University of Nevada, Reno

Exhaustive Mapping of an Autonomous System

mcanbaz@unr.edu Muhammed Abdullah Canbaz

Comparative Analysis of Internet Topology Data sets

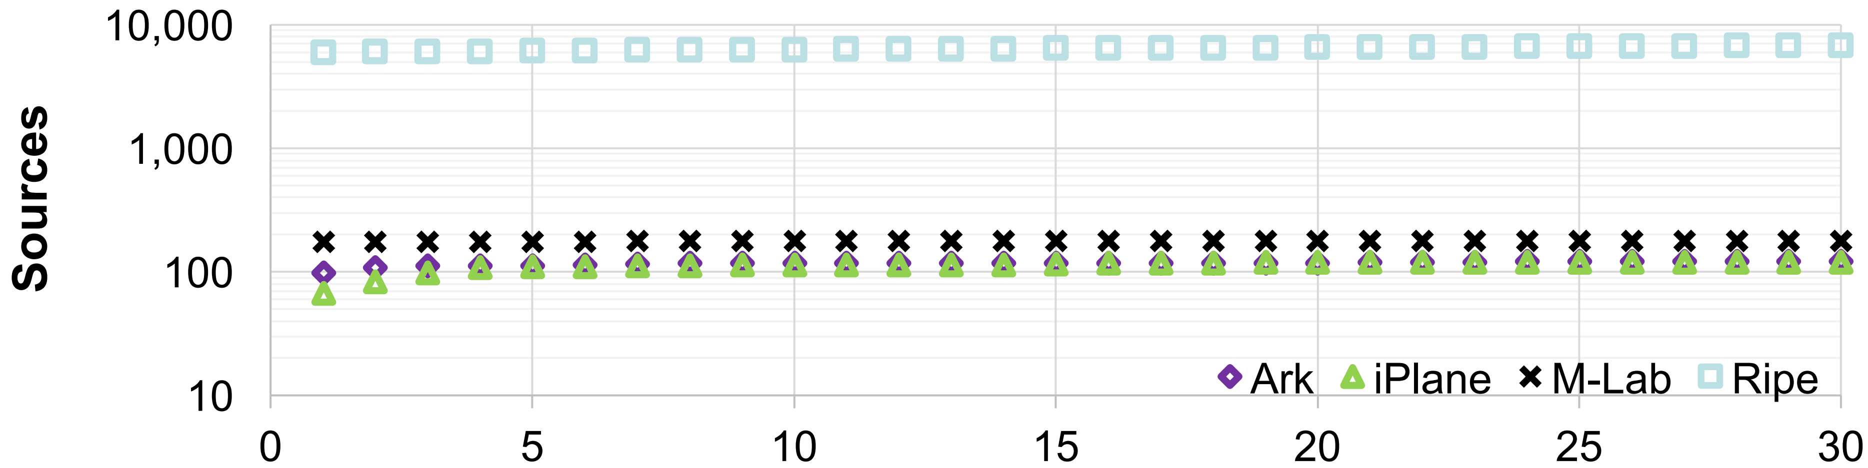

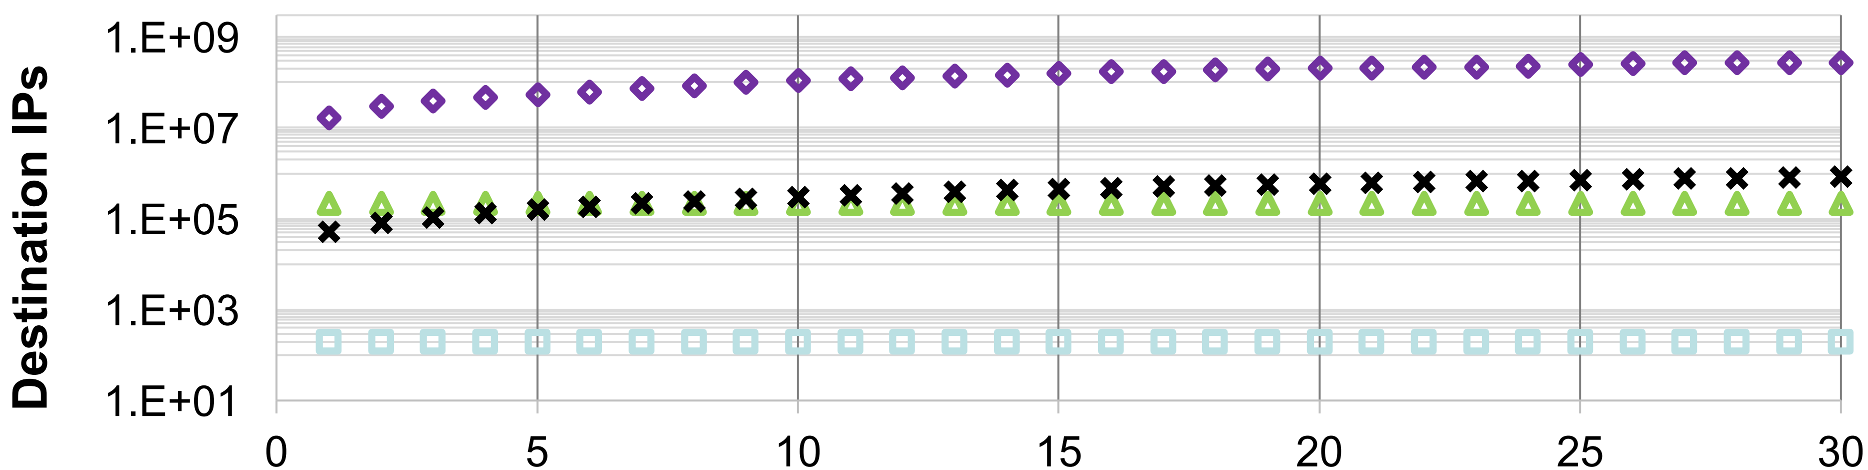

Comparative Analysis of Internet Topology Data sets

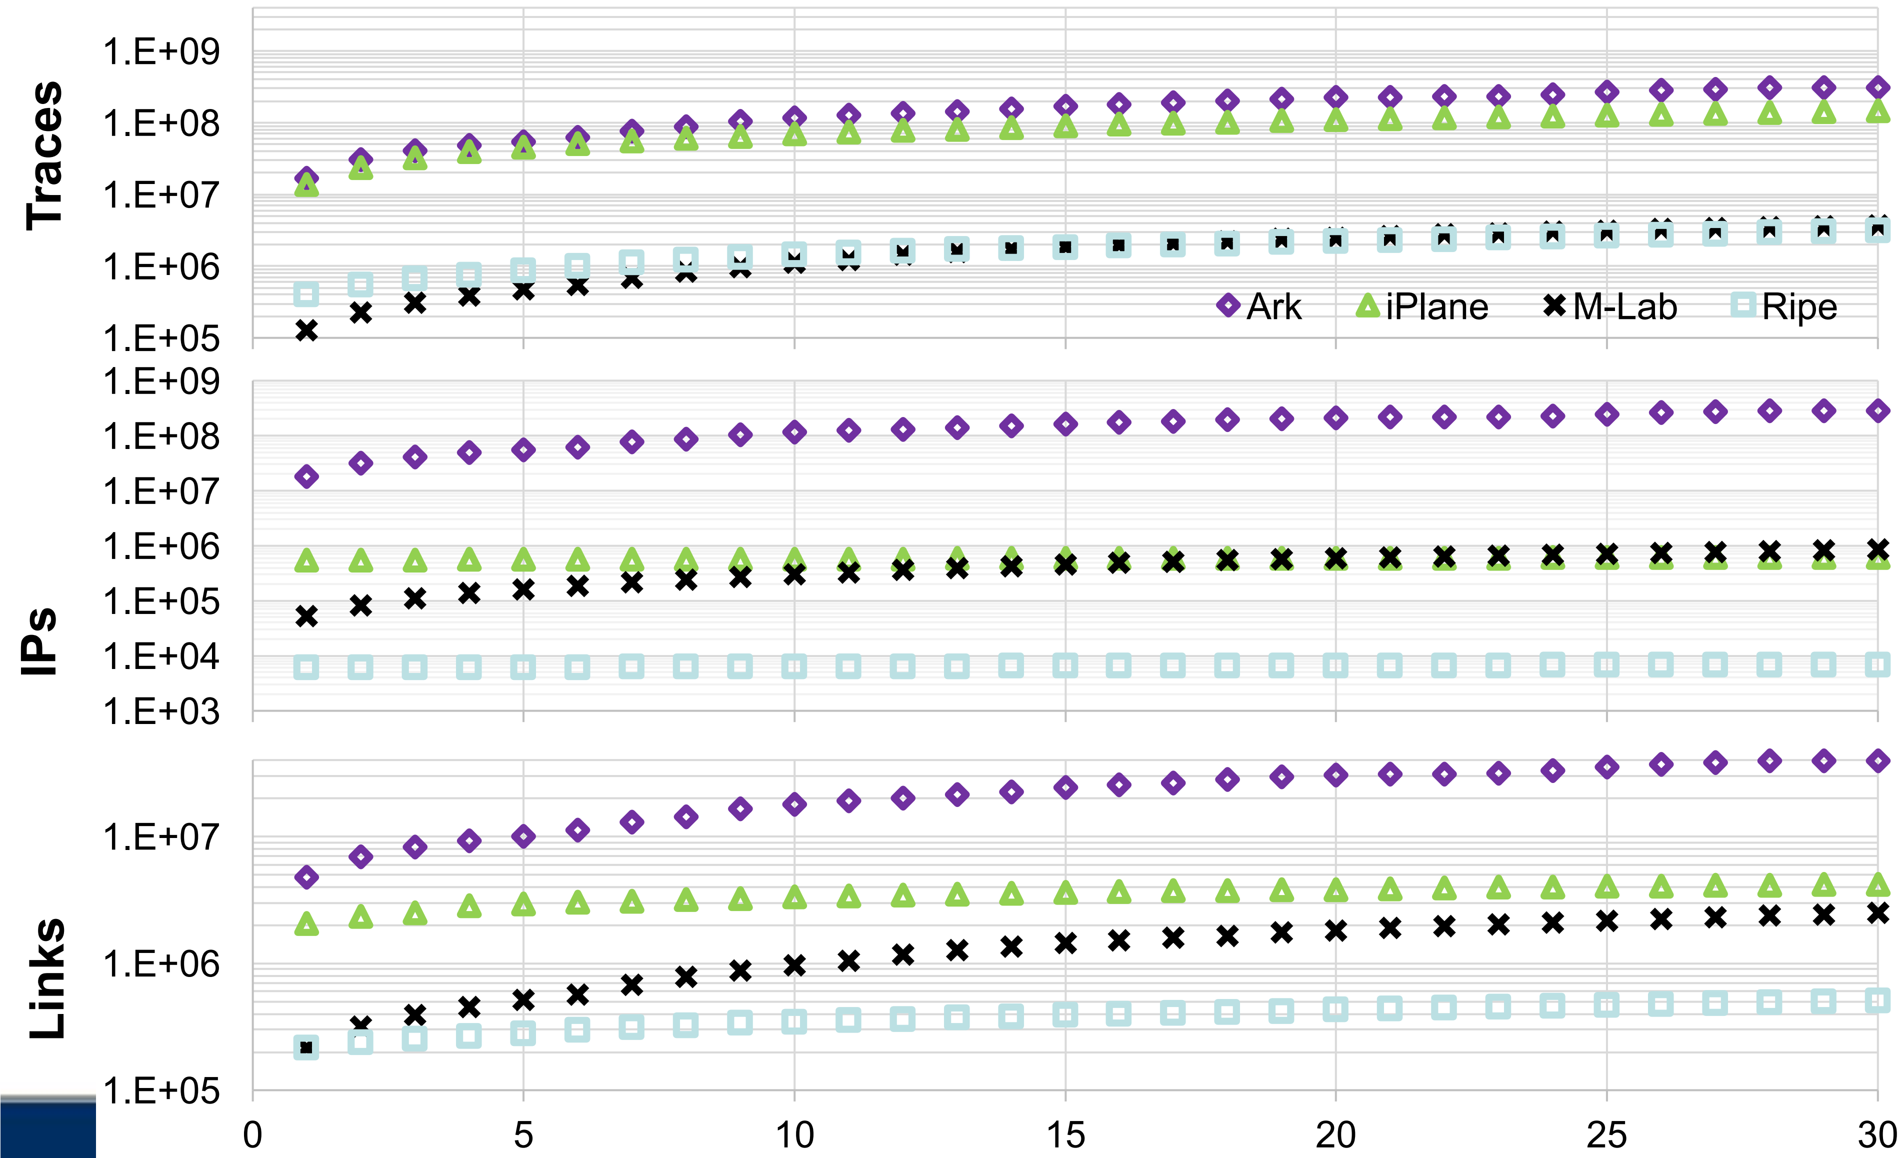

Data coverage (5 days)

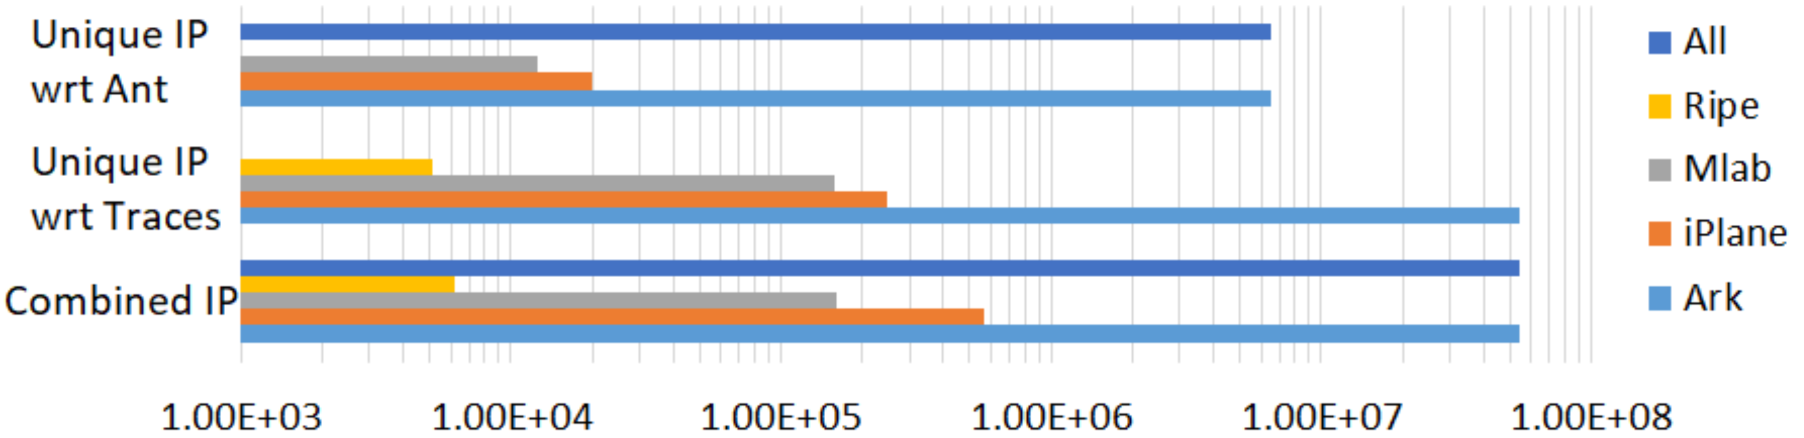

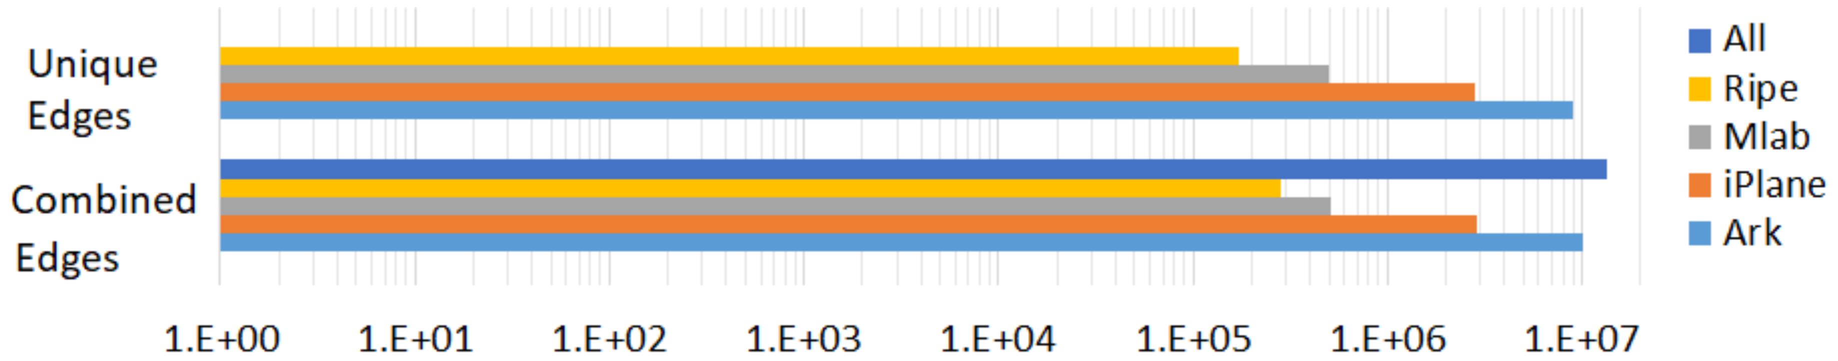

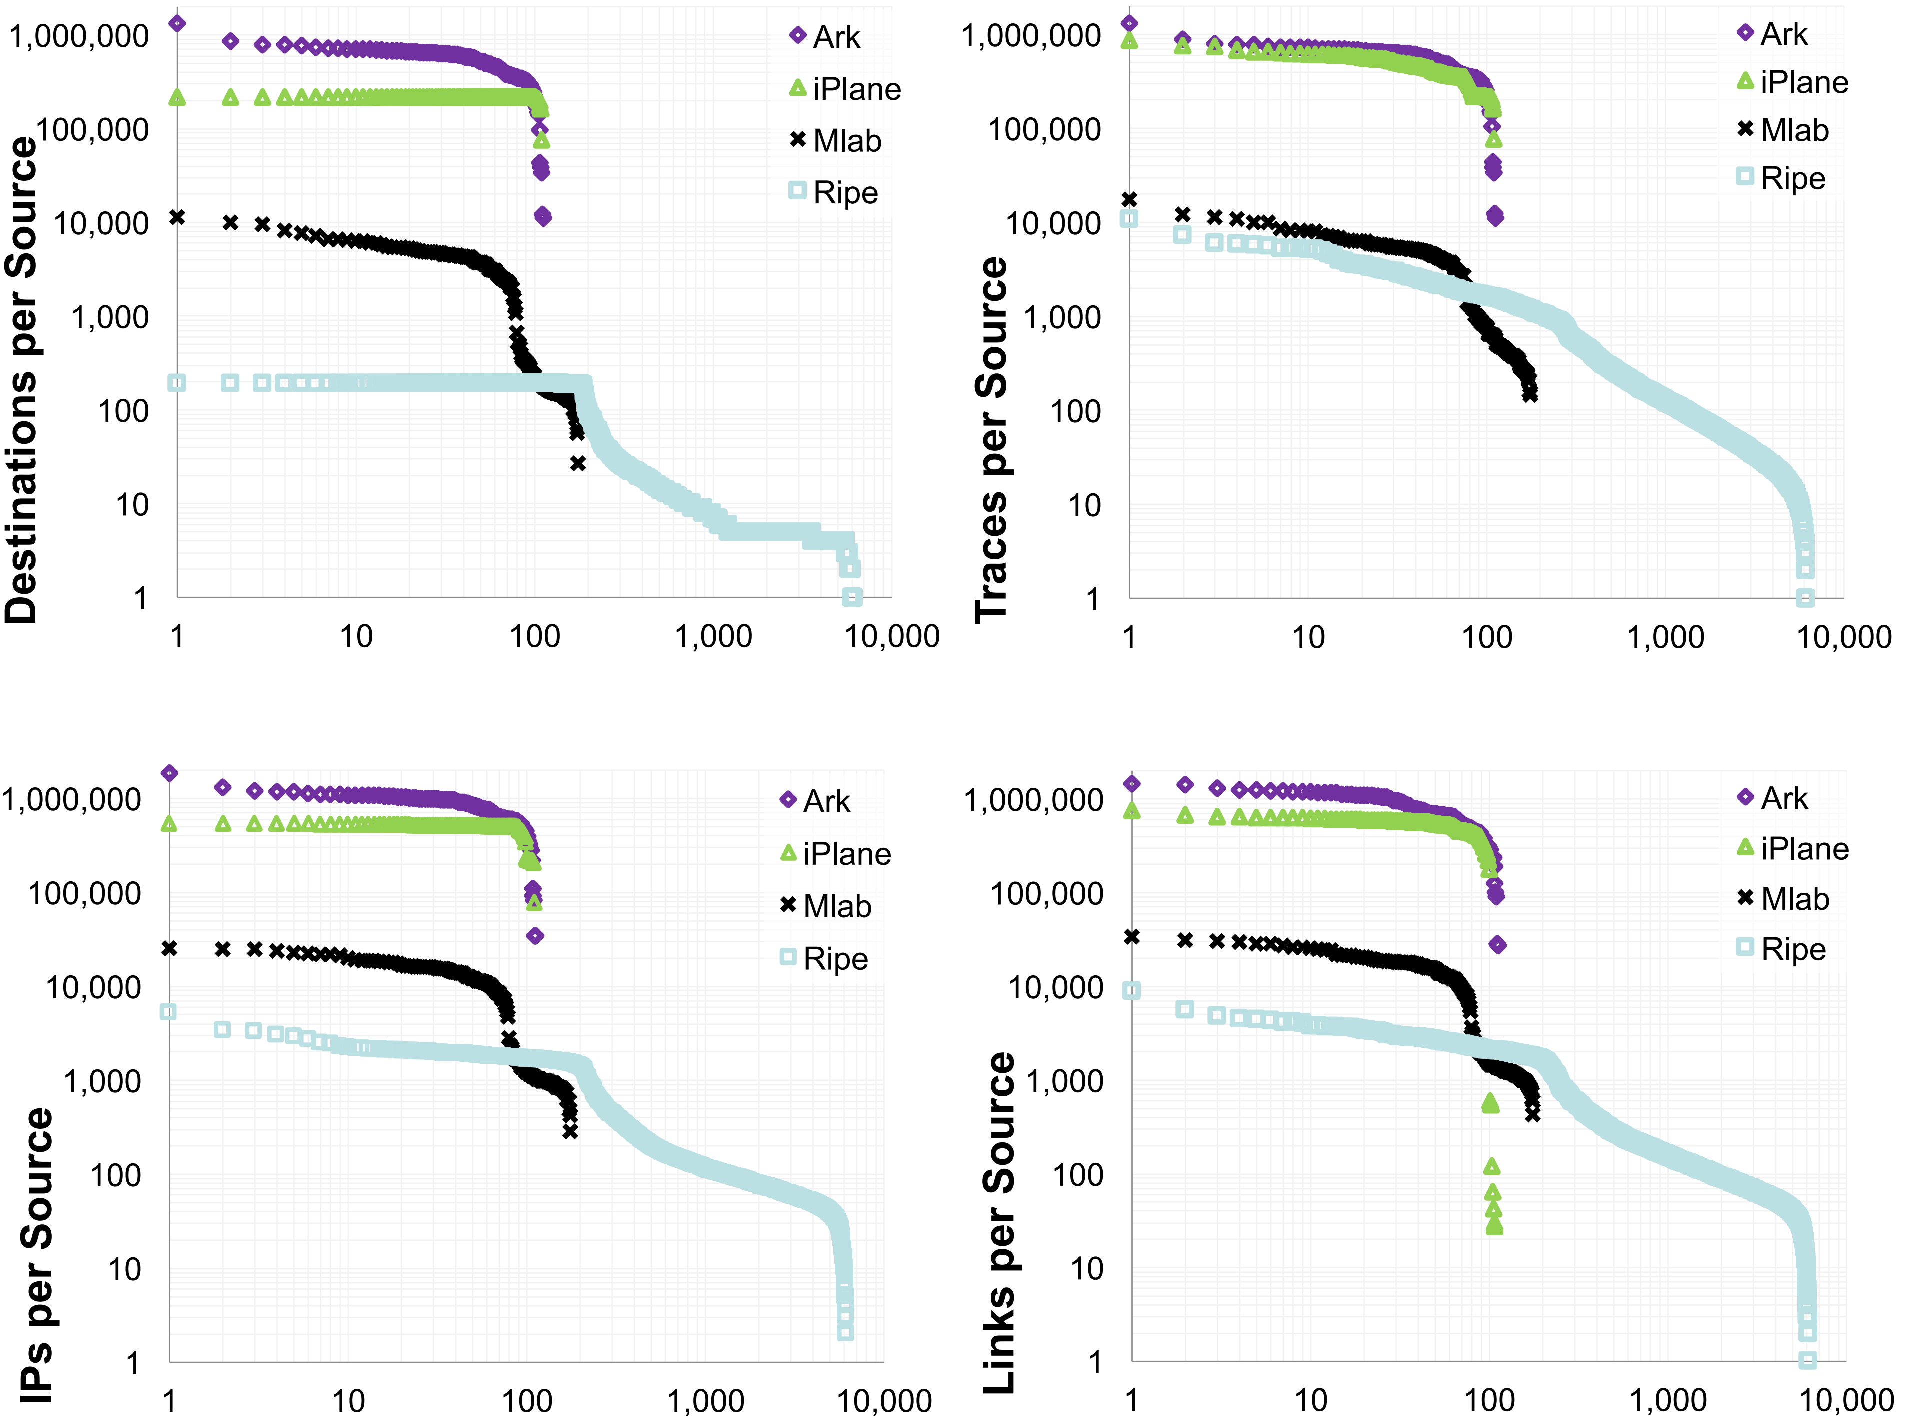

Destinations per Source
Traces per Source
IPs per Source
Links per Source
Ark
iPlane
Mlab
Ripe
1,000,000
100,000
10,000
1,000
100
10
1
1
10
100
1,000
10,000

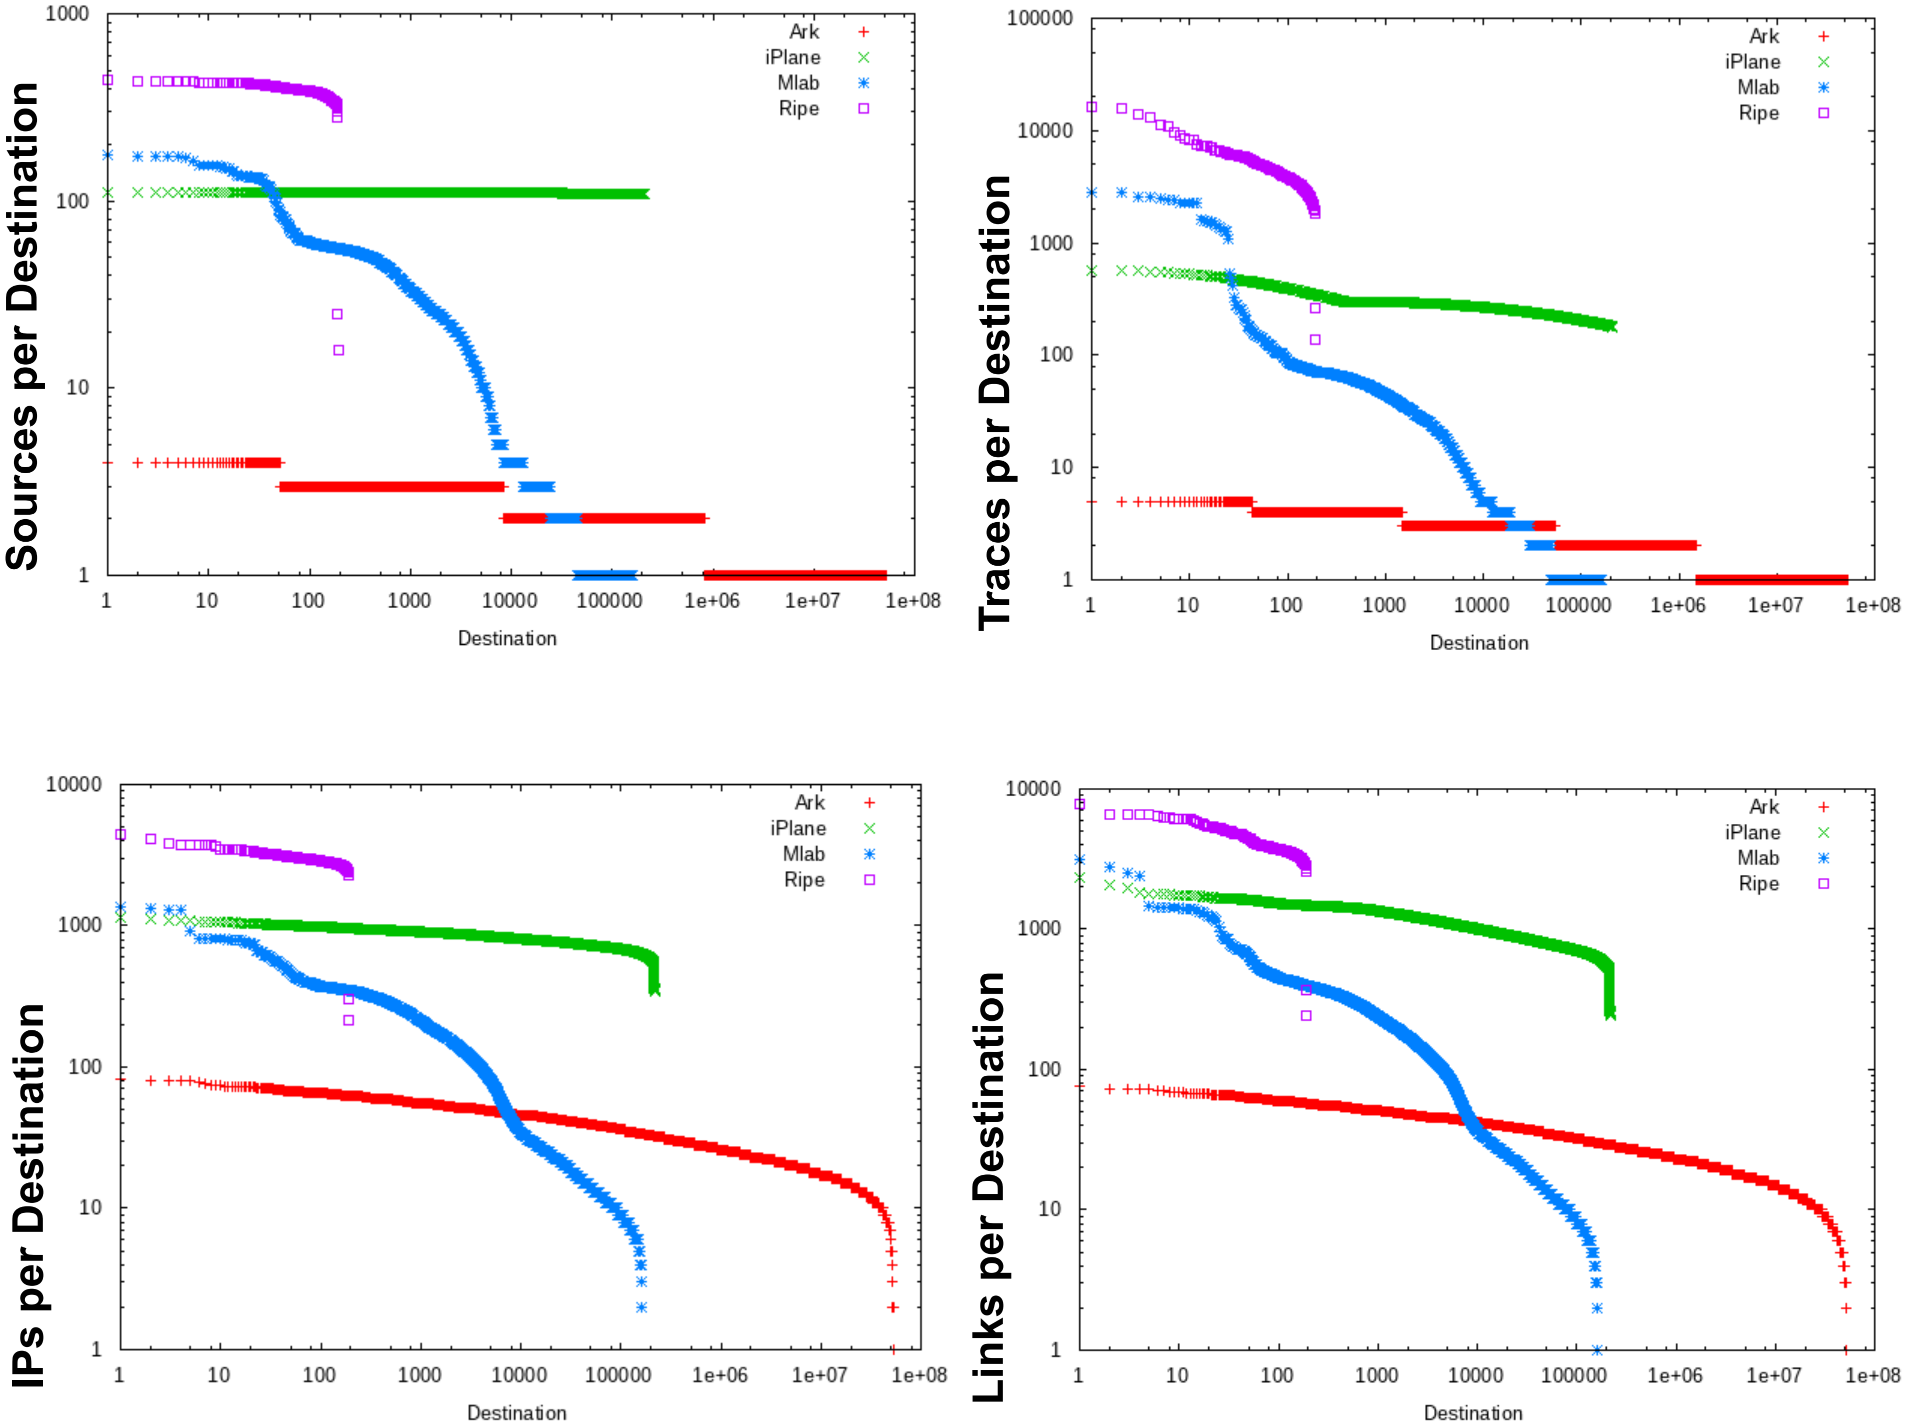

Sources per Destination
Traces per Destination
IPs per Destination
Links per Destination
Destination
Ark
iPlane
Mlab
Ripe
1
10
100
1000
10000
100000
1e+06
1e+07
1e+08

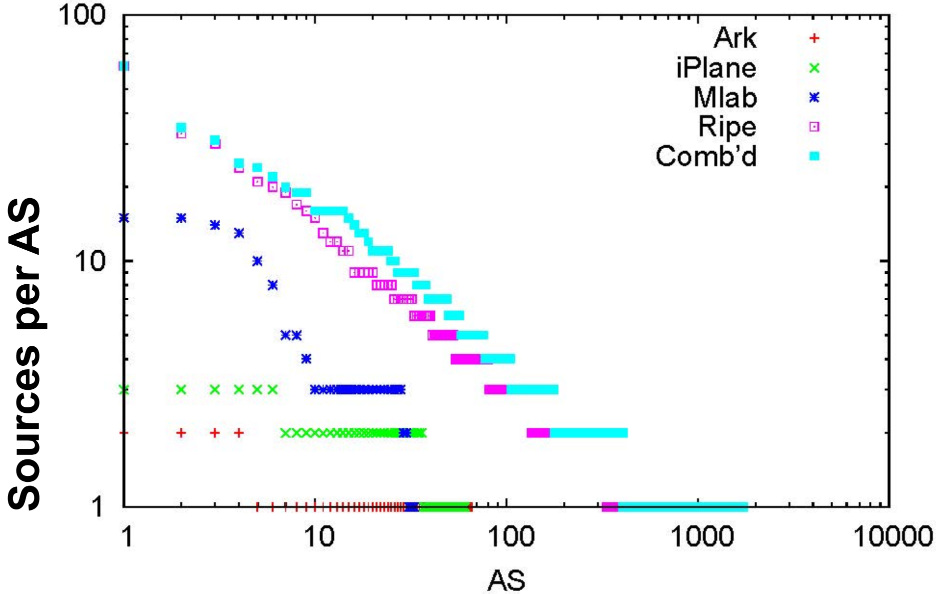
Sources per AS
AS
Ark
iPlane
Mlab
Ripe
Comb'd
1
10
100
1000
10000

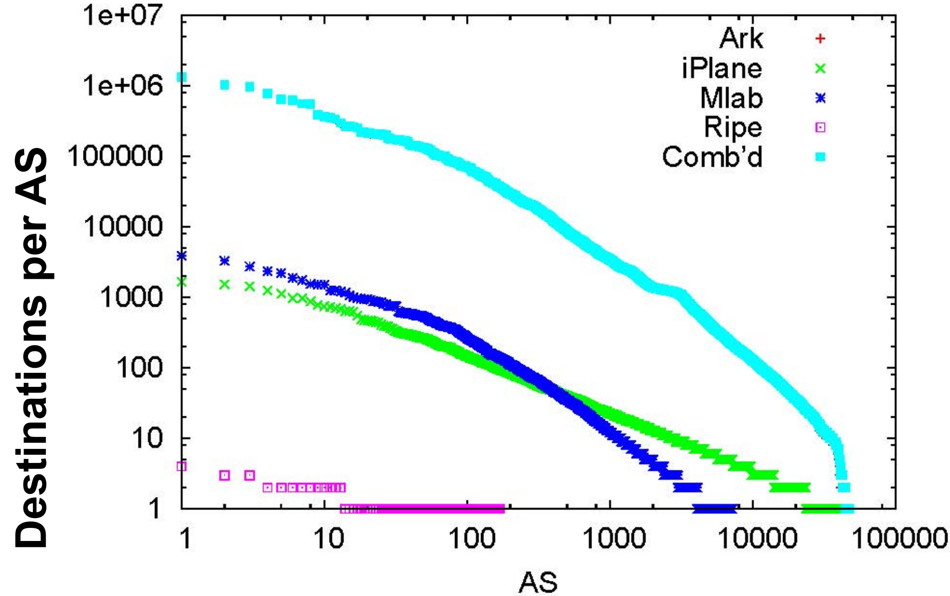
Destinations per AS
AS
Ark
iPlane
Mlab
Ripe
Comb'd
1e+07
1e+06
100000
10000
1000
100
10
1
10
100
1000
10000
100000

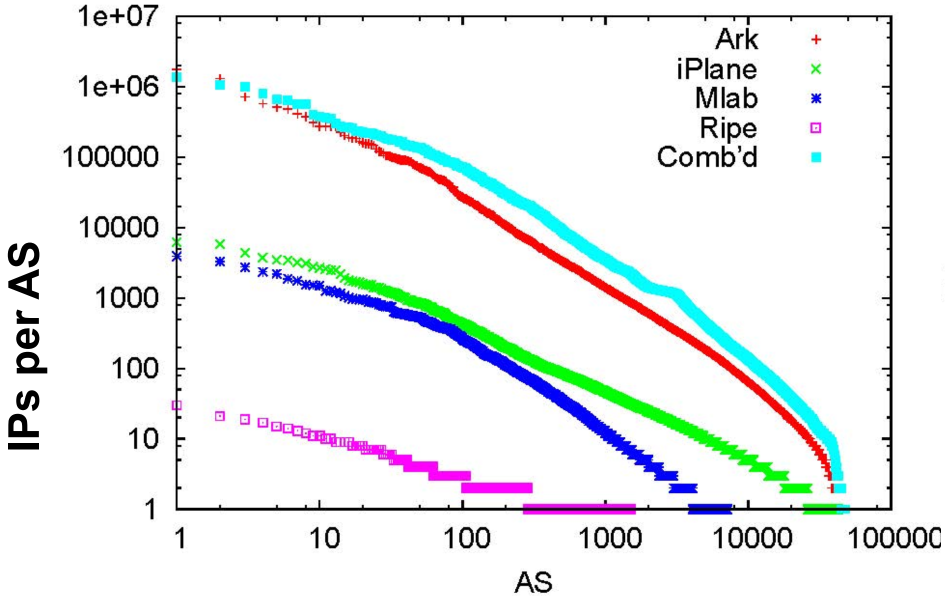
IPs per AS
AS
Ark
iPlane
Mlab
Ripe
Comb'd
1e+07
1e+06
100000
10000
1000
100
10
1
10
100
1000
10000
100000

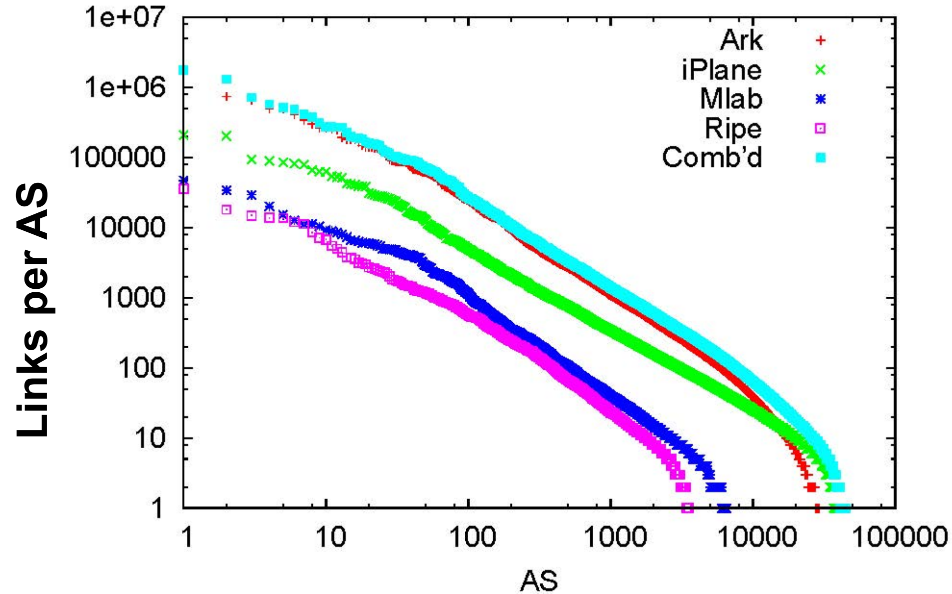
Links per AS
AS
Ark
iPlane
Mlab
Ripe
Comb'd
1e+07
1e+06
100000
10000
1000
100
10
1
10
100
1000
10000
100000

Summary

- The public measurement data sets provided by Ant census, Ark, CAIDA, CIDR, iPlane, IRL, M-Lab, and Ripe is analyzed
  - *Each dataset provided a unique topological perspective*
    - *links discovered by different platforms are often unique*

When probing:

- Ark targets a randomized and increasing set of destinations
  - the most comprehensive coverage of the public measurement platforms we survey
- If focused on AS reachability and node/link discovery, Ark provides the most diverse traces.
- IPlane uses same set of destinations every day (RIP iPlane)
- Ripe would be beneficial in studies focusing on network latency, throughput, and bandwidth measurements

In terms of vantage points:

- iPlane and Mlab both utilize nodes from the PlanetLab for vantage points
  - from 1100 nodes at its peak, to less than 100 nodes presently
- Ripe consists of about 60 times more nodes than the others
  - but has limited topology coverage due to limited set of destinations

Internet Trace Dynamics

Path traces from RIPE Atlas

- **for a month (May 1-31, 2017)**
- **219,768,505** path traces from **5,901** source nodes towards **183** destinations
- IP connectivity information of **2,237 ASes**

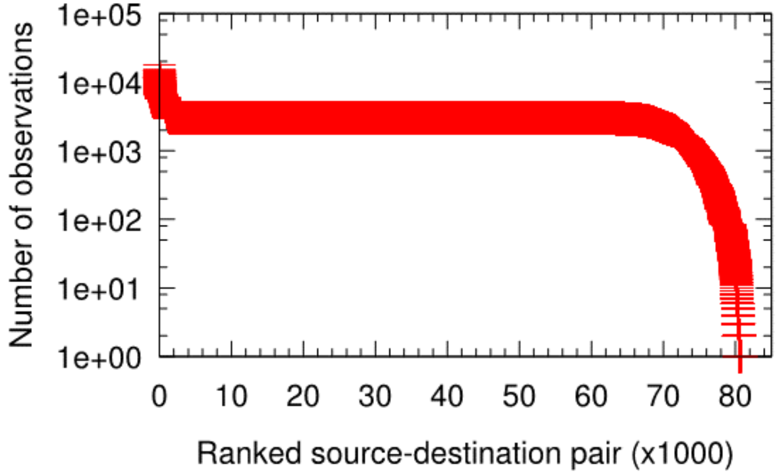

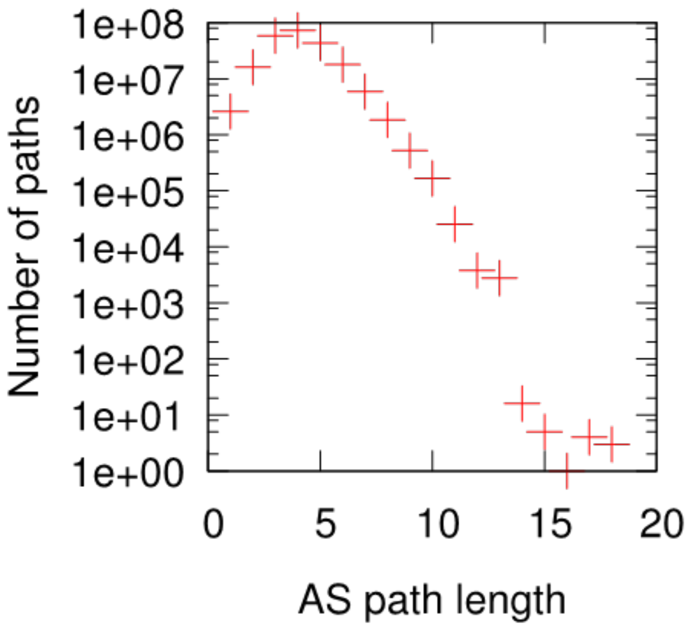

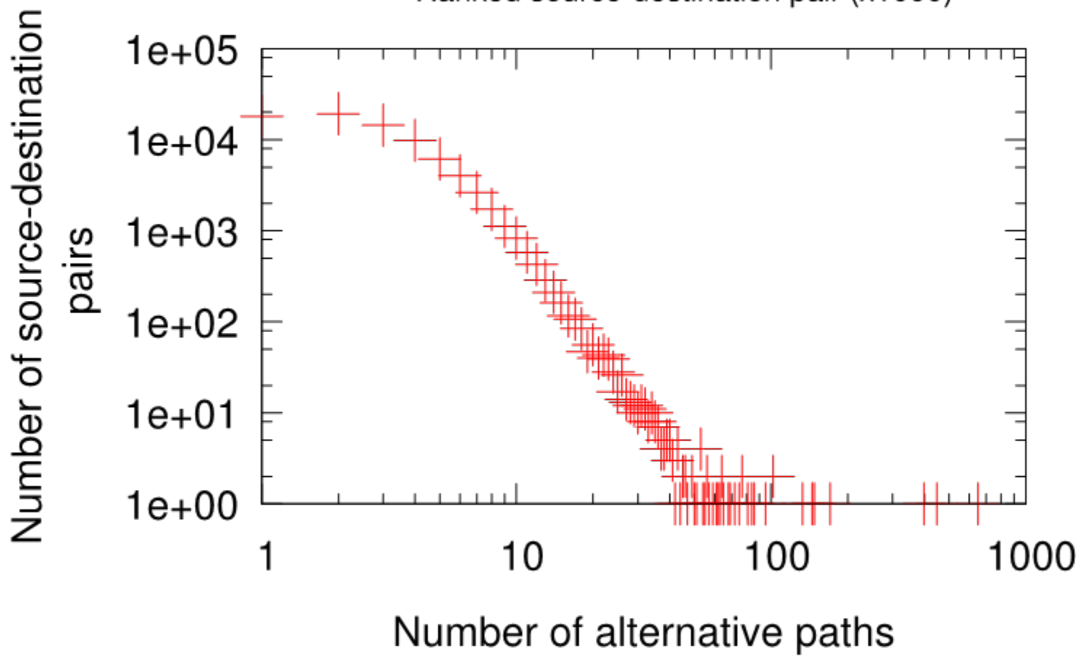

Anomalies (loops and repeats)

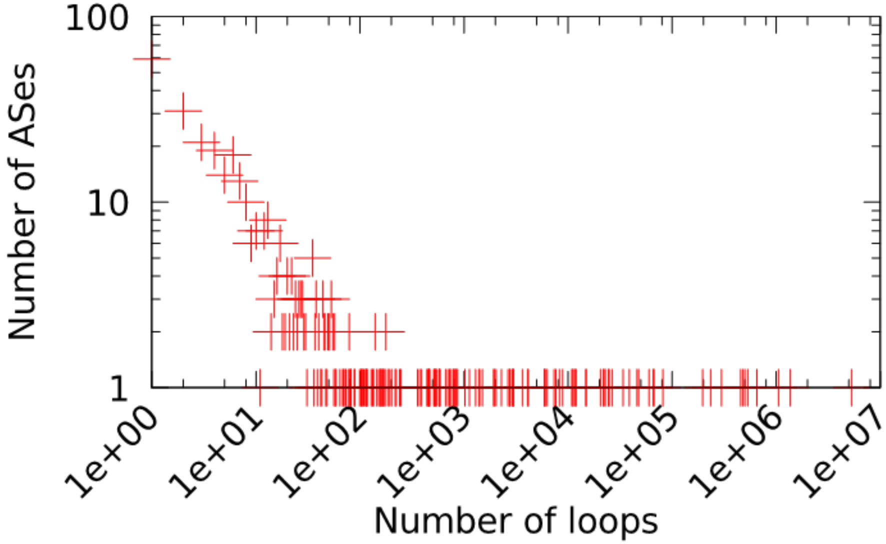

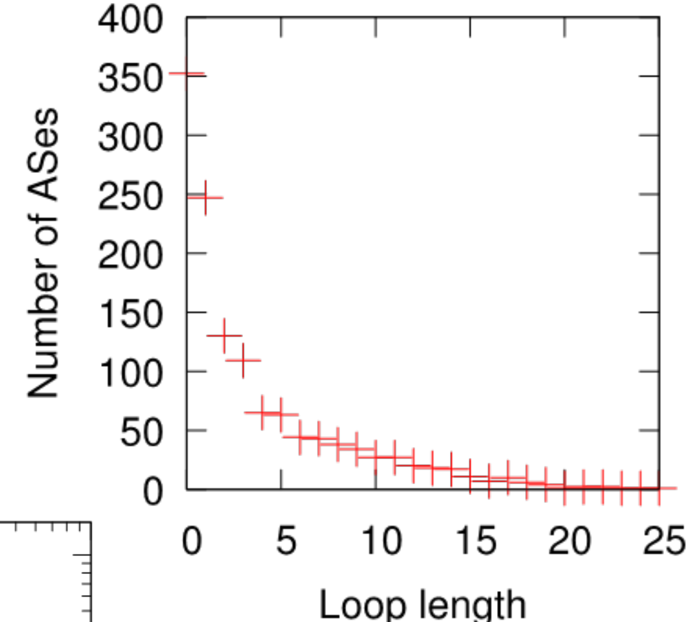

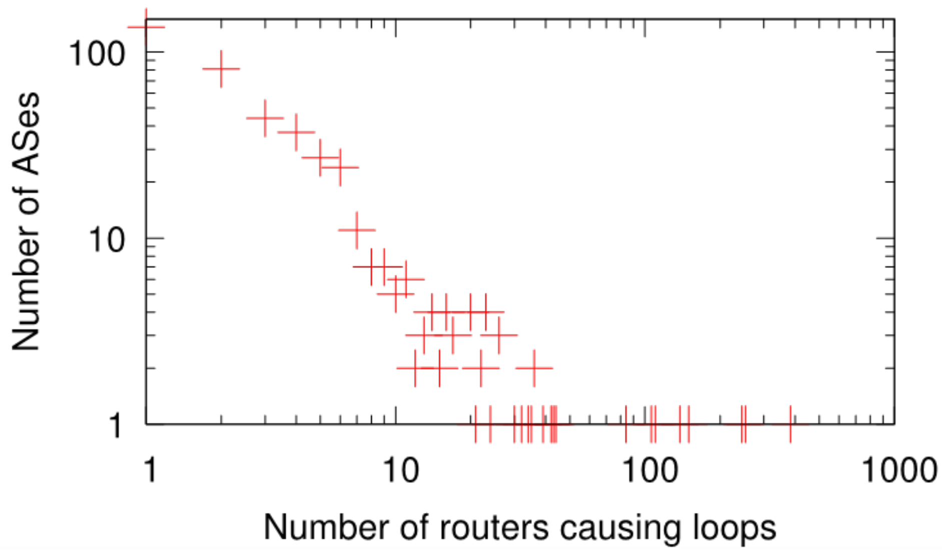

Path Observation

By frequency

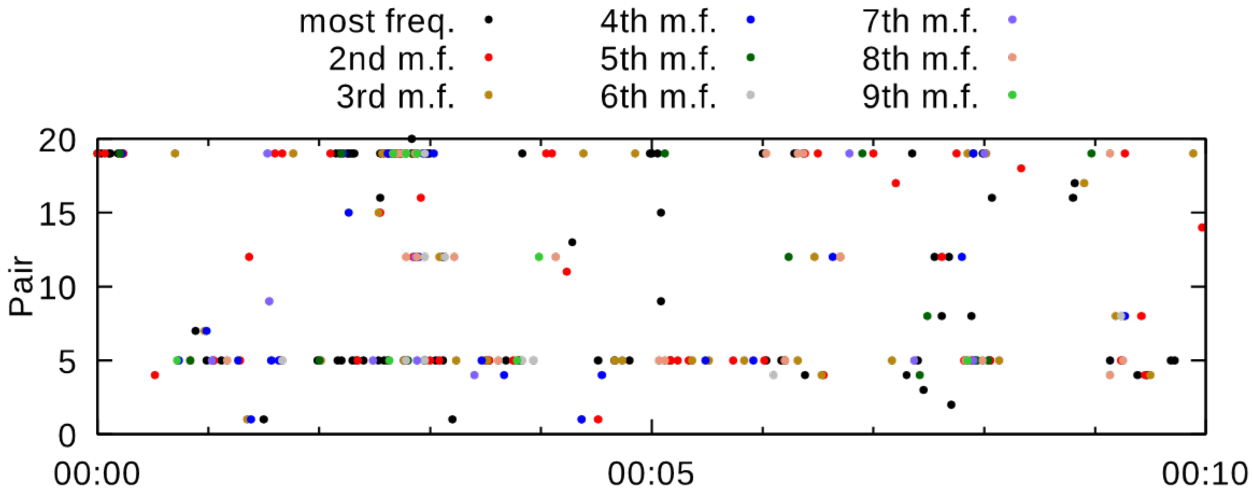

By path length

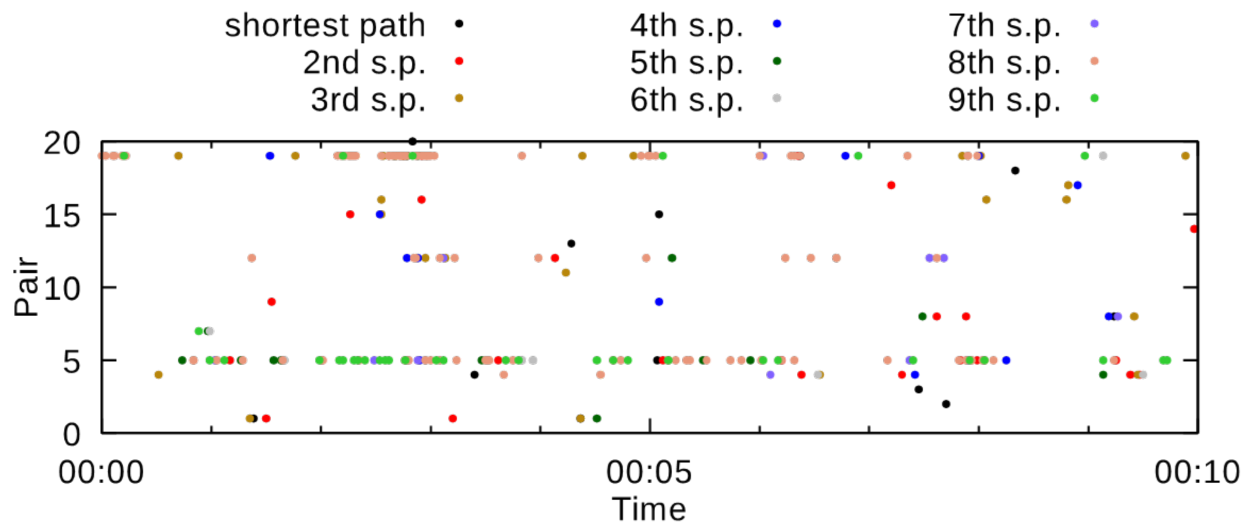

AS 2603

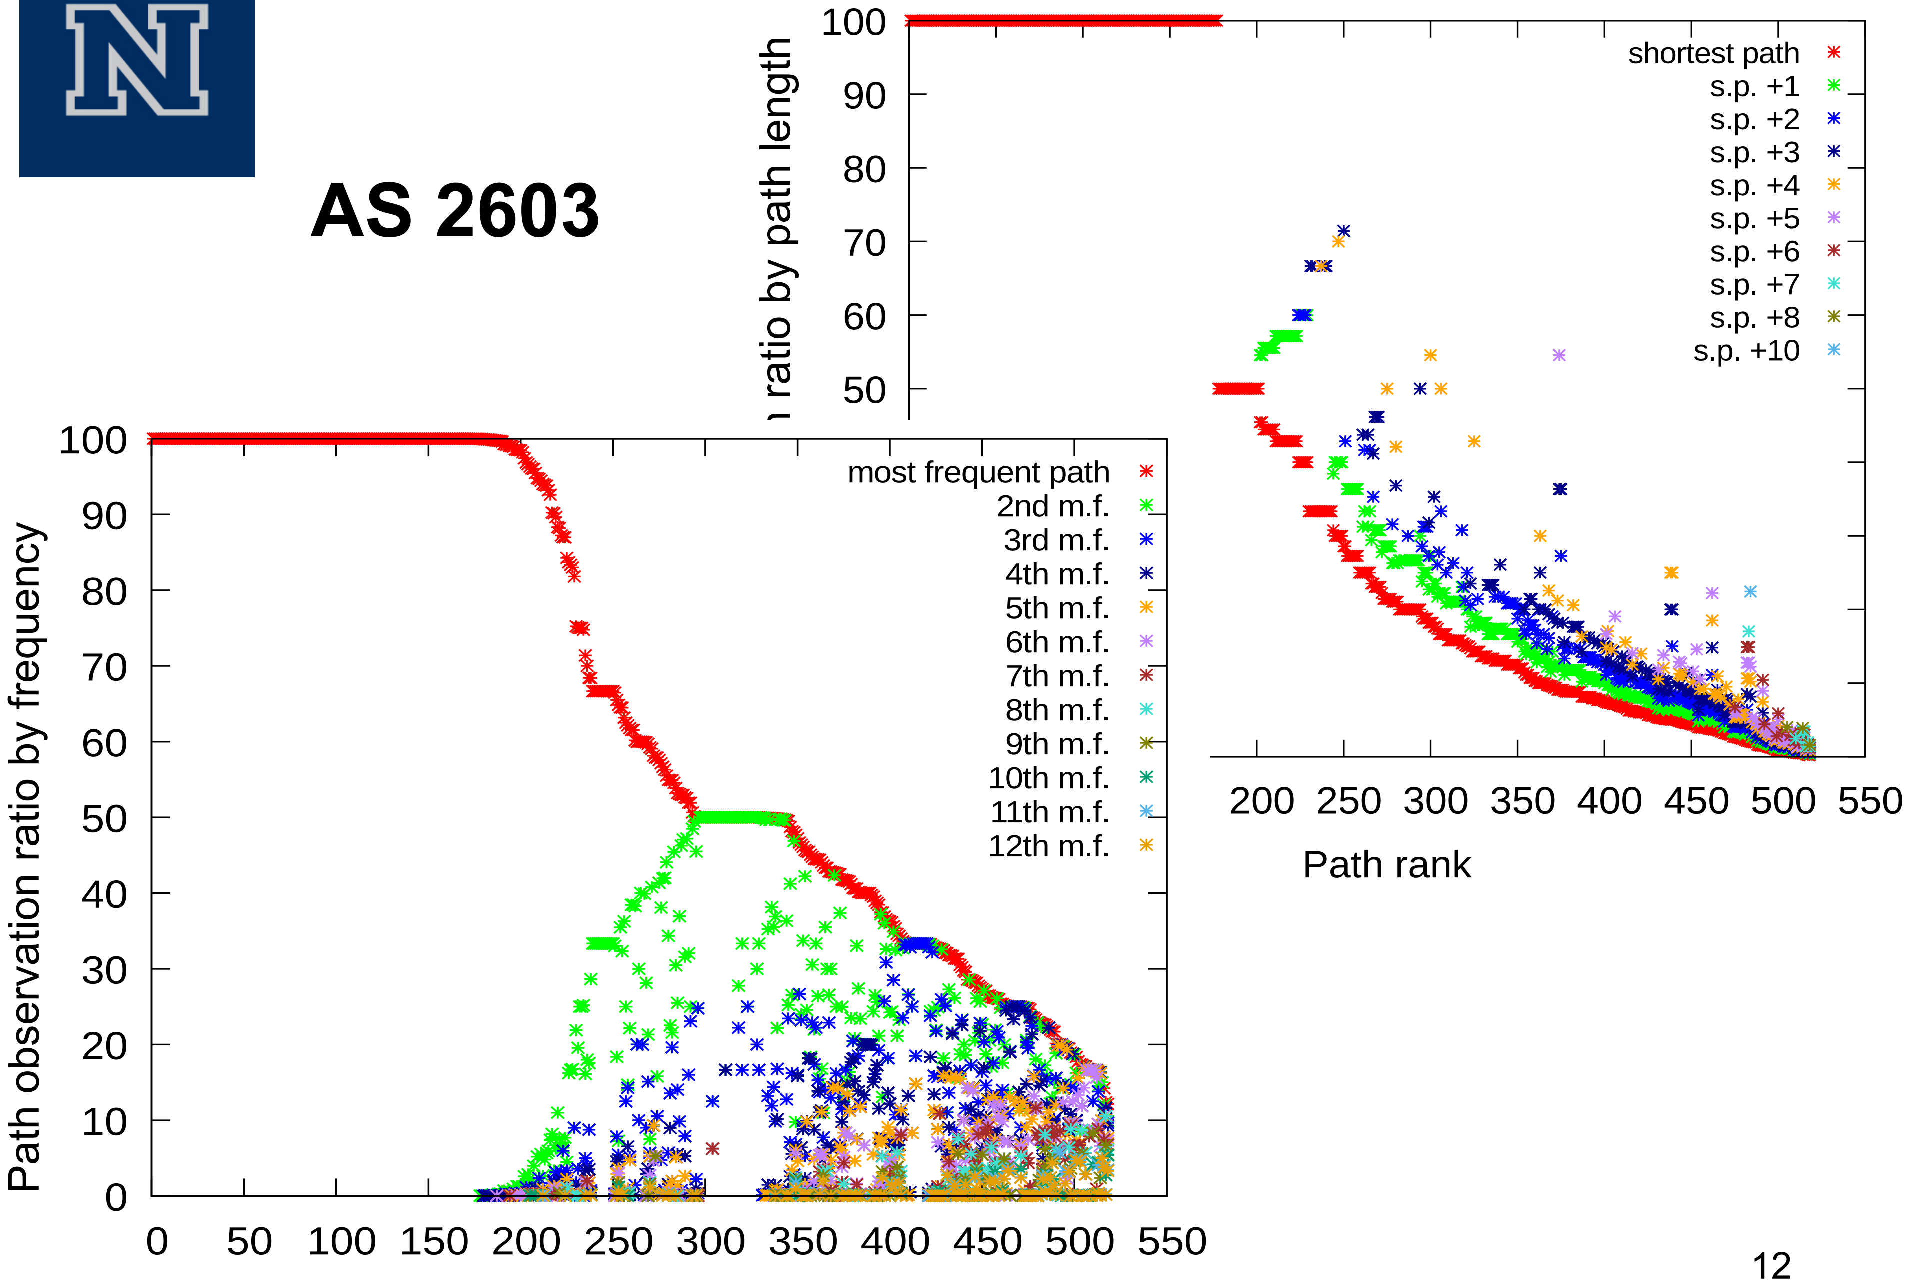

ratio by path length
100
90
80
70
60
50
shortest path
s.p. +1
s.p. +2
s.p. +3
s.p. +4
s.p. +5
s.p. +6
s.p. +7
s.p. +8
s.p. +10
200 250 300 350 400 450 500 550
Path rank
Path observation ratio by frequency
100
90
80
70
60
50
40
30
20
10
0
most frequent path
2nd m.f.
3rd m.f.
4th m.f.
5th m.f.
6th m.f.
7th m.f.
8th m.f.
9th m.f.
10th m.f.
11th m.f.
12th m.f.
0 50 100 150 200 250 300 350 400 450 500 550

AS Rank by Path Observation

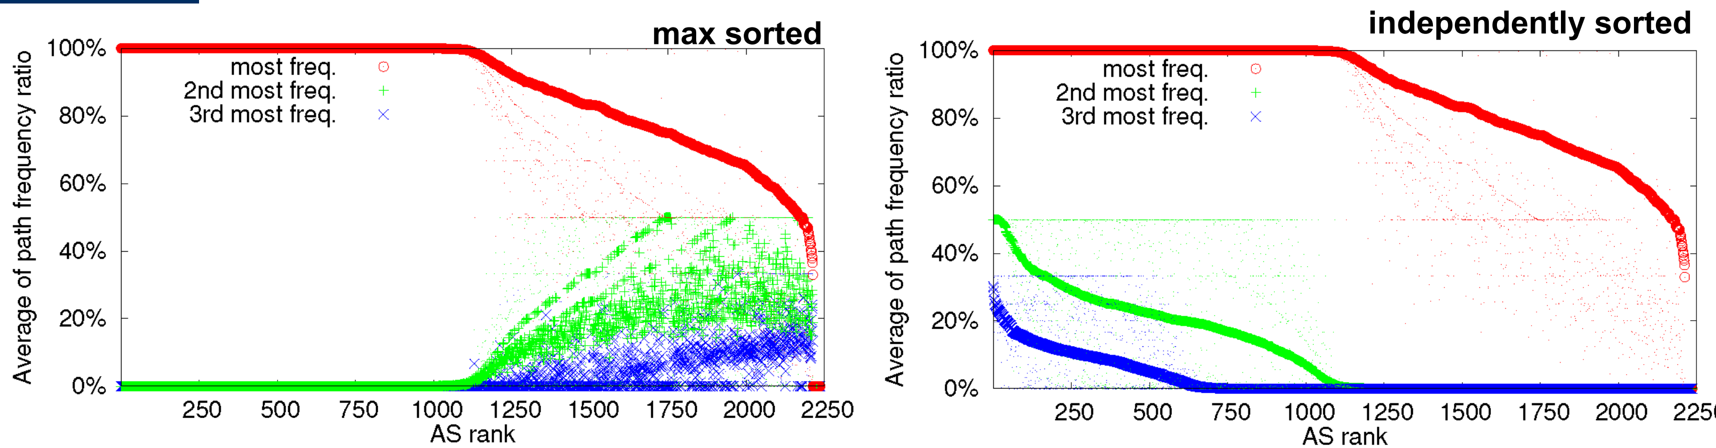

Path observation ratio averages of frequent paths for all ASes by observation frequency

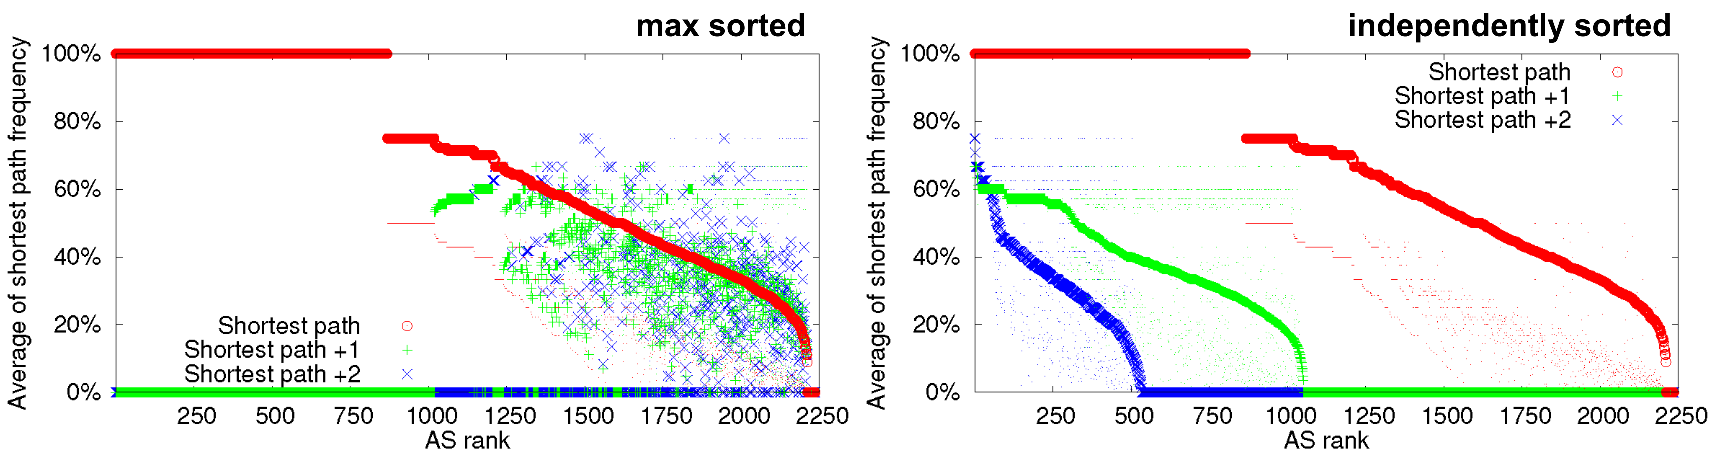

Path observation ratio averages of shortest paths for all ASes by observation frequency

AS Rank by Path Observation

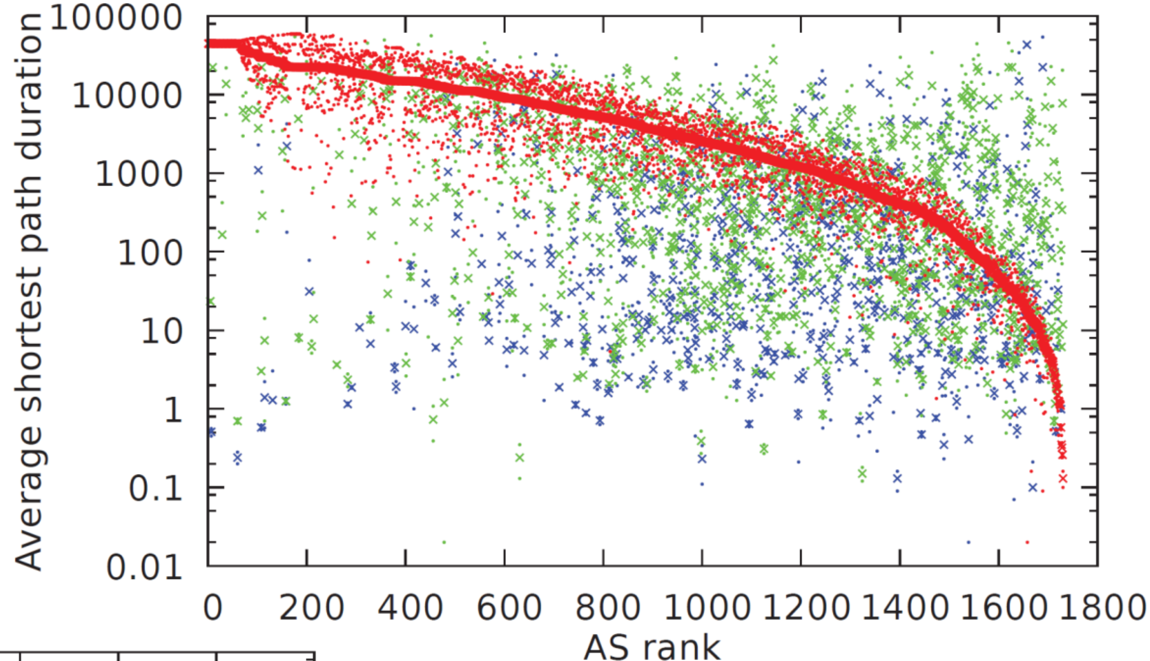

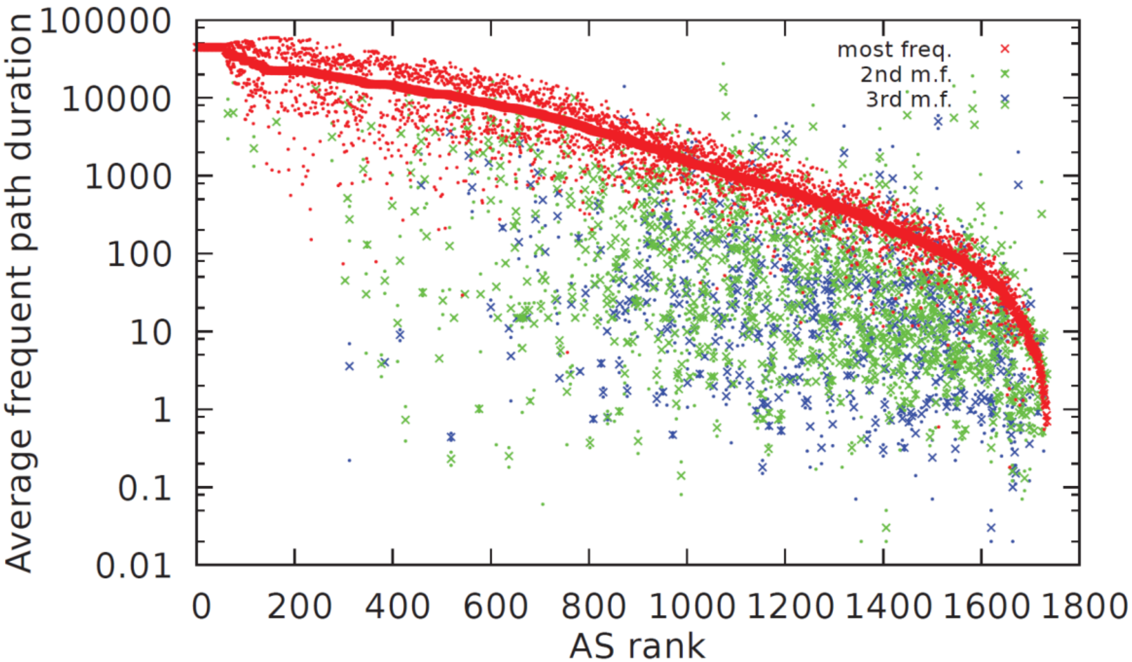

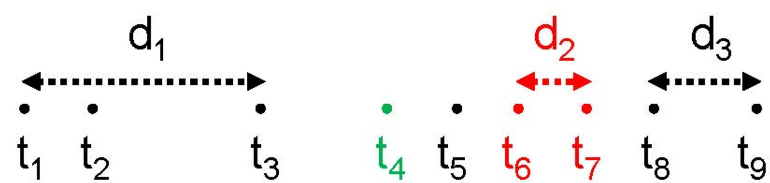

Path duration calculation

Summary

- Intra-Domain Paths: About half the ASes have a single path between all of its ingress/source and egress/destination pairs
- Inter-Domain Paths: Only 22.4% paths cross the same ASes
- ASes have more stable paths internally, end-to-end paths are more dynamic
  - Hot potato routing
  - Earlier studies focused on end-to-end paths and found *highly dynamic paths*
    - This behavior is due to the BGP level dynamism between ASes rather than the *router level dynamism* within AS networks
- Path selection and durations vary significantly
  - from observation period of *a month* to less than *a minute*
- Many ASes distribute the traffic between the source and destination over multiple paths
  - Only 3.3% of ASes have the same path
  - Only 3.7% have the same path length
- Majority of path segments through a network domain are not the shortest paths
  - Some paths are much longer than the observed shortest path in many ASes

Router Level Topologies of Autonomous Systems

- **Path traces from all**
  - October 15-20, 2016 (189M trace through 38,566 AS)
  - March 1-5, 2017 (195M trace through 39,101 AS)
- **IP to AS mapping**
  - BGPstream
  - Sister ASes
- **IP alias resolution**
  - routers have multiple interfaces
    - the path traces contain different IP addresses of a router
  - analytical and probe methods
- **Total of 19,614 ASes (in both)**
  - ignored ASes with less than 10 nodes
- **data available at https://im.cse.unr.edu/data/**

Network Size

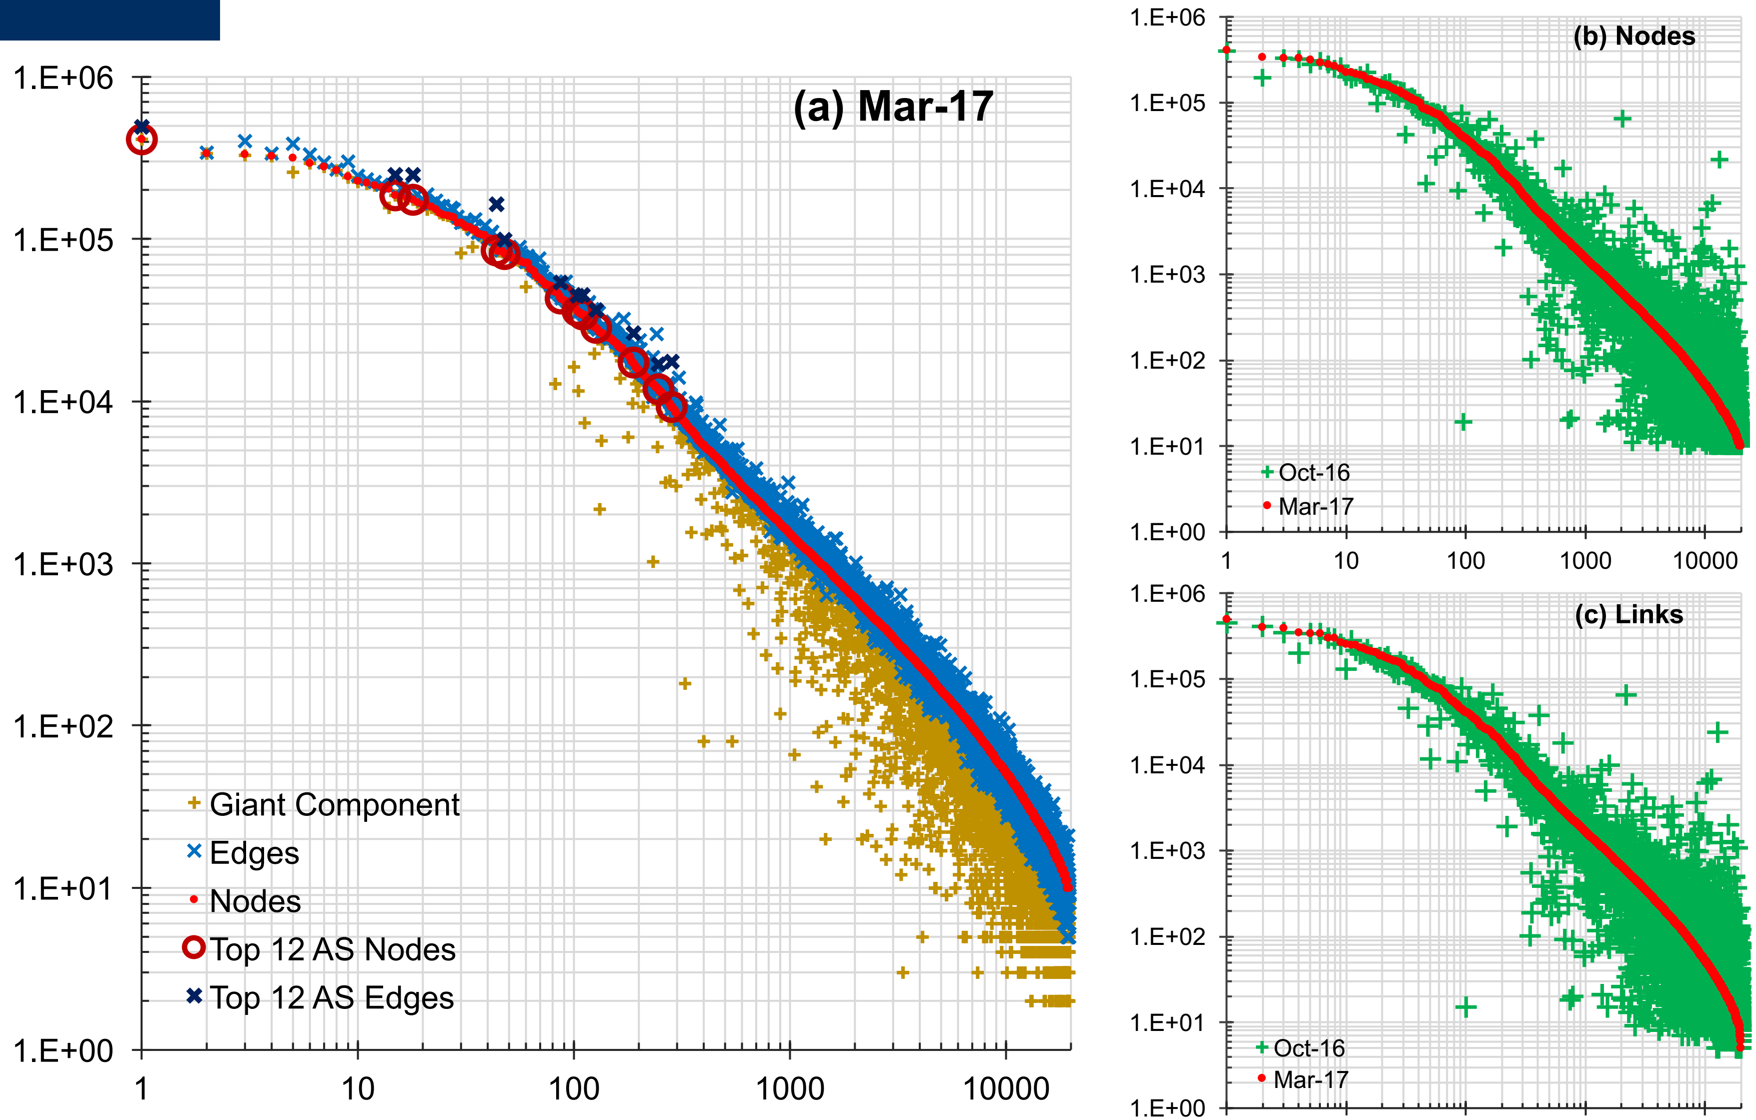

Node Degrees

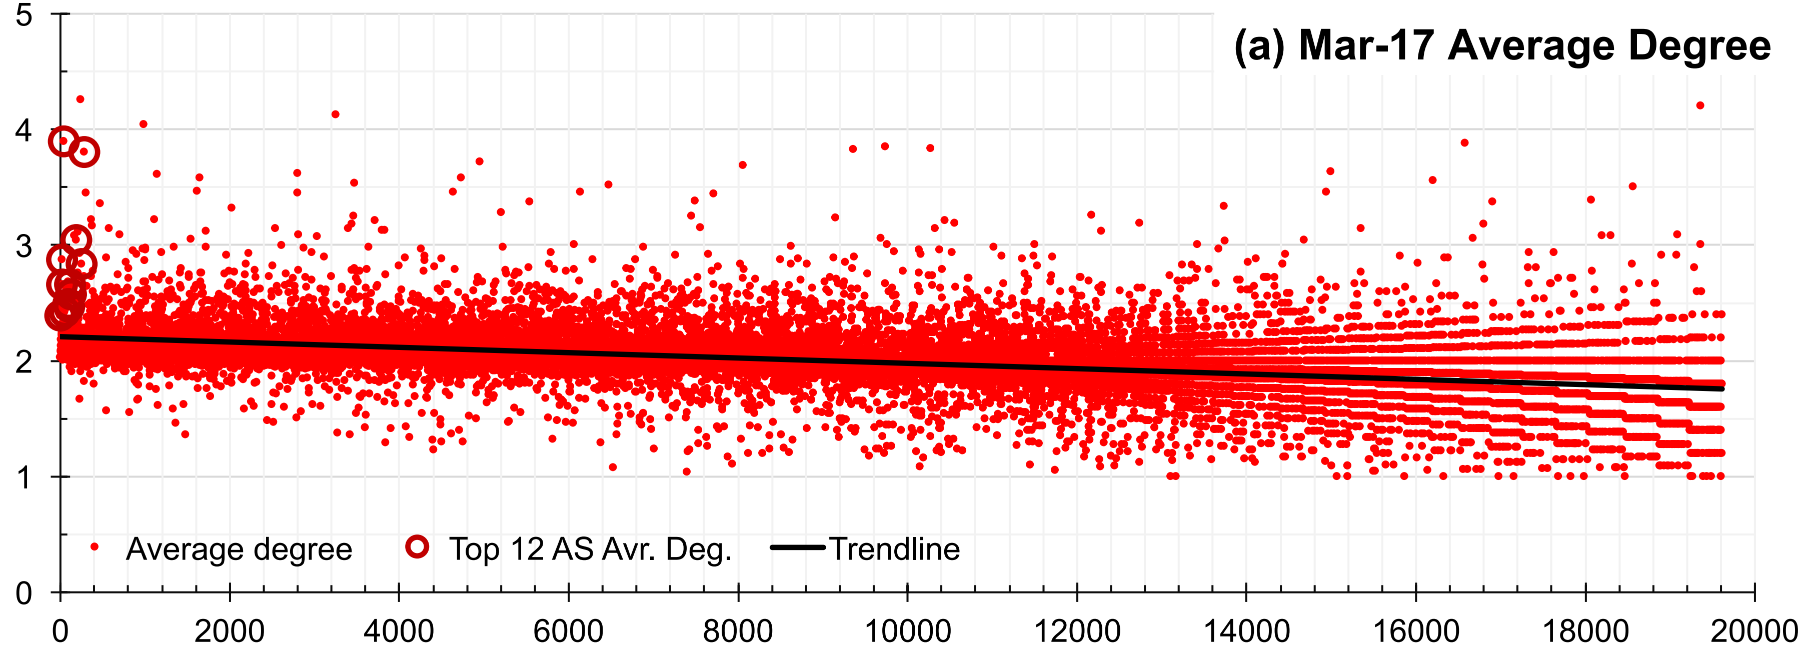

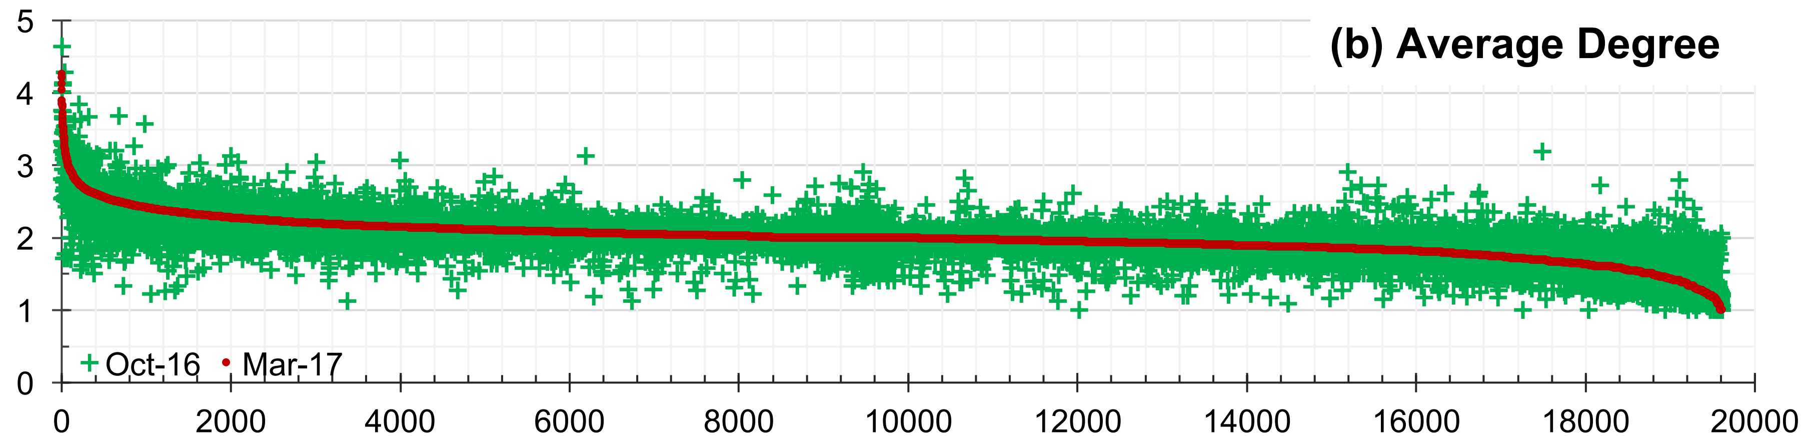

Maximum Degrees

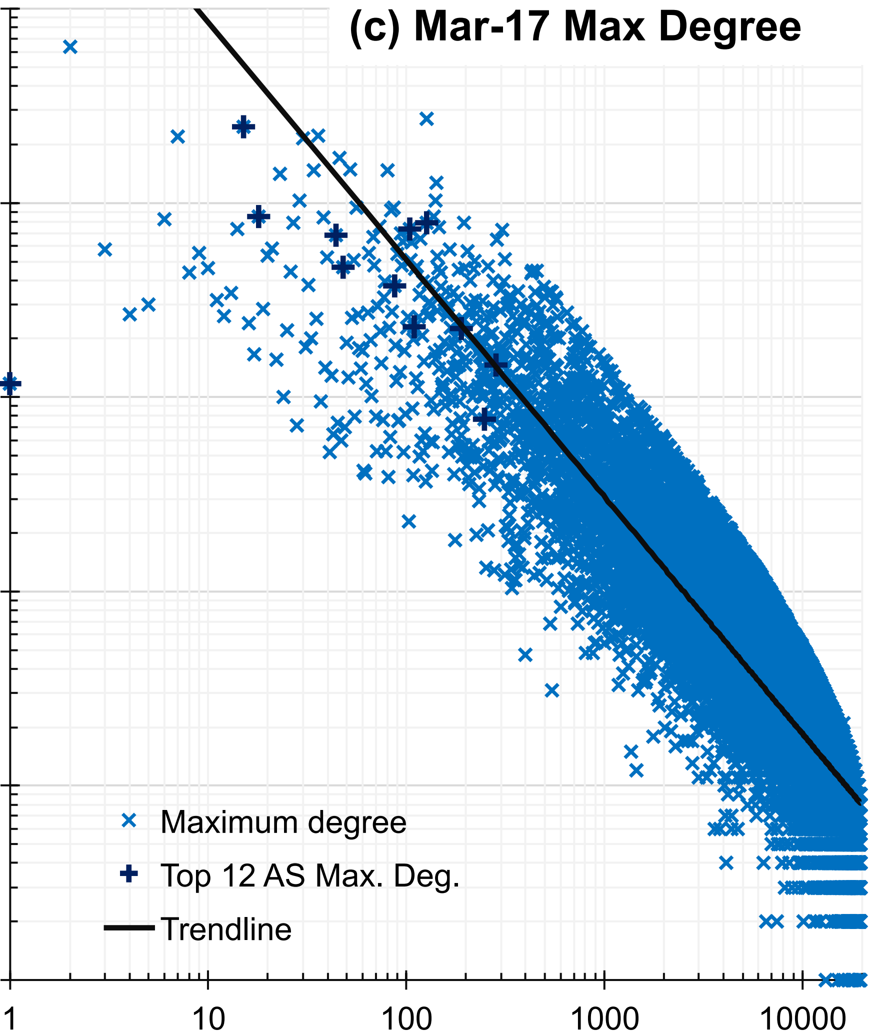

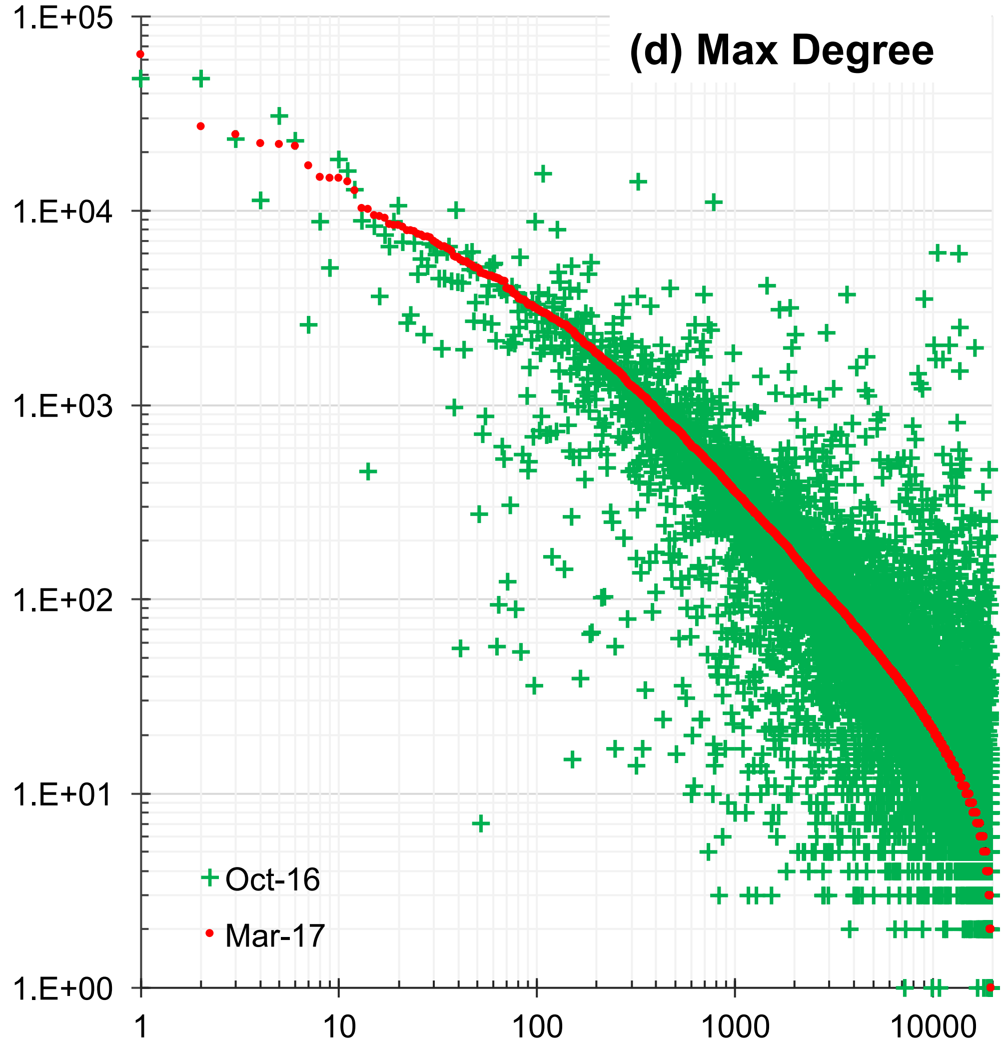

Assortativity

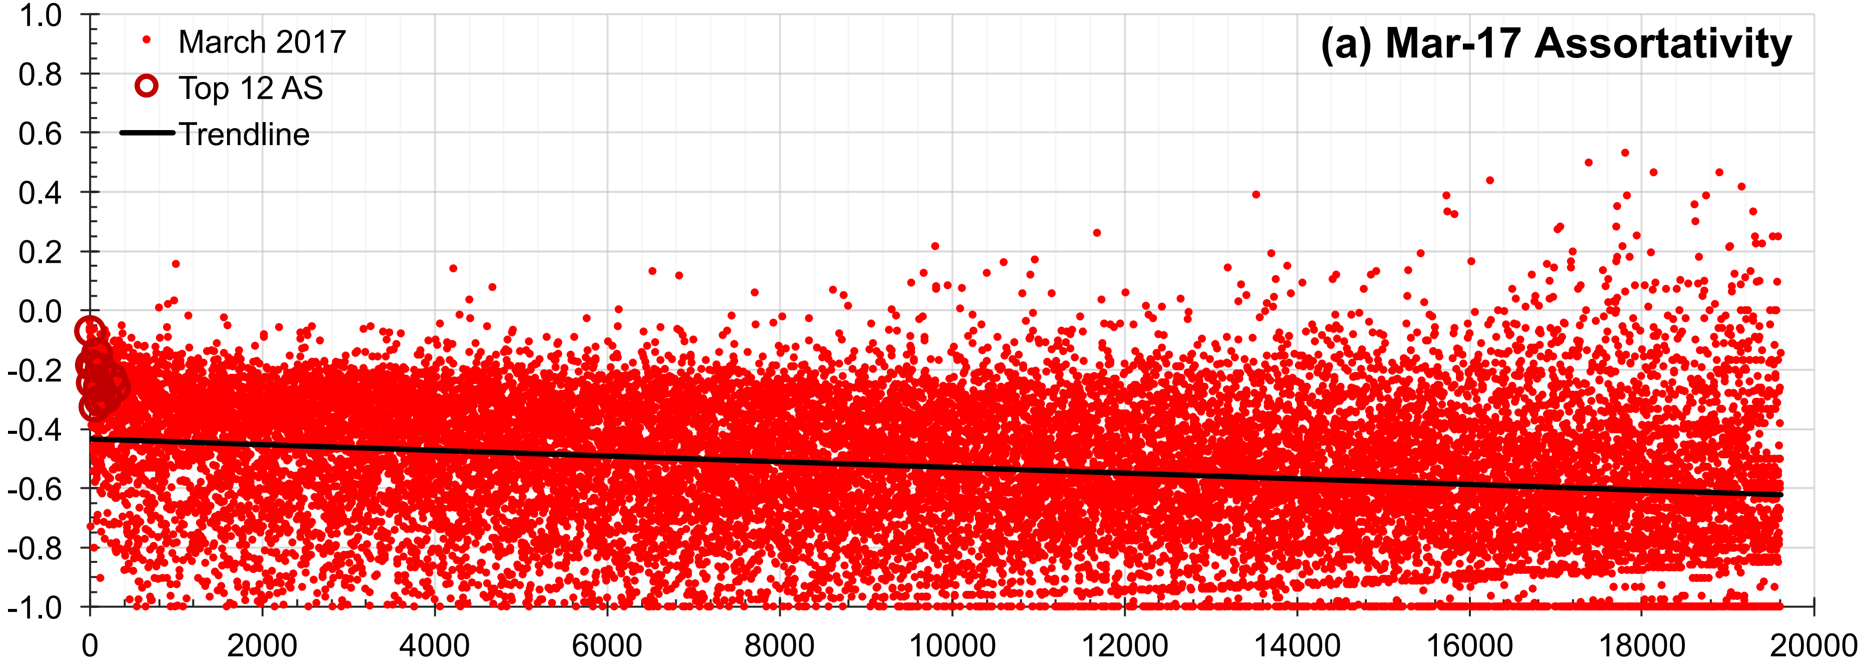

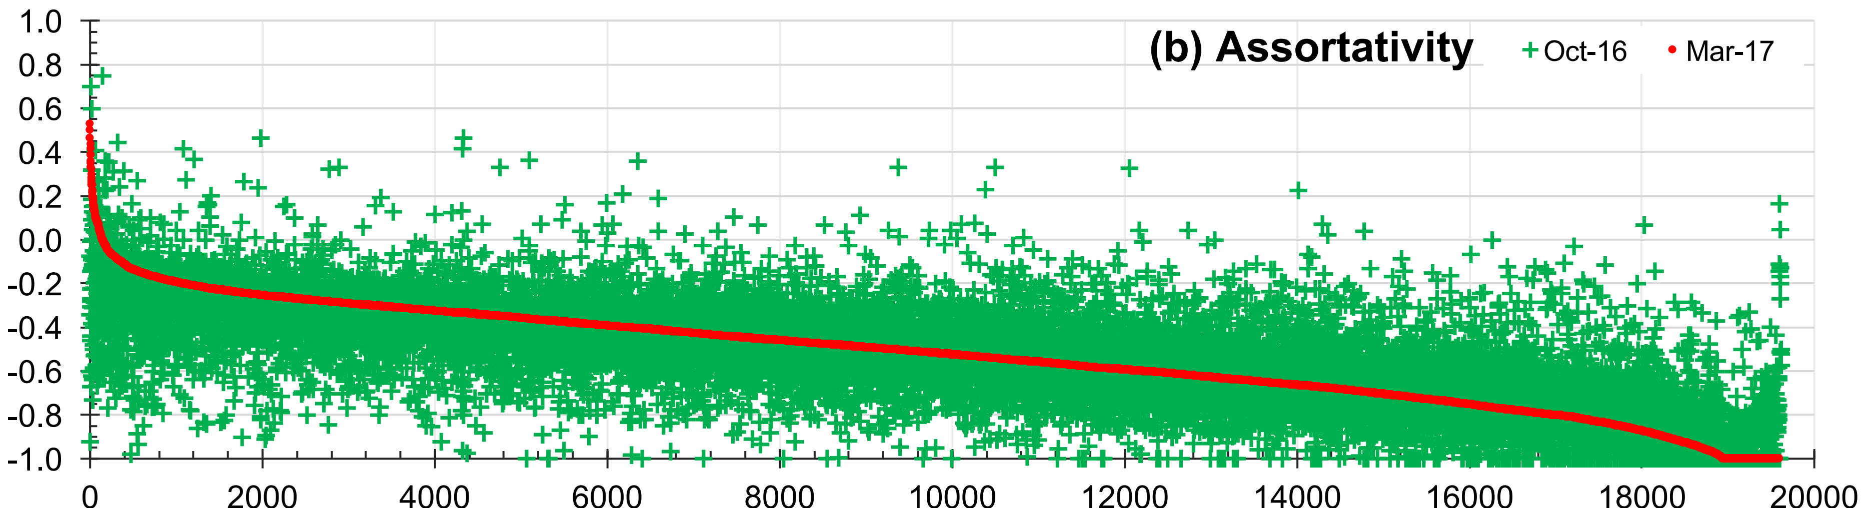

Clustering Coefficient

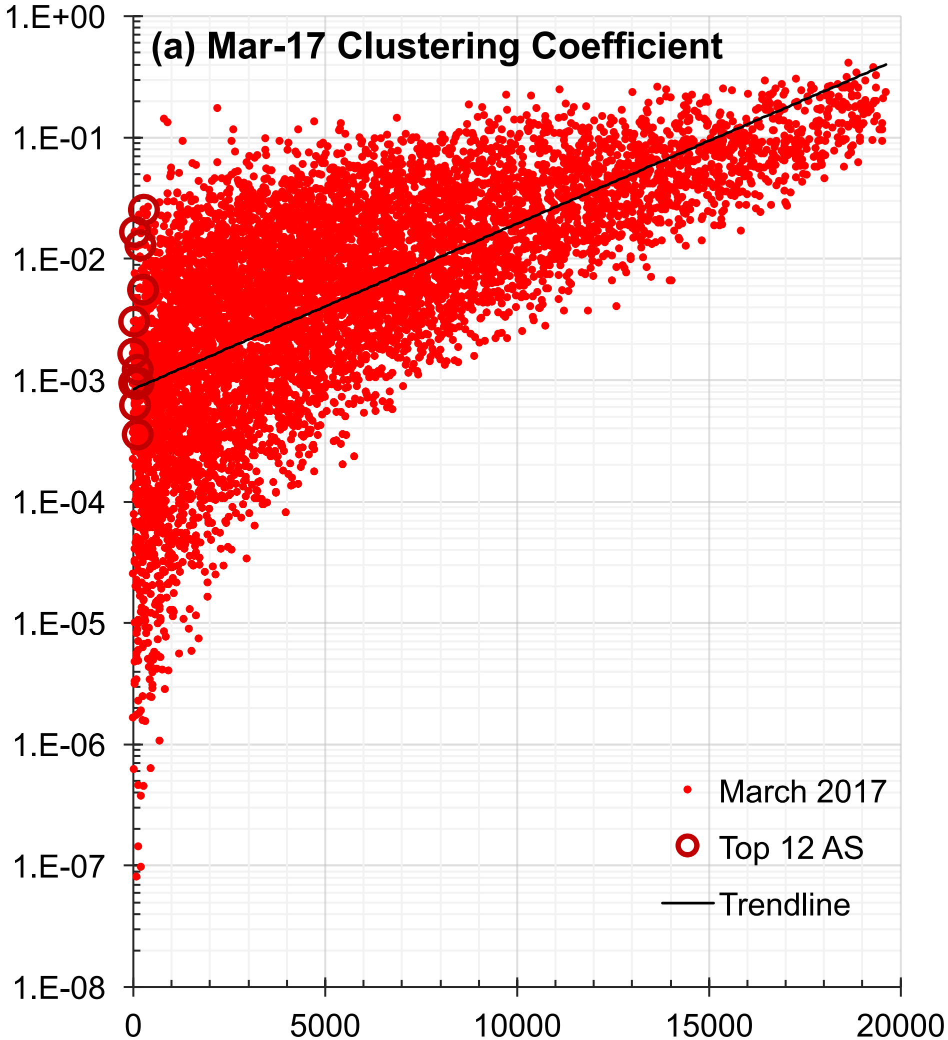

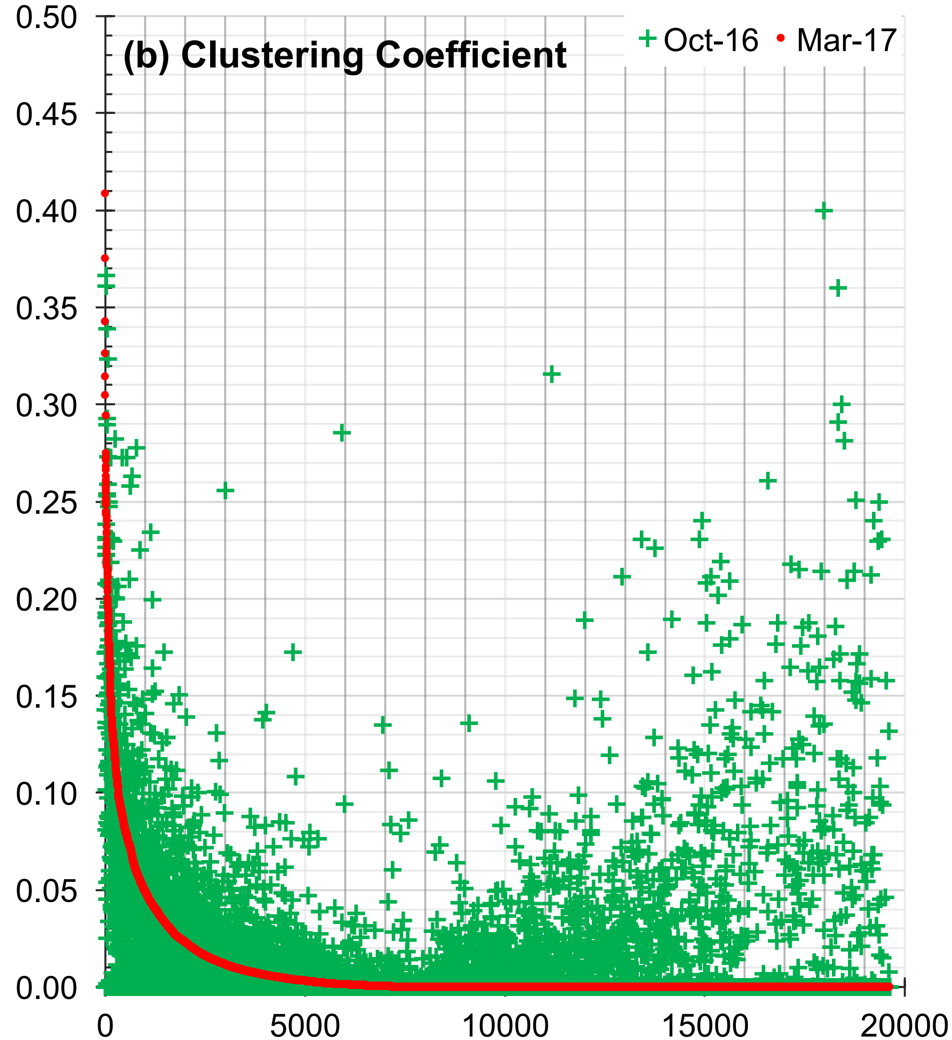

Network Characteristics of the top ranked ASes

AS number	Nodes	Links	Giant Component %	Degree: Average	Degree: Maximum	Density	Assortativity	Clustering: Node Average	Clustering: Global	Average Path Length	Diameter	Number of Communities	Modularity
3356	84,812	164,992	98.87	3.89	6,820	4.59E-05	-0.243	0.0602	0.0167	3.78	17	427	0.592
1299	186,588	248,112	99.85	2.66	24,698	1.43E-05	-0.068	0.0402	0.0010	3.64	12	139	0.800
2914	81,446	98,683	99.46	2.42	4,685	2.98E-05	-0.227	0.0168	0.0006	4.52	16	192	0.876
174	174,517	250,469	99.71	2.87	8,549	1.64E-05	-0.194	0.0548	0.0030	4.89	14	280	0.757
6453	34,206	45,501	99.40	2.66	2,303	7.78E-05	-0.181	0.0212	0.0012	4.45	15	81	0.802
6762	17,287	26,293	98.58	3.04	2,250	1.76E-04	-0.143	0.0290	0.0130	3.86	9	128	0.719
6939	36,219	44,670	99.67	2.47	7,329	6.81E-05	-0.323	0.0319	0.0004	4.07	15	82	0.813
2828	43,146	53,725	99.10	2.49	3,756	5.77E-05	-0.187	0.0220	0.0009	5.00	13	161	0.865
3491	9,234	17,551	99.26	3.80	1,462	4.12E-04	-0.257	0.0553	0.0252	3.56	12	70	0.613
701	410,701	489,928	99.89	2.39	1,175	5.81E-06	-0.303	0.0362	0.0016	4.37	16	640	0.876
1239	11,938	16,928	99.05	2.84	772	2.38E-04	-0.292	0.0638	0.0056	4.89	14	84	0.795
1273	28,365	36,524	98.99	2.58	7,939	9.08E-05	-0.253	0.0457	0.0010	3.87	17	102	0.796
mean	**93,205**	**124,448**	**99.32**	**2.84**	**5,978**	**1.03E-04**	**-0.223**	**0.0398**	**0.0059**	**4.24**	**14**	**199**	**0.775**
median	**39,683**	**49,613**	**99.33**	**2.66**	**4,221**	**6.29E-05**	**-0.235**	**0.0382**	**0.0014**	**4.22**	**14**	**134**	**0.798**

K-core layouts

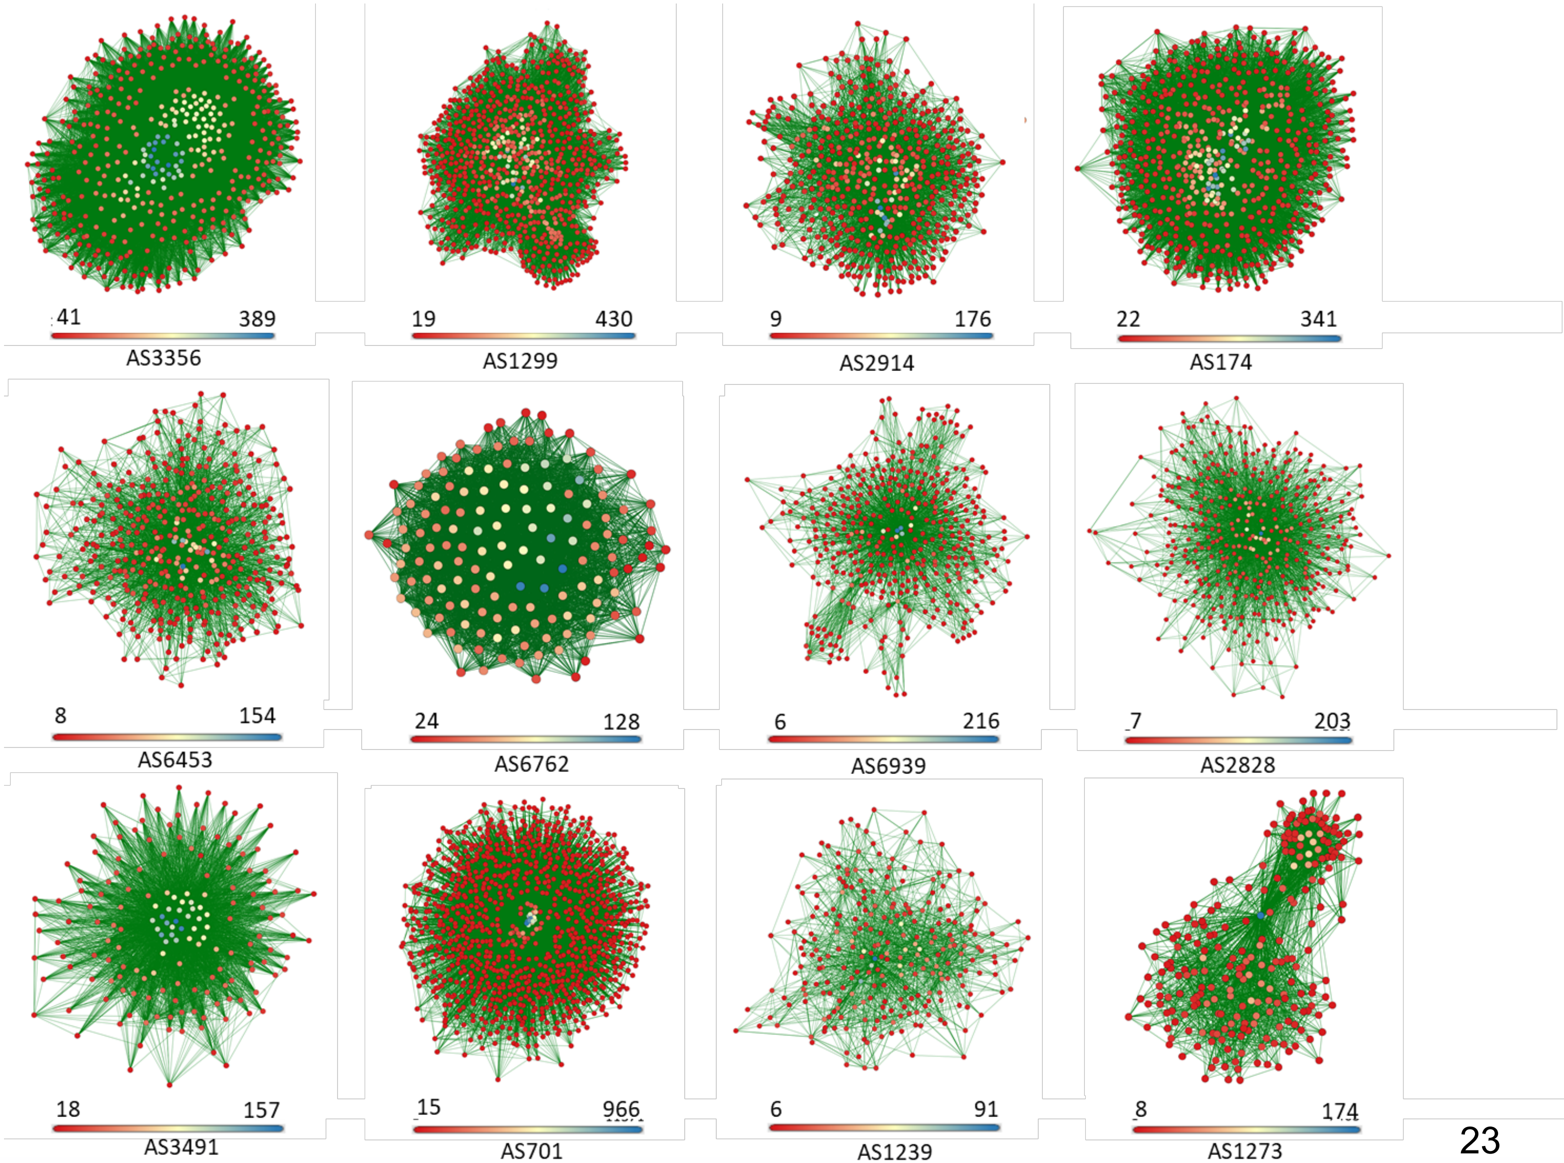

Degree Distribution

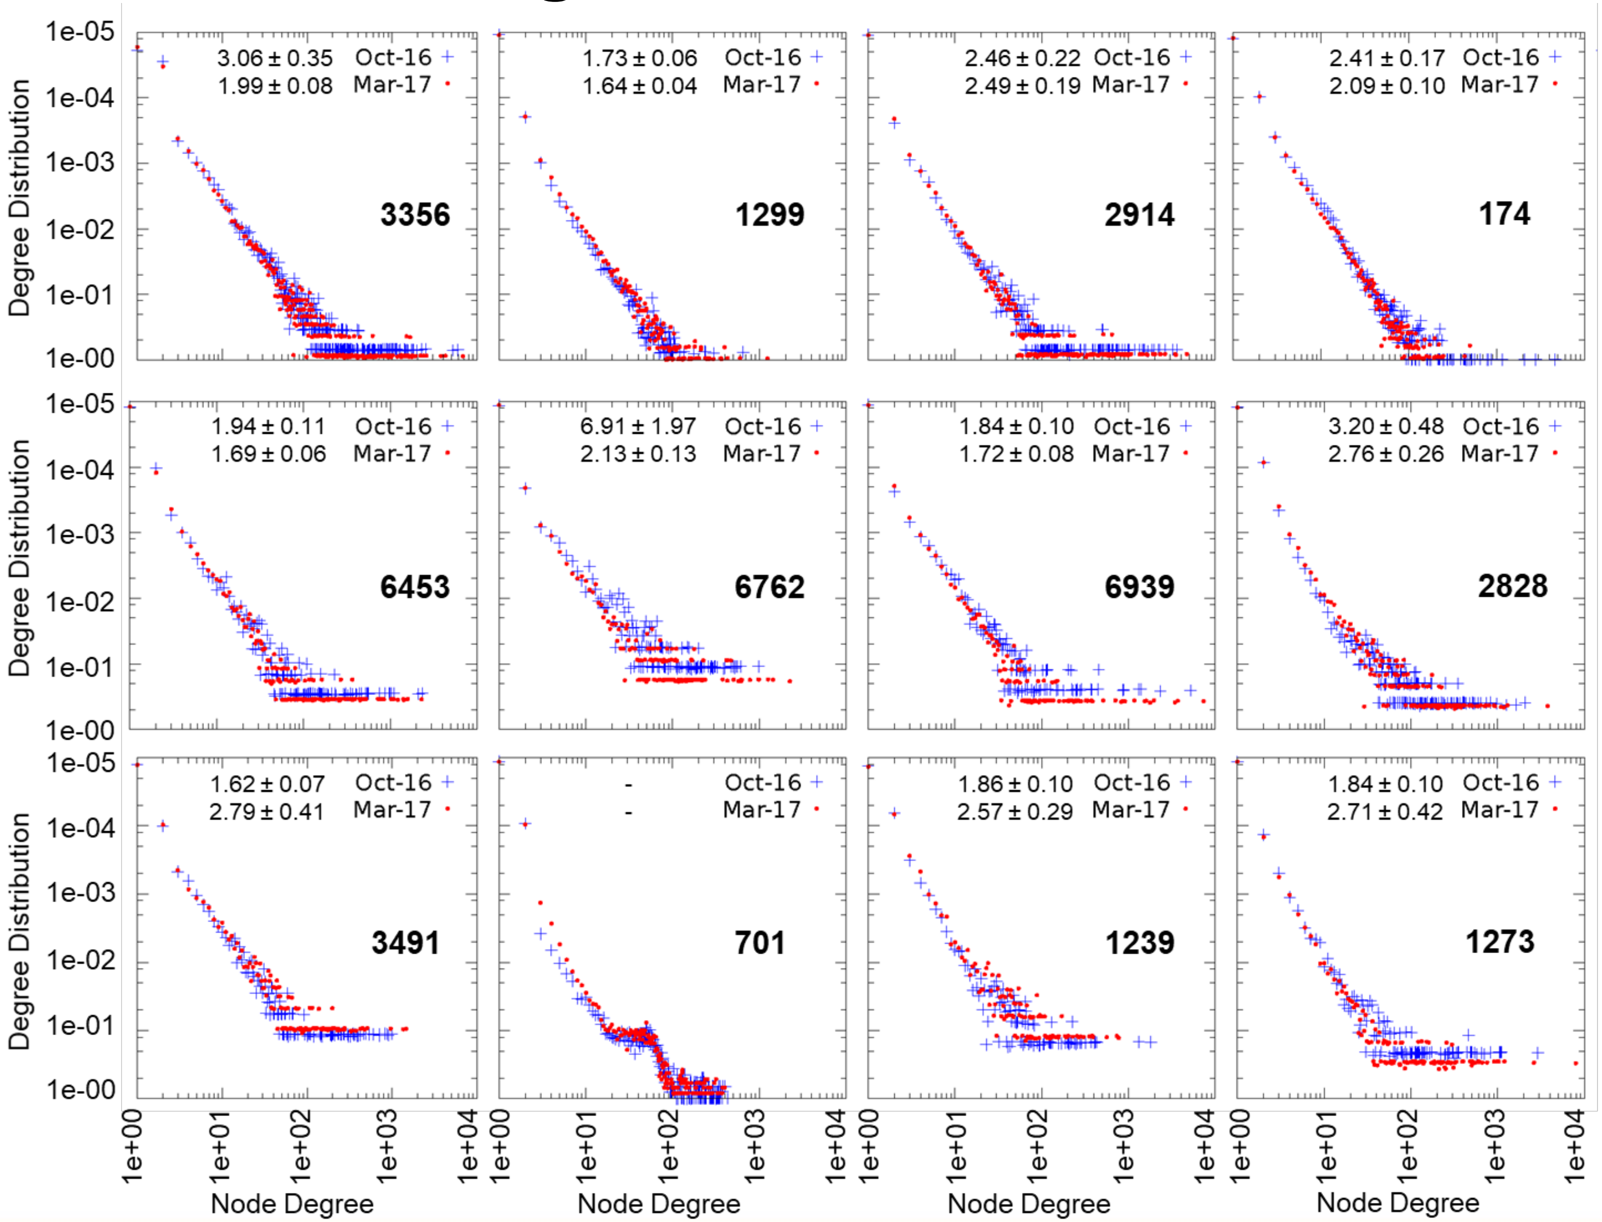

Clustering Distribution

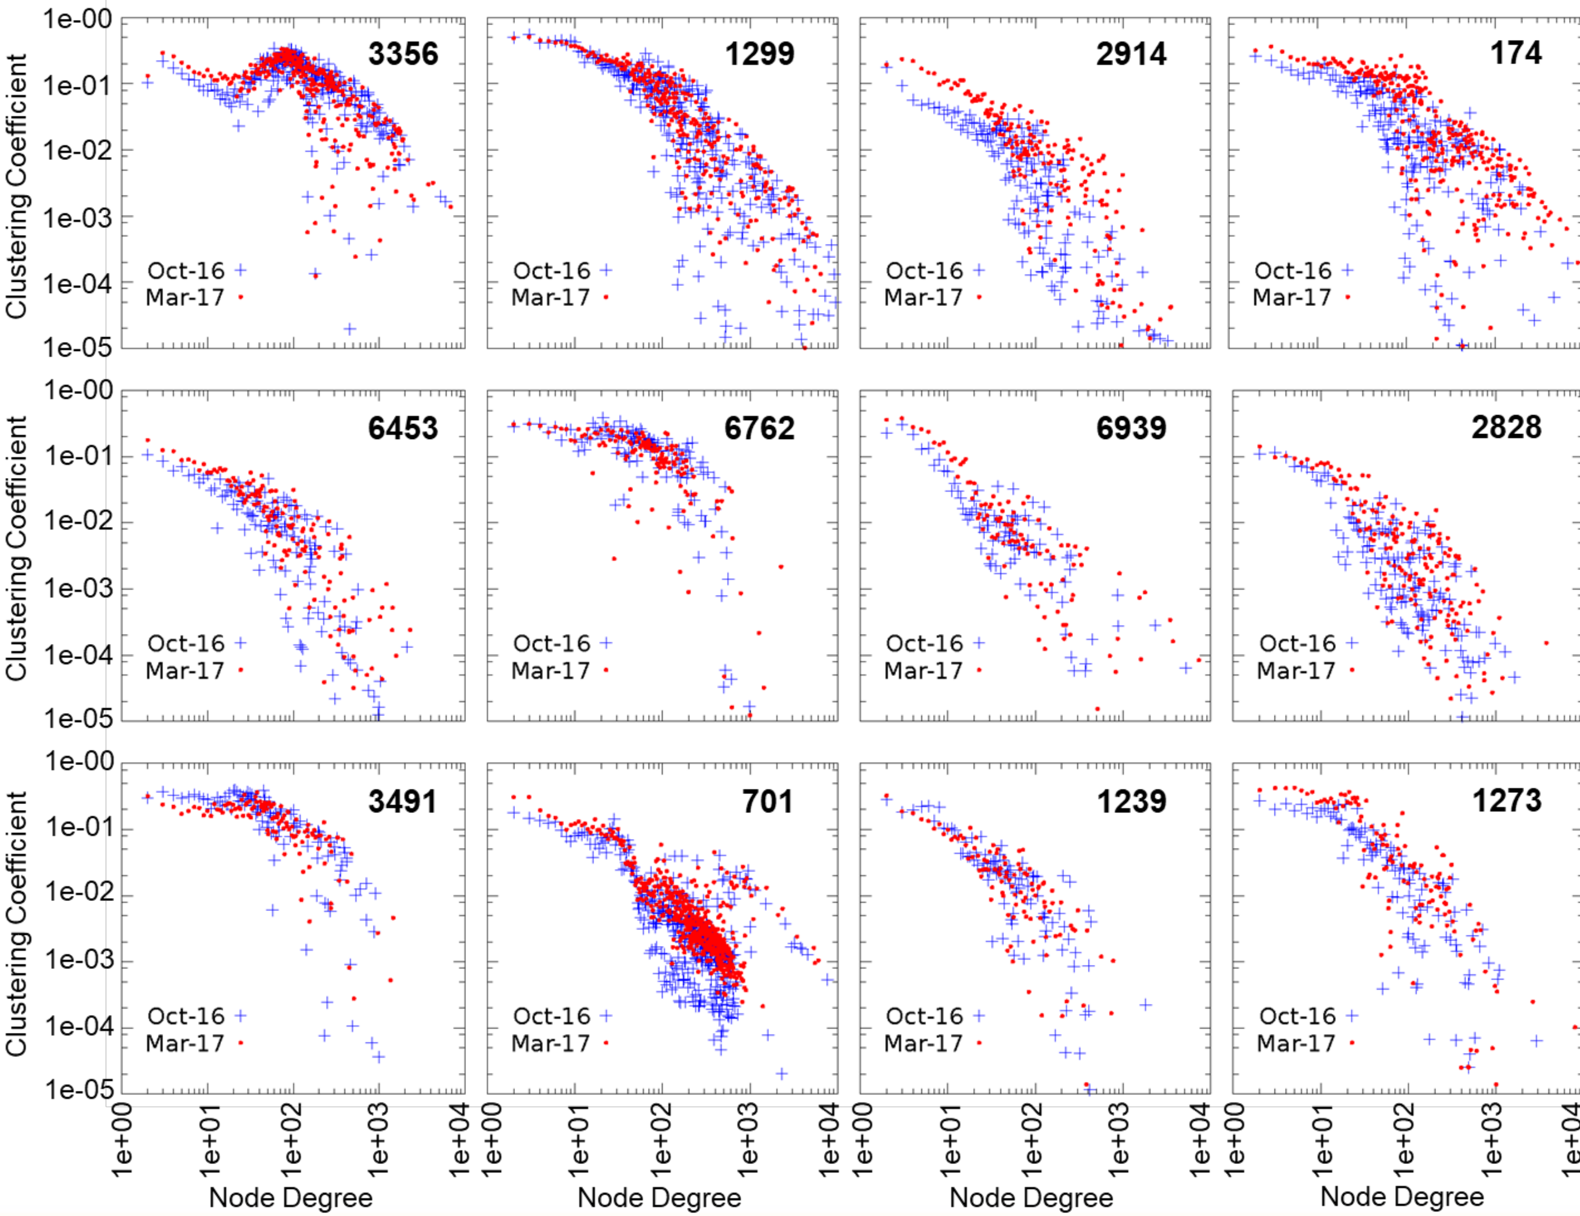

Summary

- Mapped and analyzed the backbone AS topologies
  - ***19,614 Ases*** (majority transit AS)
- Majority of analyzed ASes are disassortative
  - only few are non-assortative
  - star-like topologies
    - high degree hubs connect low degree nodes
  - assortativity of graph is independent of its size
- Majority of the *top ranked ASes* have similar graph structures
  - A *well-connected core* and *hierarchical* or *mesh based peripheries*
  - Most have power-law degree distributions
- All of the top ranked ASes are small worlds networks
  - High clustering and low average path length

Internet Measurement (IM) Platform

- Measurement network of single board computers (Odroids)
- Vantage points for future Internet measurements

Looking for volunteers to expand across the globe

im.cse.unr.edu

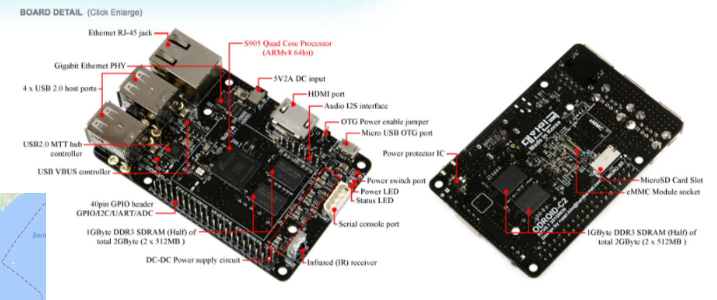

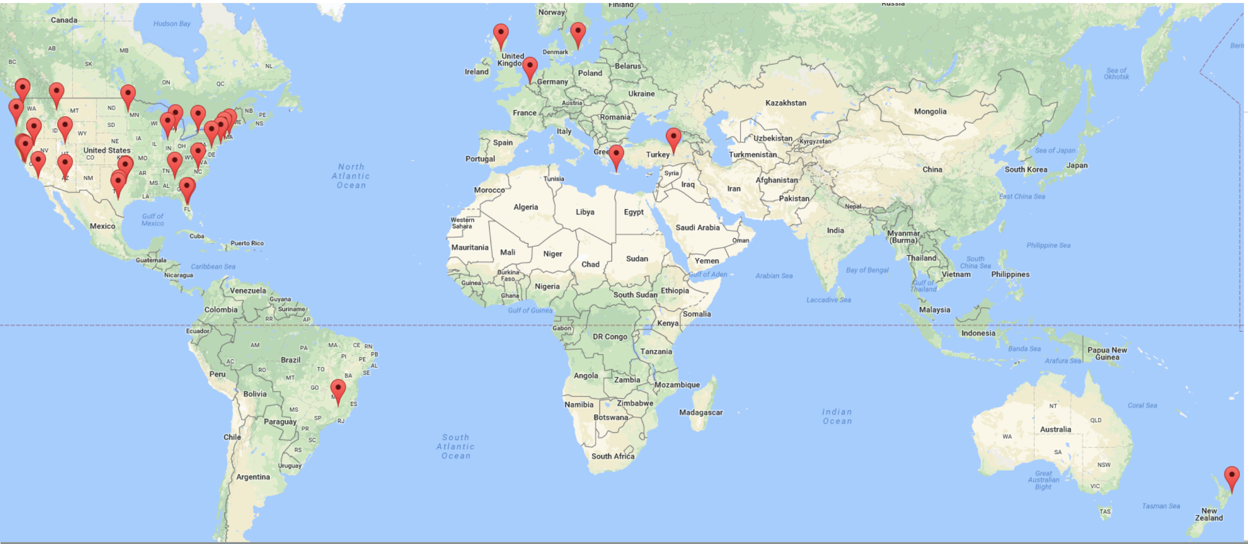

Analytical Subnet and IP Alias Resolution (ASIAR)

- **Analytical tool for Subnet & Alias Resolution**
  - perform link and router inference using the common IP assignment practices

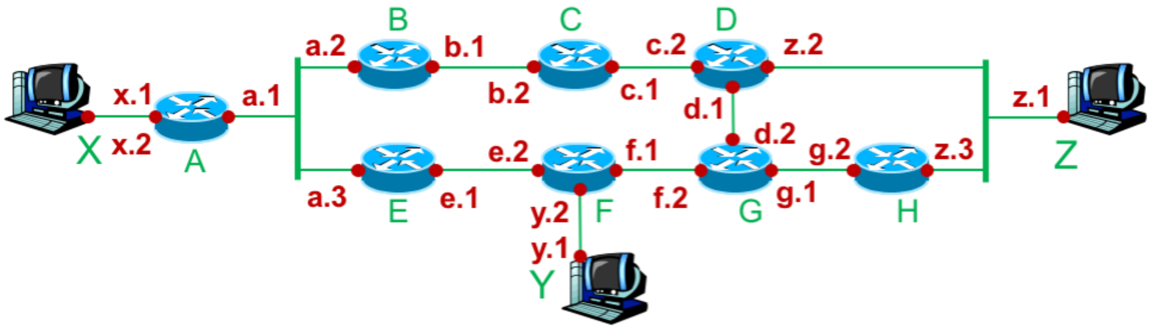

- **Subnetwork resolution**
  - Distance condition
    - Subnets are determined based on IP distances observed by the ingresses and egresses
  - Completeness condition
    - A minimum ratio of potential subnet IPs should be observed
  - Reference point completeness condition *(new)*
    - IPs should be observed from a minimum number of reference points, i.e., ingresses and egresses

Analytical Subnet and IP Alias Resolution (ASIAR)

- **IP alias resolution**
  - No-loop condition *(optional)*
    - Alias IPs should not appear in the same trace, except consecutive
  - Neighbor condition *(updated)*
    - Identifies IP aliases between 2 subnet pairs without the need of storing IP triplets
  - Distance condition *(updated)*
    - A minimum number of reference points should have matching distances

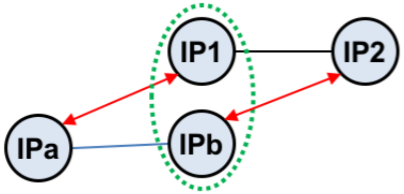

- **Parameter optimization**
  - Genetic Algorithms was used to optimize the IP completeness, reference point and distance thresholds on synthetic graphs
    - Accuracy decreases with weighted graphs where observed paths are not shortest paths
    - Introduced parameters to tolerate errors on distance match

Analytical Subnet and IP Alias Resolution (ASIAR)

- **Non-weighted**

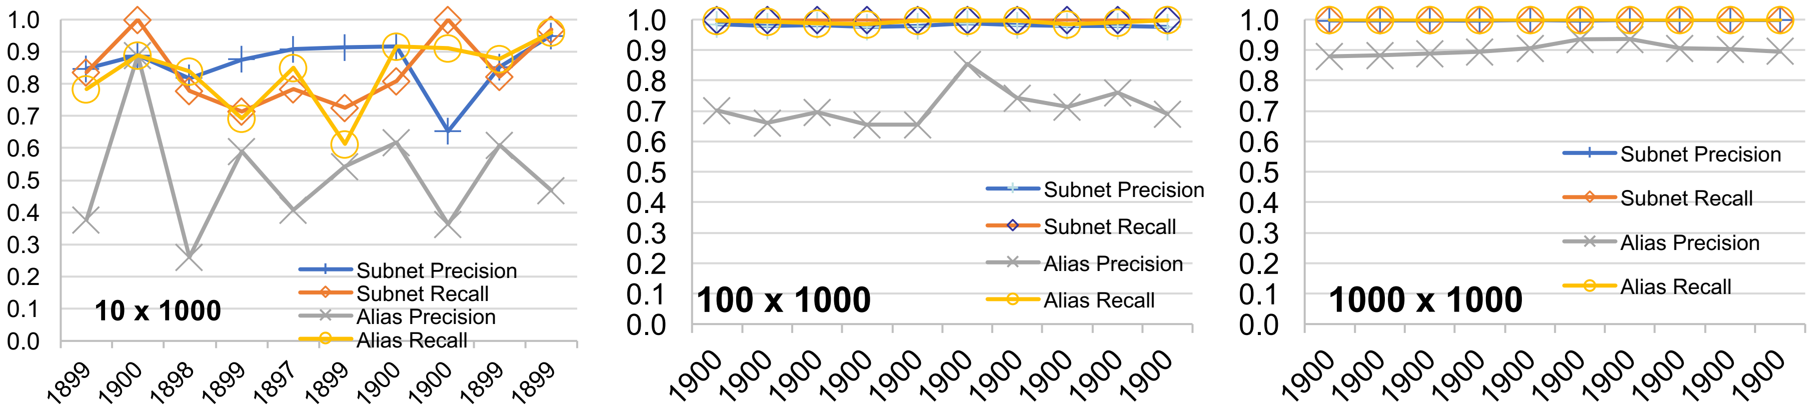

- **Weighted**

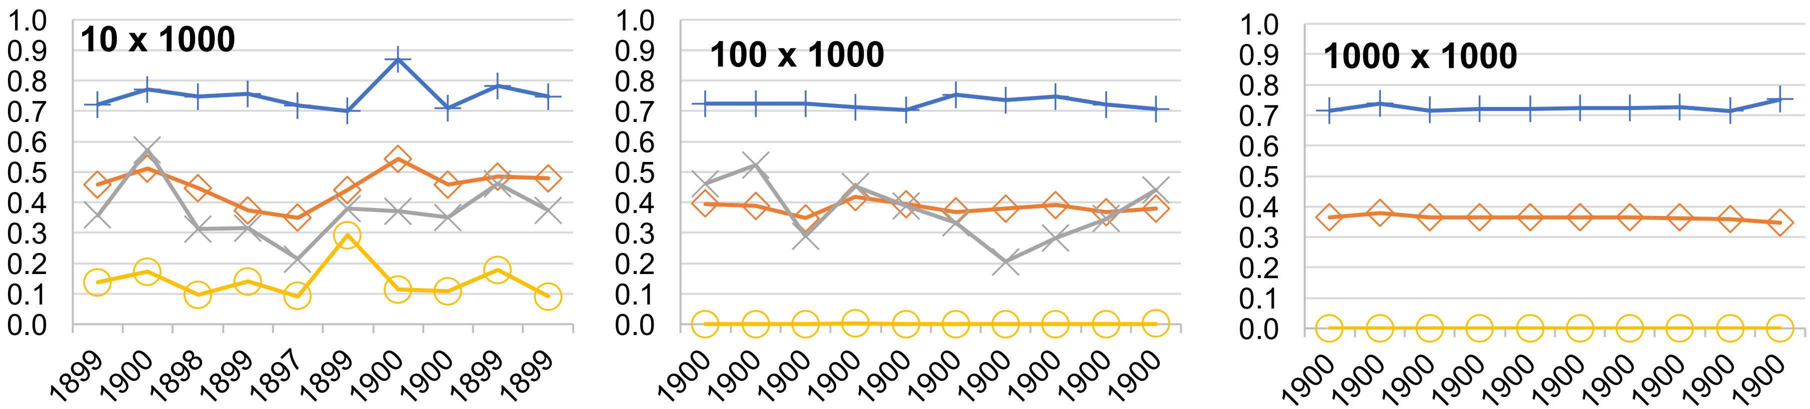

- **Code available at https://im.cse.unr.edu/**

Exhaustive Mapping of an Autonomous System

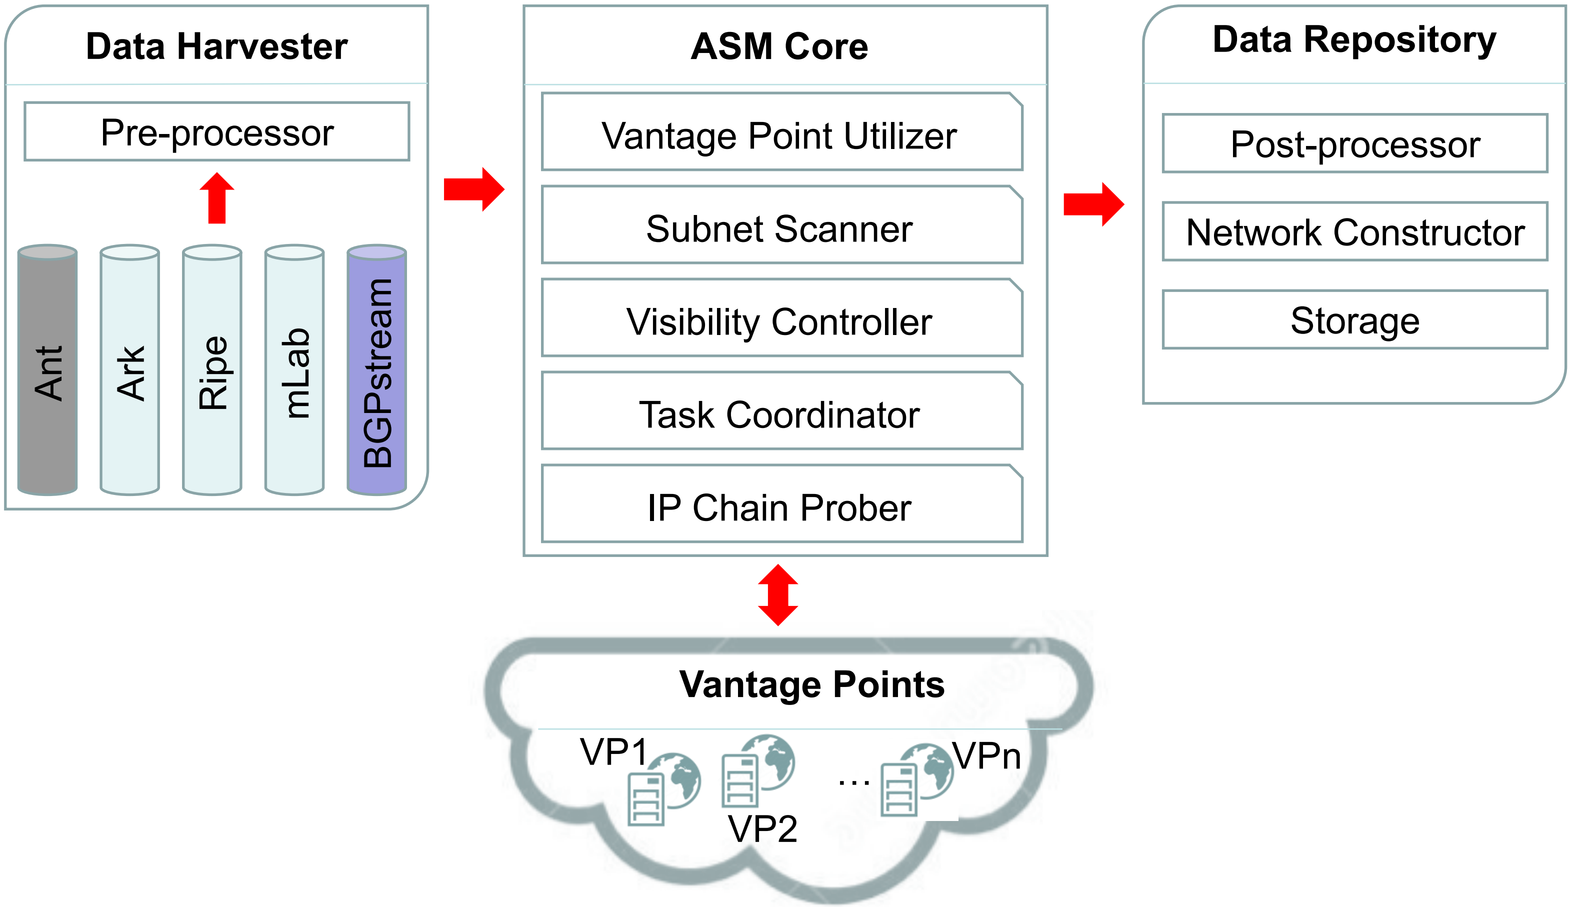

Topology discovery ranking of ingresses

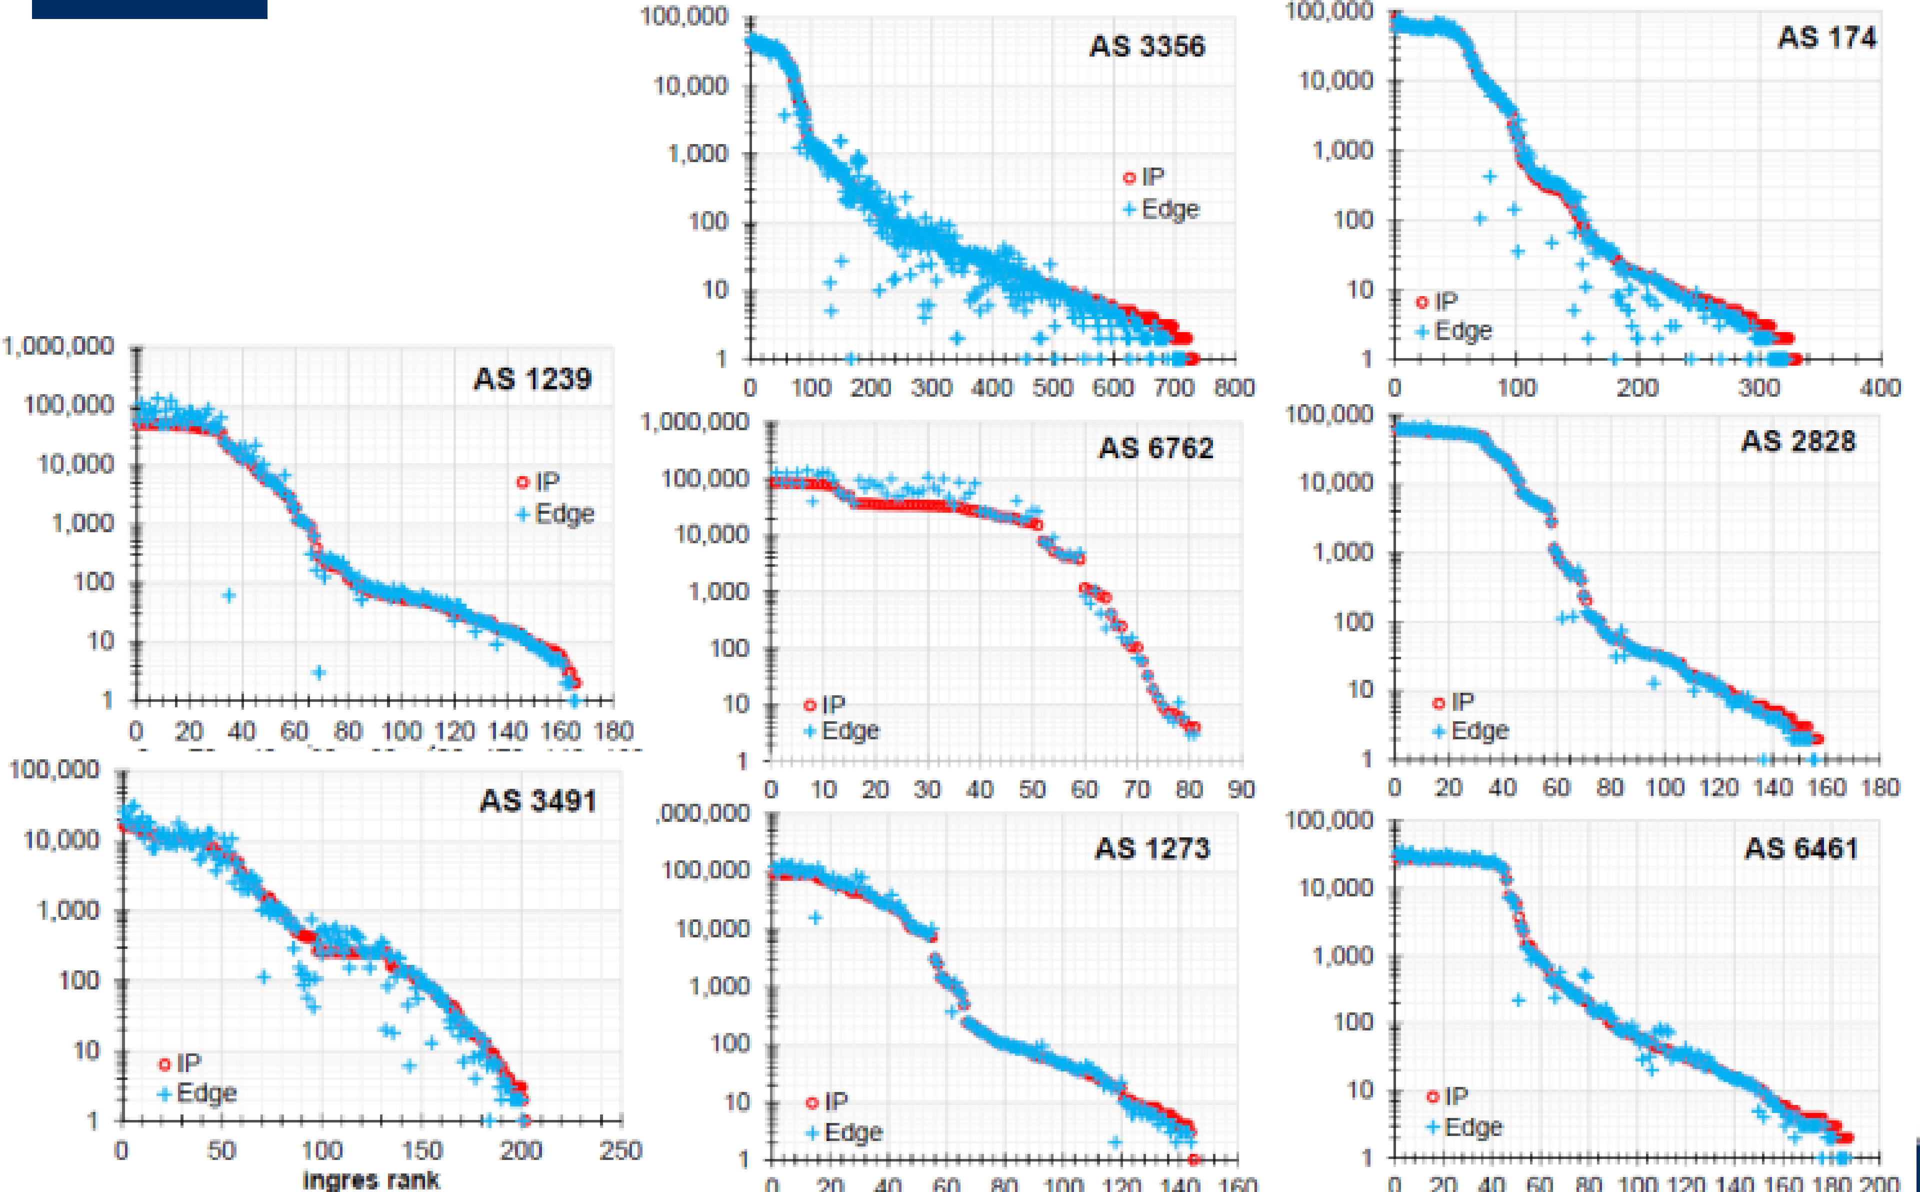

Identifying Ingresses

Iteration 1

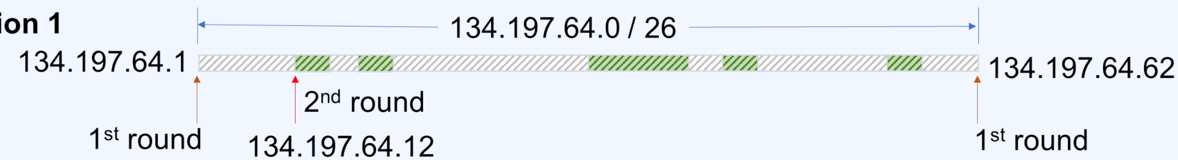

Round #	Target	Ingress	Hop D.
1	134.197.64.1	No Ingress	N/A
2	134.197.64.12	224.13.168.11	8
1	134.197.64.62	173.181.1.15	6

There is no match

Iteration 2

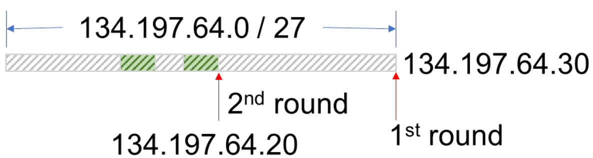

Round #	Target	Ingress	Hop D.
0	134.197.64.1	No Ingress	N/A
0	134.197.64.12	224.13.168.11	8
1	134.197.64.30	No Ingress	N/A
2	134.197.64.20	224.13.168.11	8

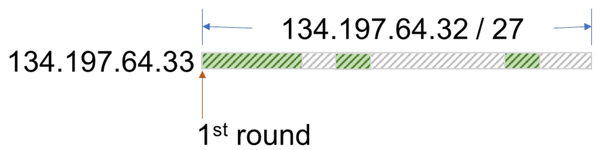

Round #	Target	Ingress	Hop D.
0	134.197.64.62	173.181.1.15	6
1	134.197.64.33	173.181.1.15	7

Results

Target: 134.197.64.0 / 26

Detected Subnet: 134.197.64.0 / 27 INGRESS: 224.13.168.11 HOP: 8

Detected Subnet: 134.197.64.32 / 27 INGRESS: 173.181.1.15 HOP: 6

Ingress to Subnet Reachability of VPs

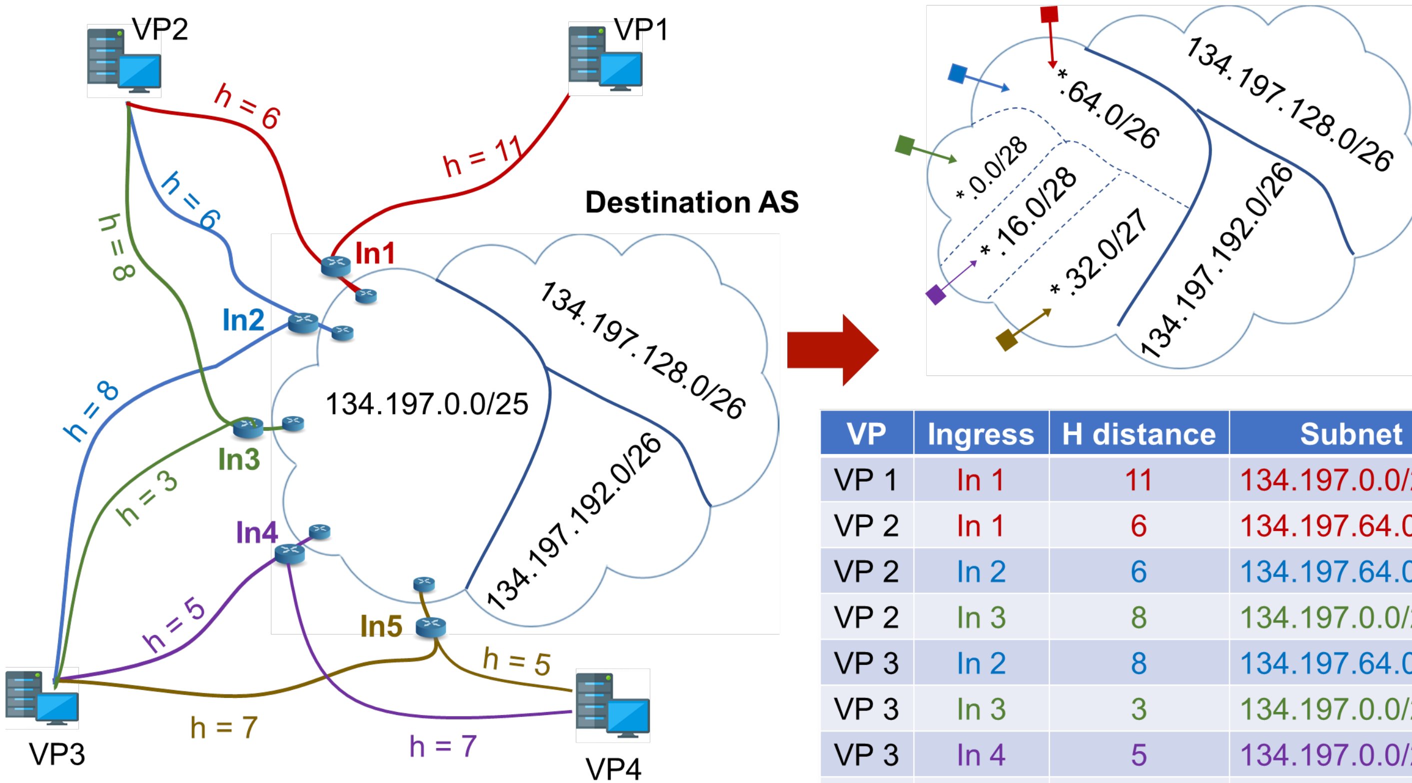

VP	Ingress	H distance	Subnet
VP 1	In 1	11	134.197.0.0/25
VP 2	In 1	6	134.197.64.0/26
VP 2	In 2	6	134.197.64.0/26
VP 2	In 3	8	134.197.0.0/26
VP 3	In 2	8	134.197.64.0/26
VP 3	In 3	3	134.197.0.0/28
VP 3	In 4	5	134.197.0.0/28
VP 3	In 5	7	134.197.32.0/27
VP 4	In 4	7	134.197.0.0/25
VP 4	In 5	5	134.197.32.0/27

Probes Generated

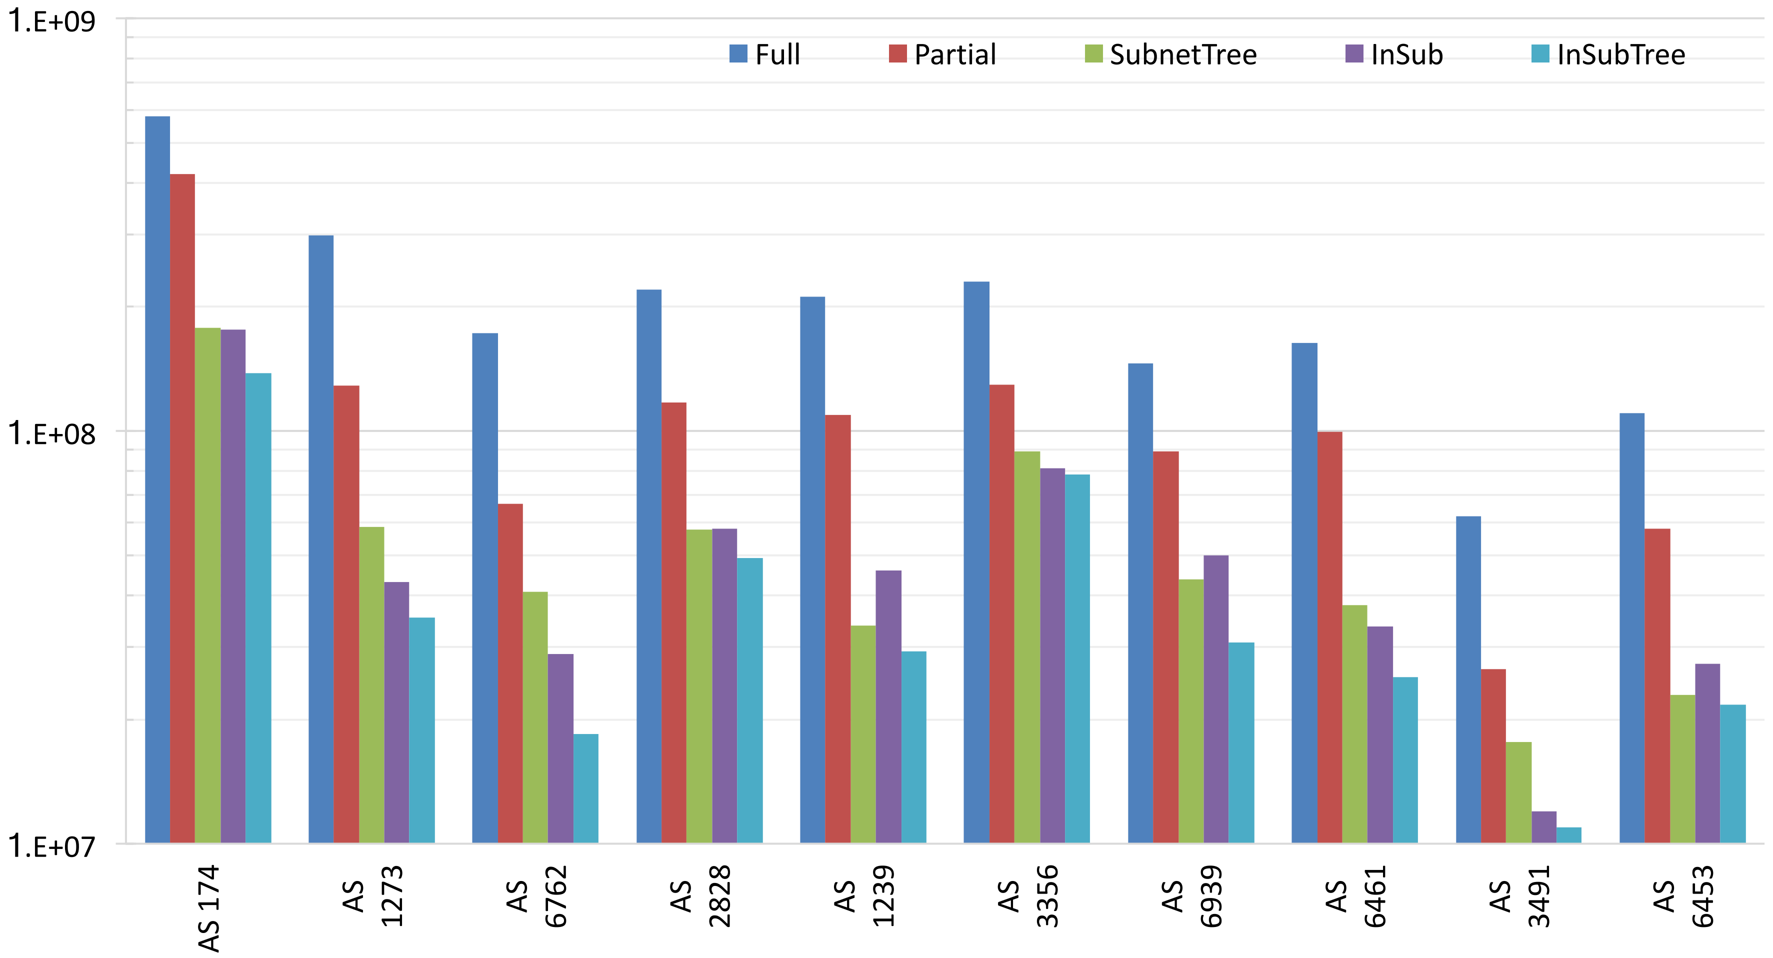

Measured IP addresses

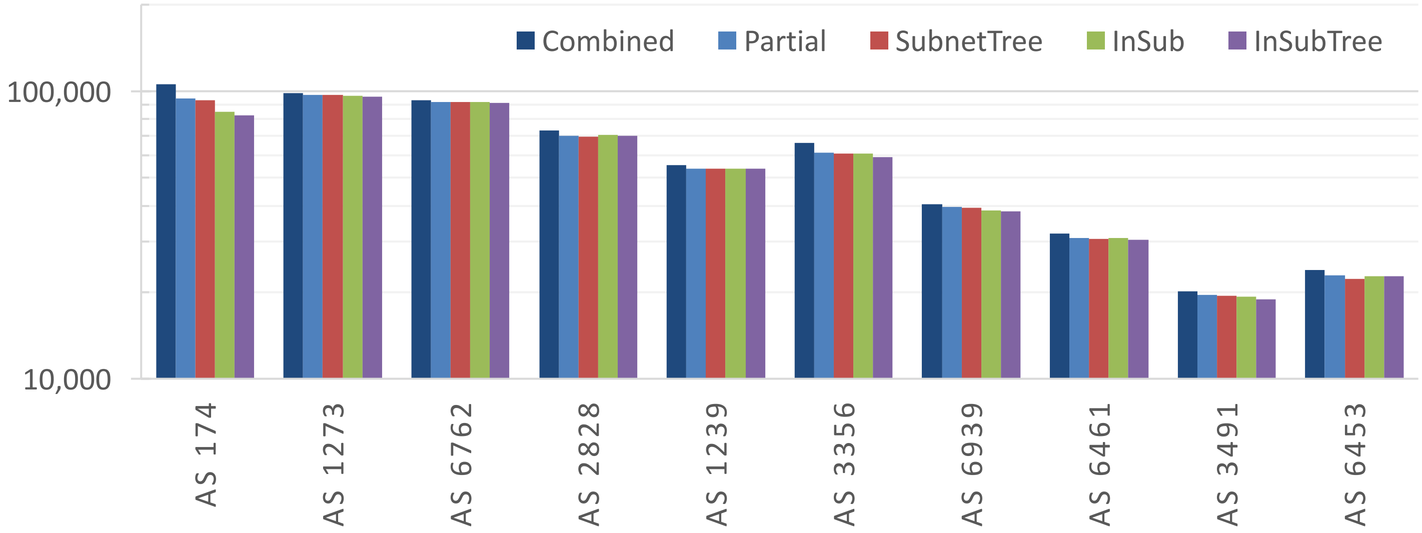

Observed IPs per AS (for each method)

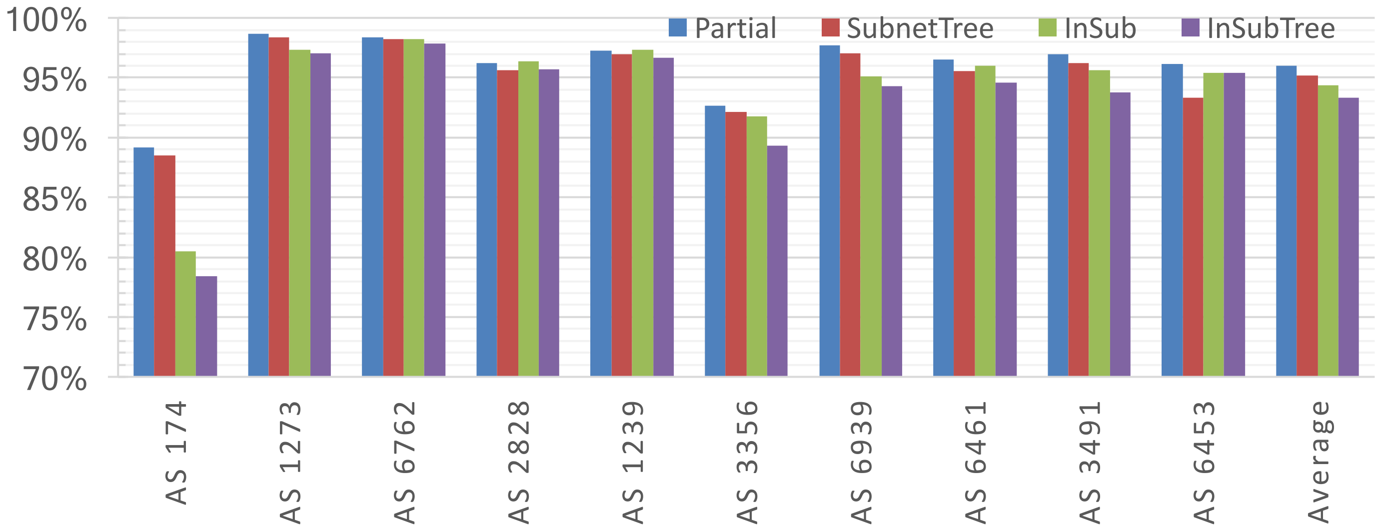

Ratio of IPs discovered by each method

Measured Links

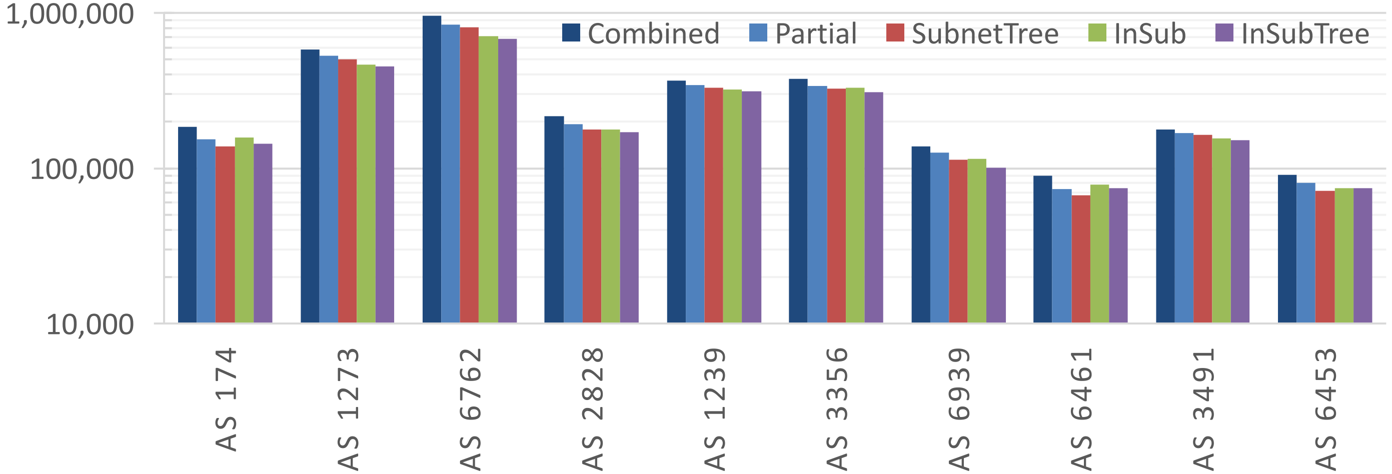

Observed links per AS (for each method)

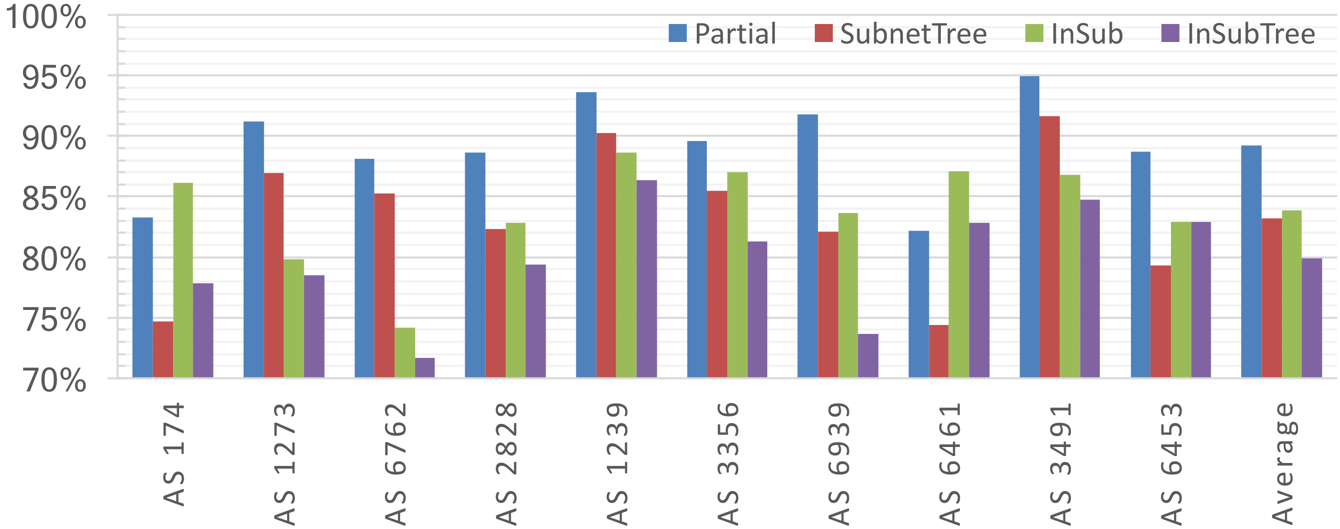

Ratio of links discovered by each method

Exhaustive Mapping of an Autonomous System

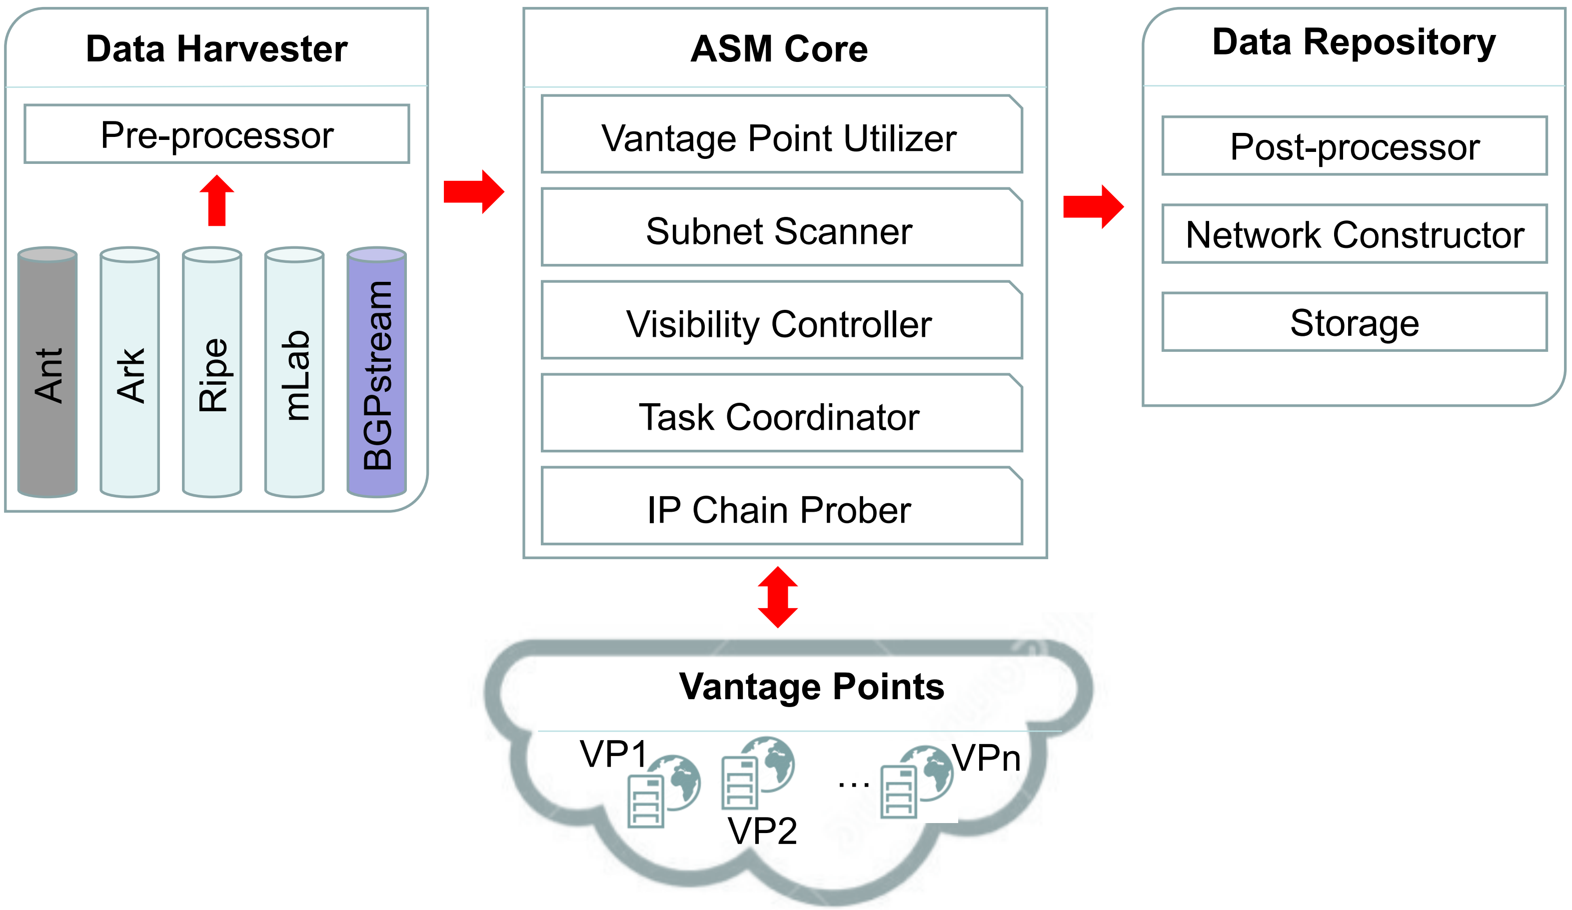

Vantage points per ingress

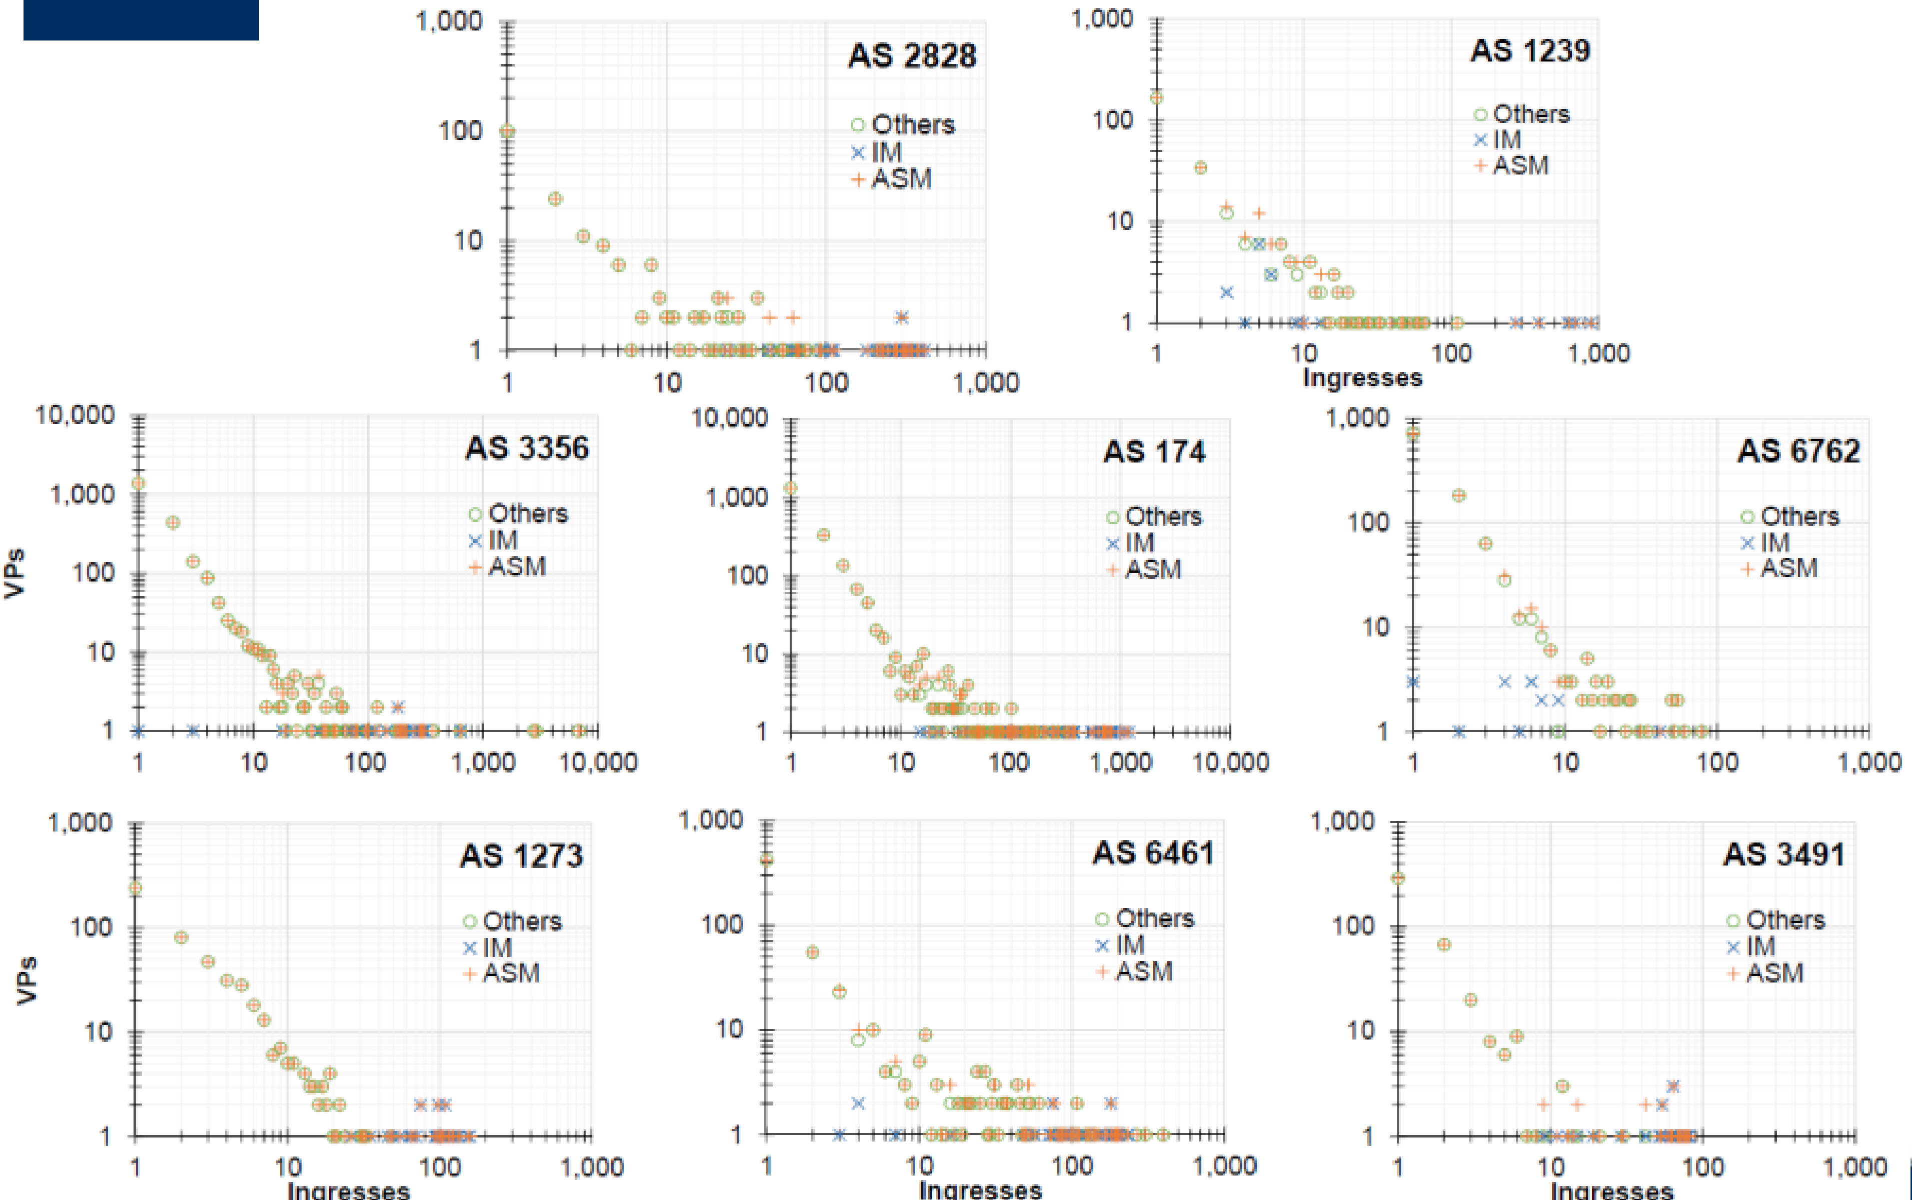

IP Discovery Per AS

	AS 1239	AS 1273	AS 174	AS 2828	AS 3356	AS 3491	AS 6461	AS 6762
Ant Census	61,645	113,262	650,910	225,742	276,942	74,040	102,338	109,384
Others	2,027	904	20,310	3,671	23,557	1,180	7,357	2,563
IM	48,032	72,522	351,857	189,178	66,801	54,844	91,797	94,974
ASM	49,419	73,152	367,025	191,722	88,447	55,439	97,815	96,567
other %	**4.10%**	**1.24%**	**5.53%**	**1.91%**	**26.63%**	**2.13%**	**7.52%**	**2.65%**
IM %	**97.19%**	**99.14%**	**95.87%**	**98.67%**	**75.53%**	**98.93%**	**93.85%**	**98.35%**
both %	1.30%	0.37%	1.40%	0.59%	2.16%	1.06%	1.37%	1.00%

Link Discovery Per AS

	AS 1239	AS 1273	AS 174	AS 2828	AS 3356	AS 3491	AS 6461	AS 6762
Ant Census	-	-	-	-	-	-	-	-
Others	3,528	1,356	48,664	6,402	65,931	3,059	15,322	10,405
IM	67,633	153,918	461,487	285,754	110,062	392,372	206,682	517,112
ASM	70,420	155,212	505,405	291,671	174,861	394,930	221,144	526,484
other %	5.01%	0.87%	9.63%	2.19%	37.70%	0.77%	6.93%	1.98%
IM %	96.04%	99.17%	91.31%	97.97%	62.94%	99.35%	93.46%	98.22%
both %	1.05%	0.04%	0.94%	0.17%	0.65%	0.13%	0.39%	0.20%

IP Discovery per Ingress

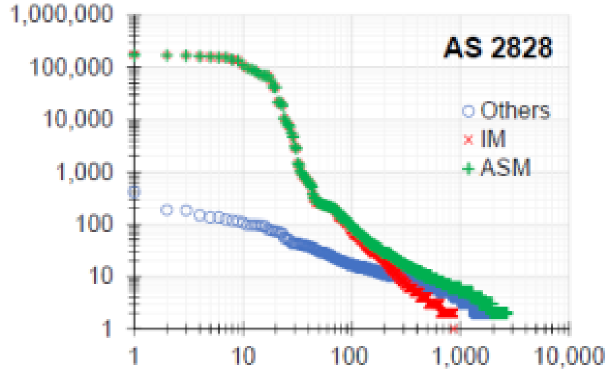

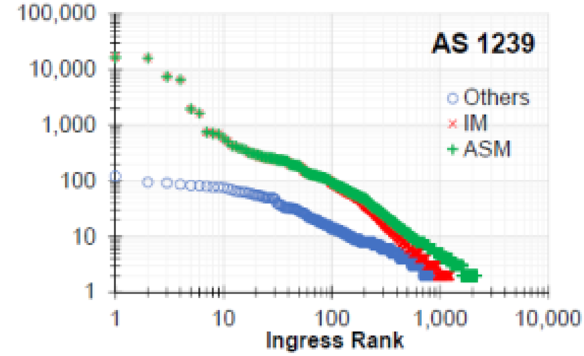

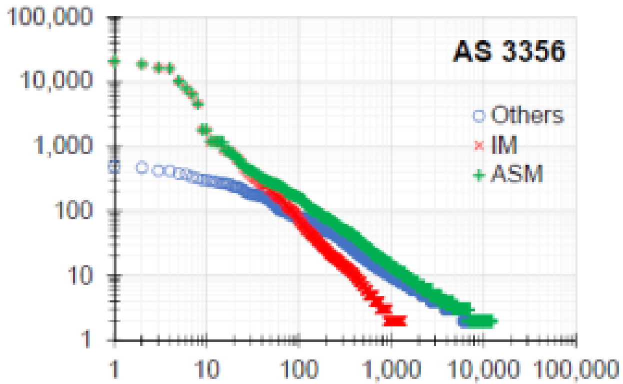

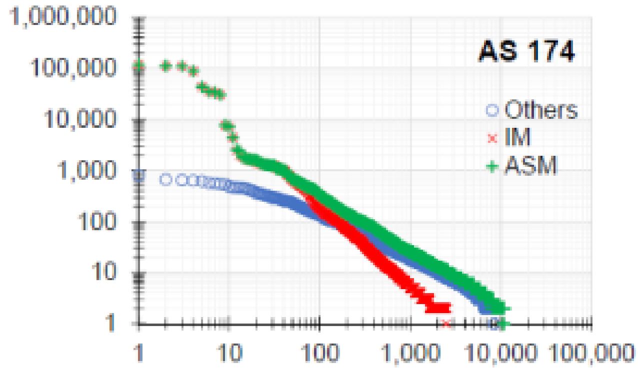

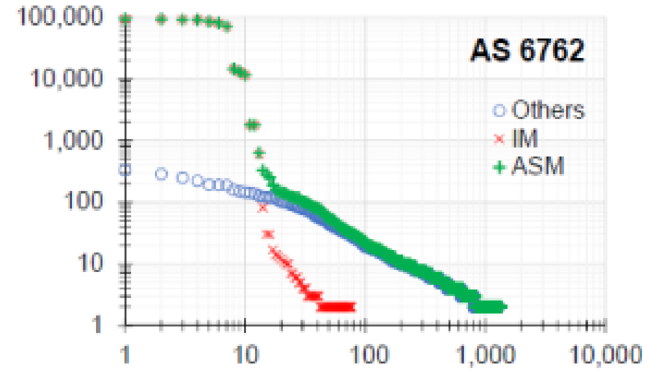

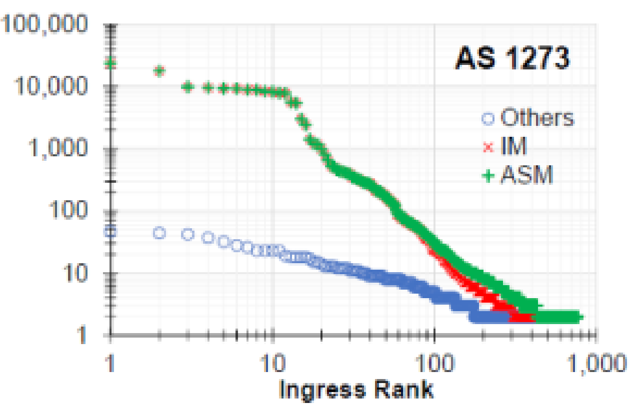

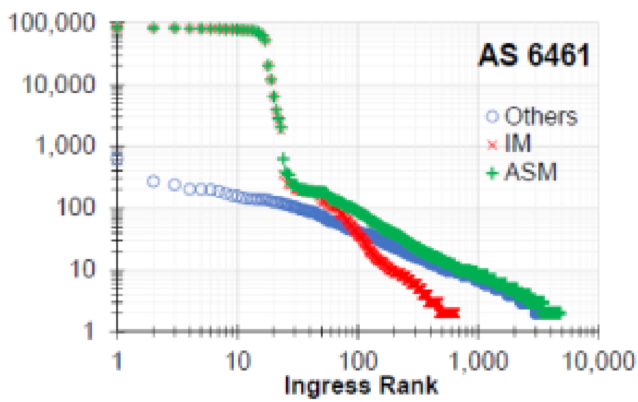

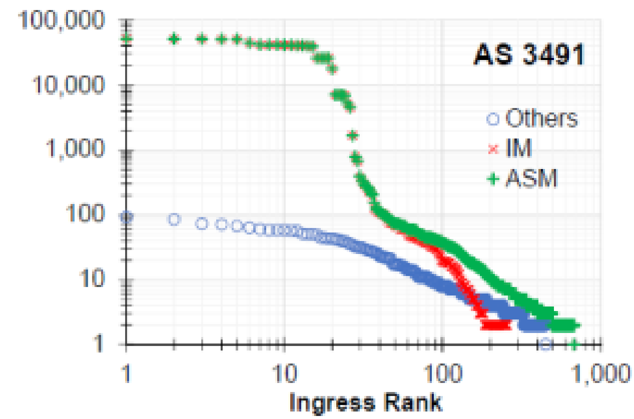

Link Discovery per Ingress

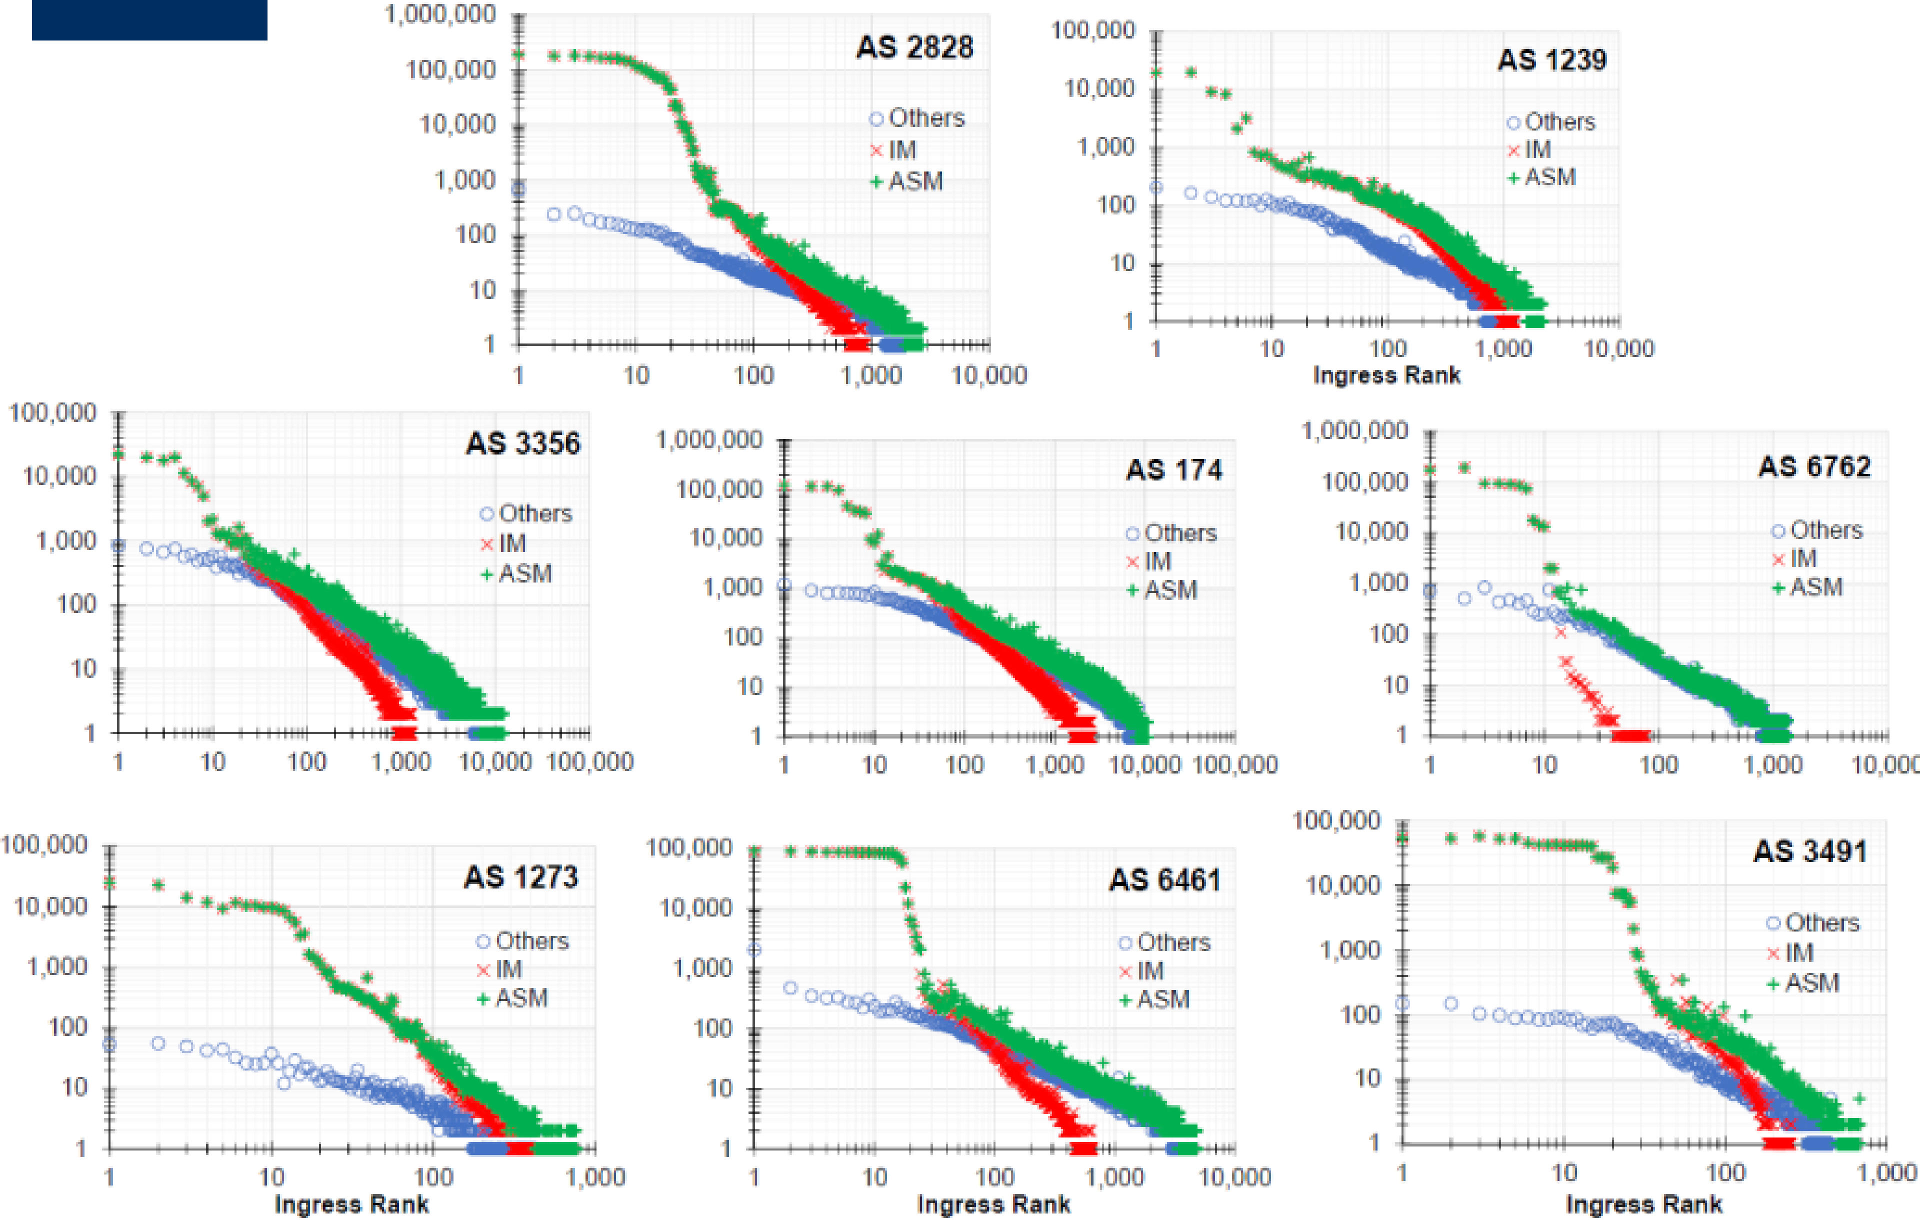

Summary

- We harvest path trace data
  - IPs, Edges, Prefixes of each AS
- Autonomous System Mapper (ASM)
  - comprehensively mapping the router-level topology of transit ASes
  - taking advantage of existing measurement platforms as a seed for ASM
  - Minimizing probing overhead by assigning VPs to Ingress-Subnet pair
    - partially trace toward subnet IPs based on their reachability through ingresses
  - Finding optimal VPs to measure from is an NP problem
    - devised a dynamic assignment mechnaism
- We observe that network discovery focusing on the ingress-subnet reachability of the vantage points considerably improves the measurements

Thanks

Principal Investigator
Mehmet Hadi Gunes
mgunes@unr.edu

Muhammed Abdullah Canbaz
mcanbaz@unr.edu

Jay Thom
jthom@cse.unr.edu

Khalid Bakhshaliyev
khalidb@nevada.unr.edu

Steven Fisher
stevenfisher@nevada.unr.edu

Ahmet Aksoy
aksoy@nevada.unr.edu

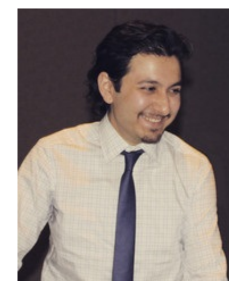

This material is based upon work supported by the National Science Foundation under grant numbers CNS-1321164 and EPS-IIA-1301726.