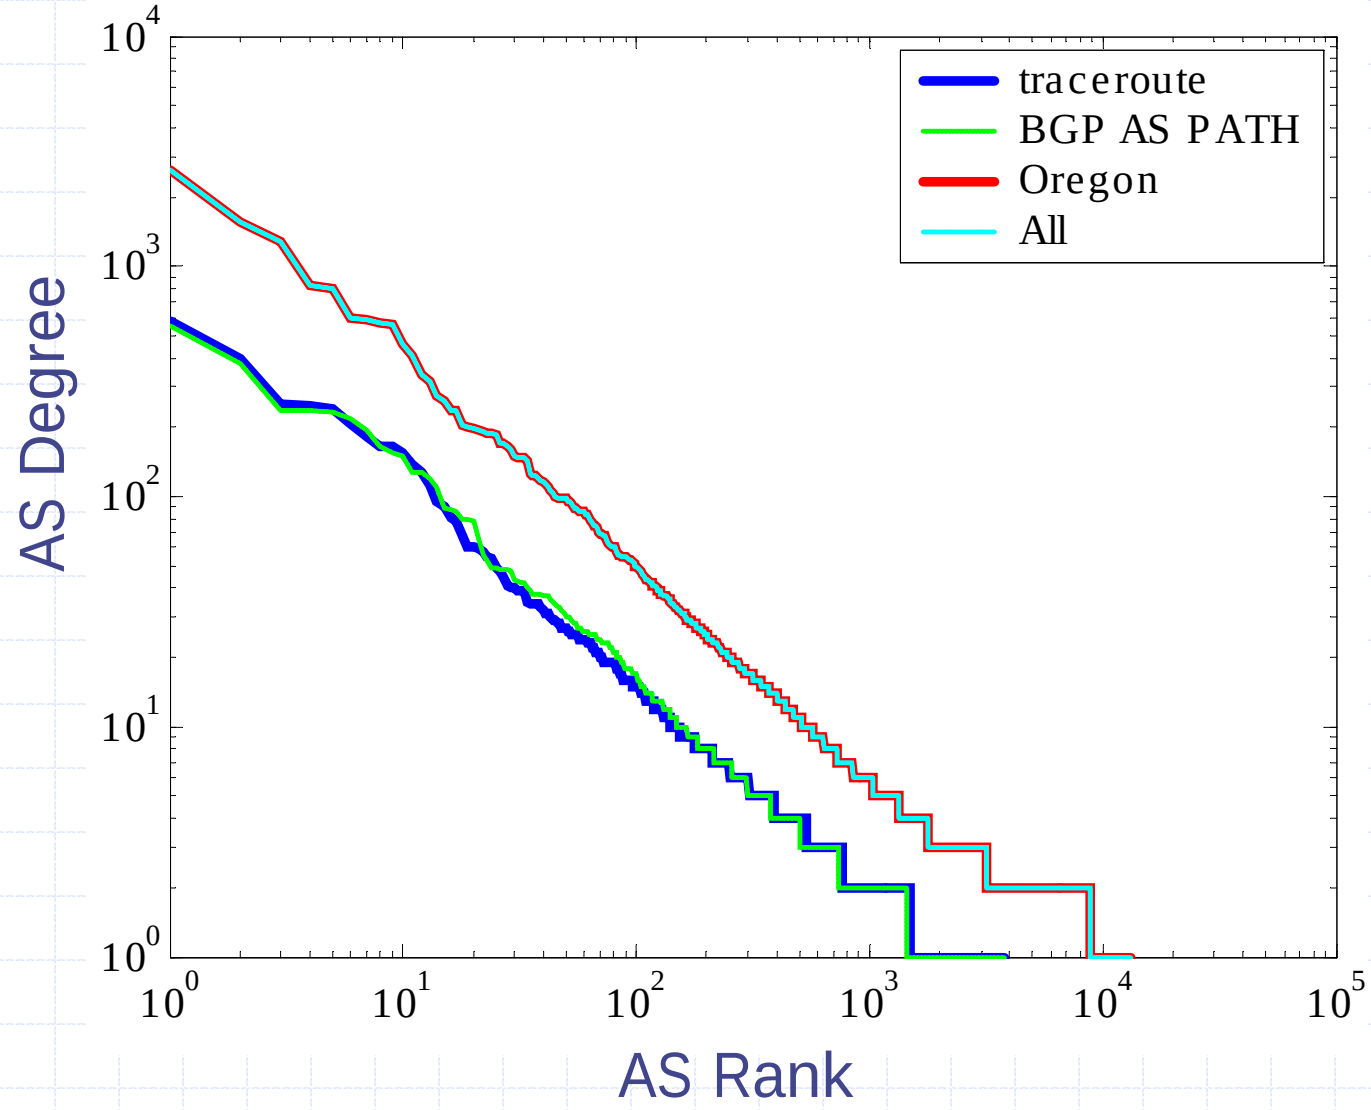

traceroute and BGP AS_PATH data almost completely overlapped

traceroute included only 18 additional nodes

BGP had ~200 edges not in `traceroute`

visual inspection: traceroute had 1 add'l node – XP.

Of 3700 BGP, traceroute nodes only 35 not in Oregon

What can we conclude?

- Advertised portion of routing policy and the traceroute behavior differ significantly
- Minimal differences in attributes representing aggregates
  - mean path length (7.5%)
  - AS degree distribution
- Current data sources not completely reliable for per path attributes
  - forward/reverse traceroute (56-69%)
  - BGP prediction of traceroute path (6-15%)