Graphs
that make the Net work

(analysis of prepending, shortest AS paths, atoms)

Andre Broido
broido@caida.org, CAIDA.SDSC.UCSD
kc claffy
kc@caida.org, UCSD/SDSC/CAIDA
9 dec 00
ISMA 0012

Graph: A set of nodes and edges/links

- Directed edges
- In– and outdegrees
- One– and two–way connectivity
- Connected components
- Shortest and longest paths
- Combinatorial core
- Giant component
- Perron–Frobenius eigenvector
- A useful model
- May not capture all properties

Nodes granularity

- MAC/Ethernet address
- IP address
- Network prefix
- BGP atom (def. below)
- AS number
- Administrative domain (all AS owned by one entity)

- Dual graph: make nodes links and vice versa

Other: host name, domain name, URL

Net as a graph: Levels 1–2

- 1. Nodes=devices (bridges, repeaters etc.); links=wires/fibers
- Example: LANs, SONET networks

- 2. Nodes=switches/hubs/routers, links=Layer 2 connections (skipped: repeaters, fiber segments)
- Examples: Ethernet, ATM

Net as a graph: Level 3 (IP)

- 3a. Nodes=IP interfaces, links=packets seen at both (no IP devices in between, TTL difference 1)
- repeaters, hubs, switches included in links
- 3b. Nodes=IP devices (routers, firewalls, caches, NAT boxes); links=observed packets (no IP devices in between, TTL difference 1)

Net as a graph: Levels 4–6

- 4. Nodes=IP network prefixes (SDSC: 132.249.0.0/16), links – packet seen at both (no networks in between)
- 5. Atoms: Nodes=groups of prefixes with equal AS paths, links – packets seen at both (no atoms in between)
- 6. AS: Nodes=Autonomous systems (SDSC: AS 195), links=packets seen at both (no AS in between)

Related graphs

- Web graph: documents and hyperlinks (see [Broder e.a. 2000])

- Domain name to IP address:
- Bipartite graph,
- Collection of K(m,n)'s
- Up to 14,400 domain names on one IP address

Dynamic properties

- All graphs change with time
- New equipment (nodes, links, firewalls) added
- New IP blocks alocated
- Renumbering out of old IP blocks
- "Death rate" of the Net
- Paths oscillate (load balancing)
- Paths fluctuate (routing instability)
- Paths flip (manual config)
- Outages, cuts, blackouts

Ambiguities

- Non-responding hops (no IP returned)
- Rate-limited response
- Private addresses
- Multicast addresses
- Addresses in 0.-2./8 blocks
- No matching BGP prefix
- Prefixes with multiple origin AS

Small connected components

giant component is 200 times larger than any other
largest of small components are /24s, /25s

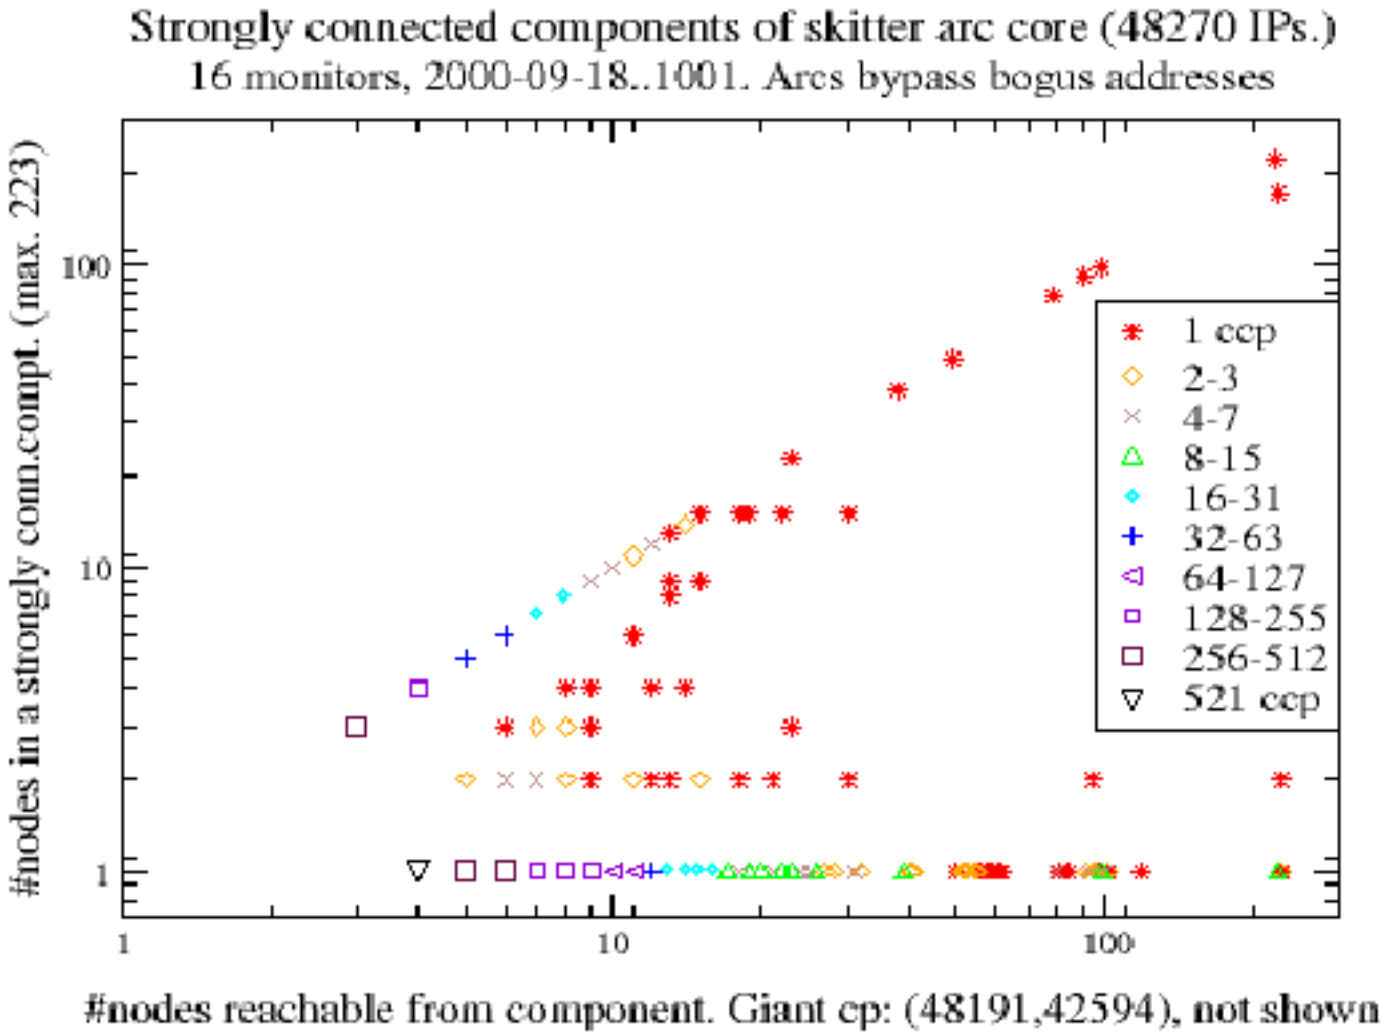

BGP data: sources

- BGP tables from David Meyer's Oregon Route Views
  - http://moat.nlanr.net/Routing/rawdata
  - MAE-East (Washington DC), MAE-West (Palo Alto), London, Amsterdam, Tokyo, Frankfurt, Ankara, Chicago, Johannesburg
- Looking glasses (BGP-enabled traceroute servers)

Analysis:

- www.telstra.net/ops/bgp-as-paths.html
- mirror.caida.org/~broido/bgp/bgp.html
- CAIDA's "Arctic views" (longitude/degree)

Uses of BGP data

- Aggregating IP to network prefixes
- Aggregating prefixes to origin AS
- Inferring contractual relations
- "Bird's eye view" of the Net – AS graph
- Predicting AS path taken by a packet???

Oregon BGP data

```
192.172.226.0/24    134.24.127.30    64   1740 195  1909 i
CAIDA network         peer (cerf)    med   cerf sdsc caida

192.172.226.0/24    204.147.128.141   –    145 195  1909 i
CAIDA network         peer (vbns)    med  vbns sdsc caida
```

AS path: 1740 195 1909
Grows from tail to head – "prepending":

CAIDA advertizes prefix to SDSC
SDSC advertizes CAIDA's prefix to CERF
ergo,
CERF "knows" how to send traffic to CAIDA

AS path length

- All AS vs. unique AS count:
- Prepending – repeating of AS
- Makes AS path look longer
- Reduces traffic via this path
- Bumps at AS path length 12,14
  - e.g., 202.183.247.0/24 3561 5400 5400 5727 4651 7568 7568 7568 7568 7568
  - 5 unique + 5 repeated ASes == AS path length 10

- Routes are chosen by Local Preference
- Other metrics:
  - Multiexit discriminator (MED)
  - Communities

but

- AS path length is used as default metric

Oregon Route Views, 2000-11-25

- 2.627M lines (prefix, peer, AS path)
- 34 contributing peers
- 28 peers with 81K-95K prefixes
- 2.534M lines in 28 tables

Out of those:

- 2,243,524 (88.5%) have no repeated AS
- 290,392 (11.5%) lines with repeated AS
- 176,576 (7%) end with repeated AS
- 2/5 repeated had it only in the middle
- 0.45% suppressed, dampened or history

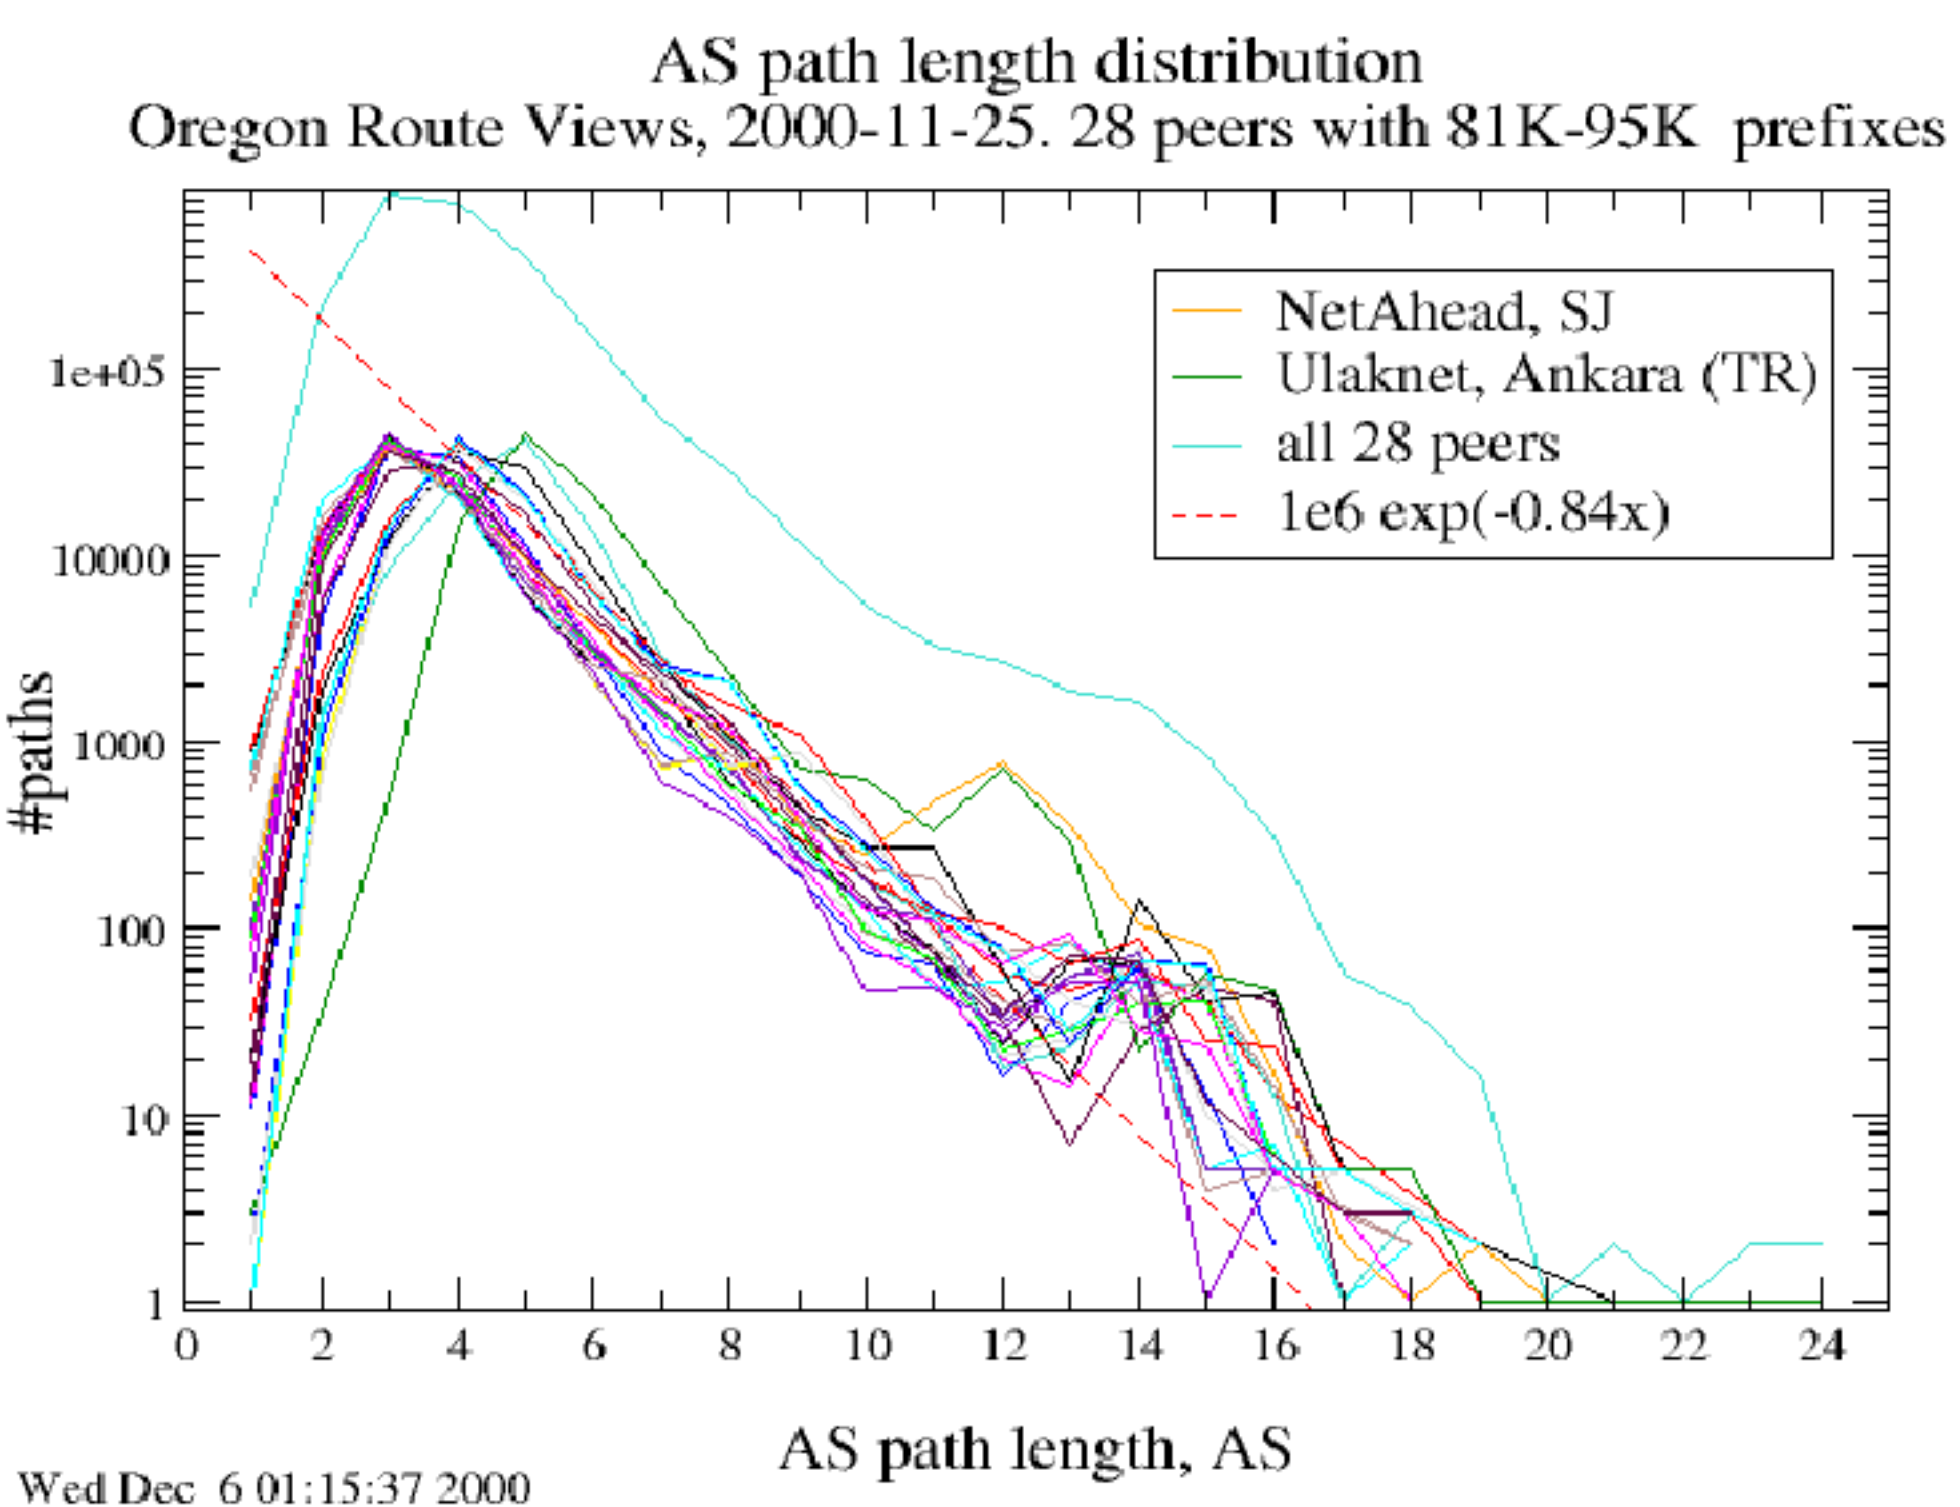
AS path length distribution
Oregon Route Views, 2000-11-25. 28 peers with 81K-95K prefixes
NetAhead, SJ
Ulaknet, Ankara (TR)
all 28 peers
1e6 exp(-0.84x)
#paths
1e+05
10000
1000
100
10
1
0
2
4
6
8
10
12
14
16
18
20
22
24
AS path length, AS
Wed Dec 6 01:15:37 2000

Path length measured in unique AS

- Much smaller tail: exp(−1.8x) vs. exp(−0.84x)
- Much smoother
- Close to Gamma distrib.

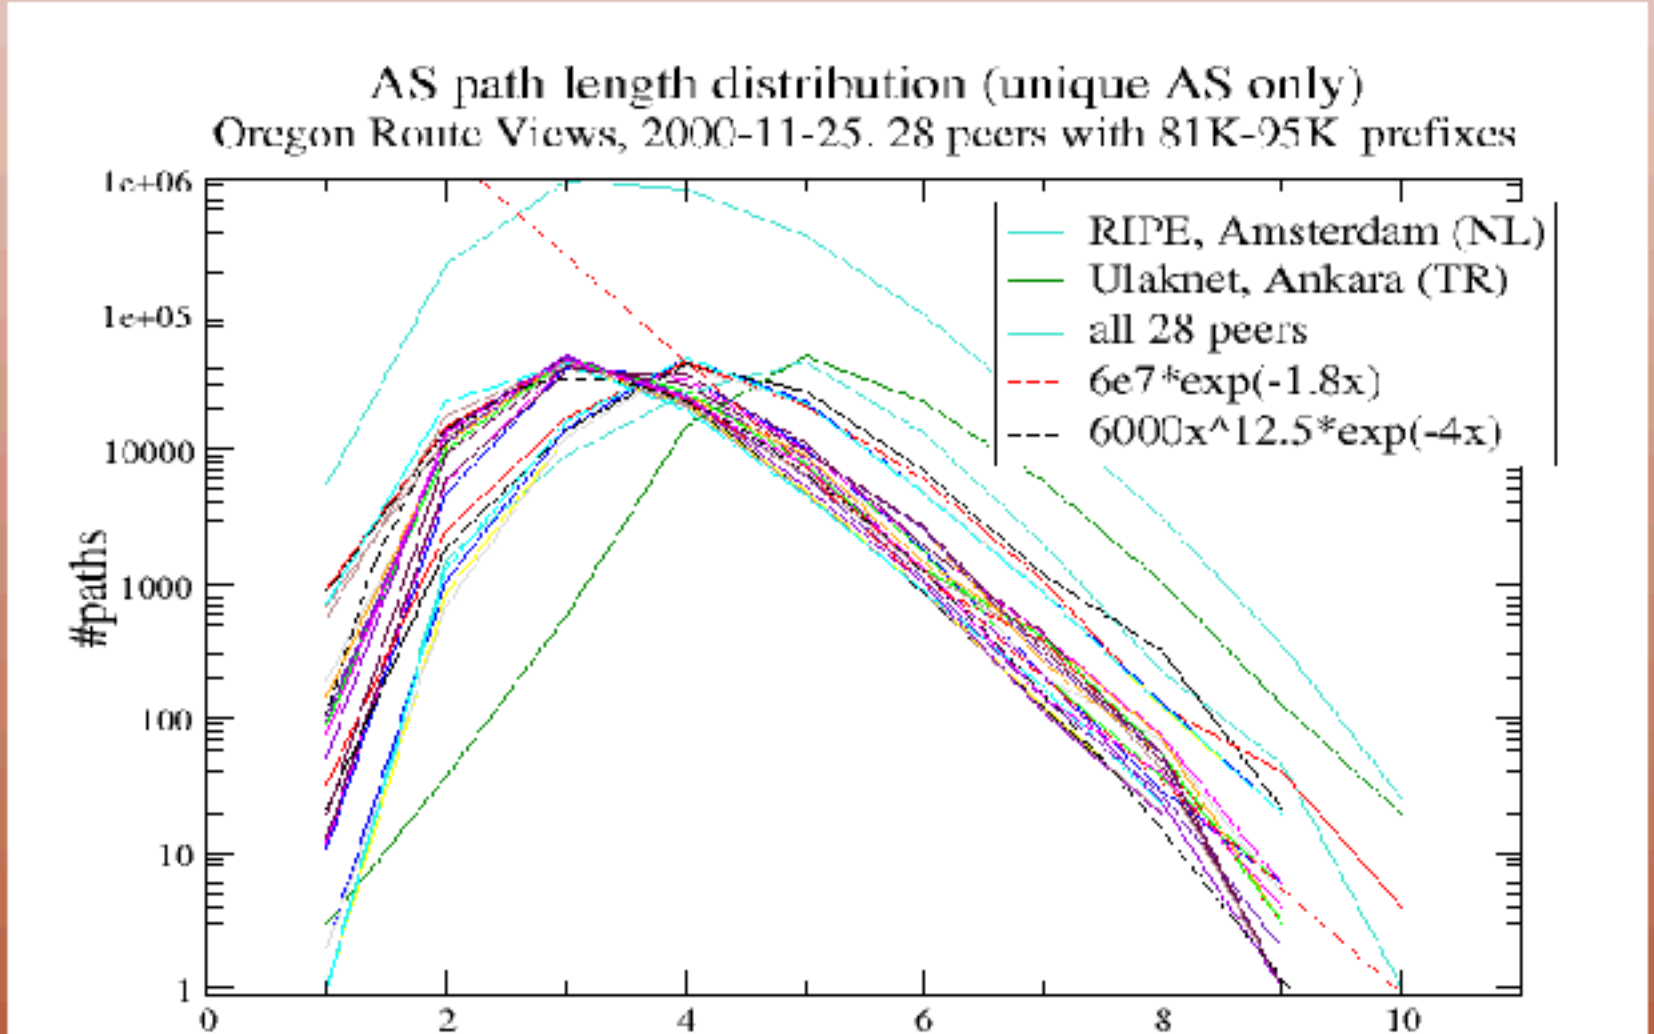

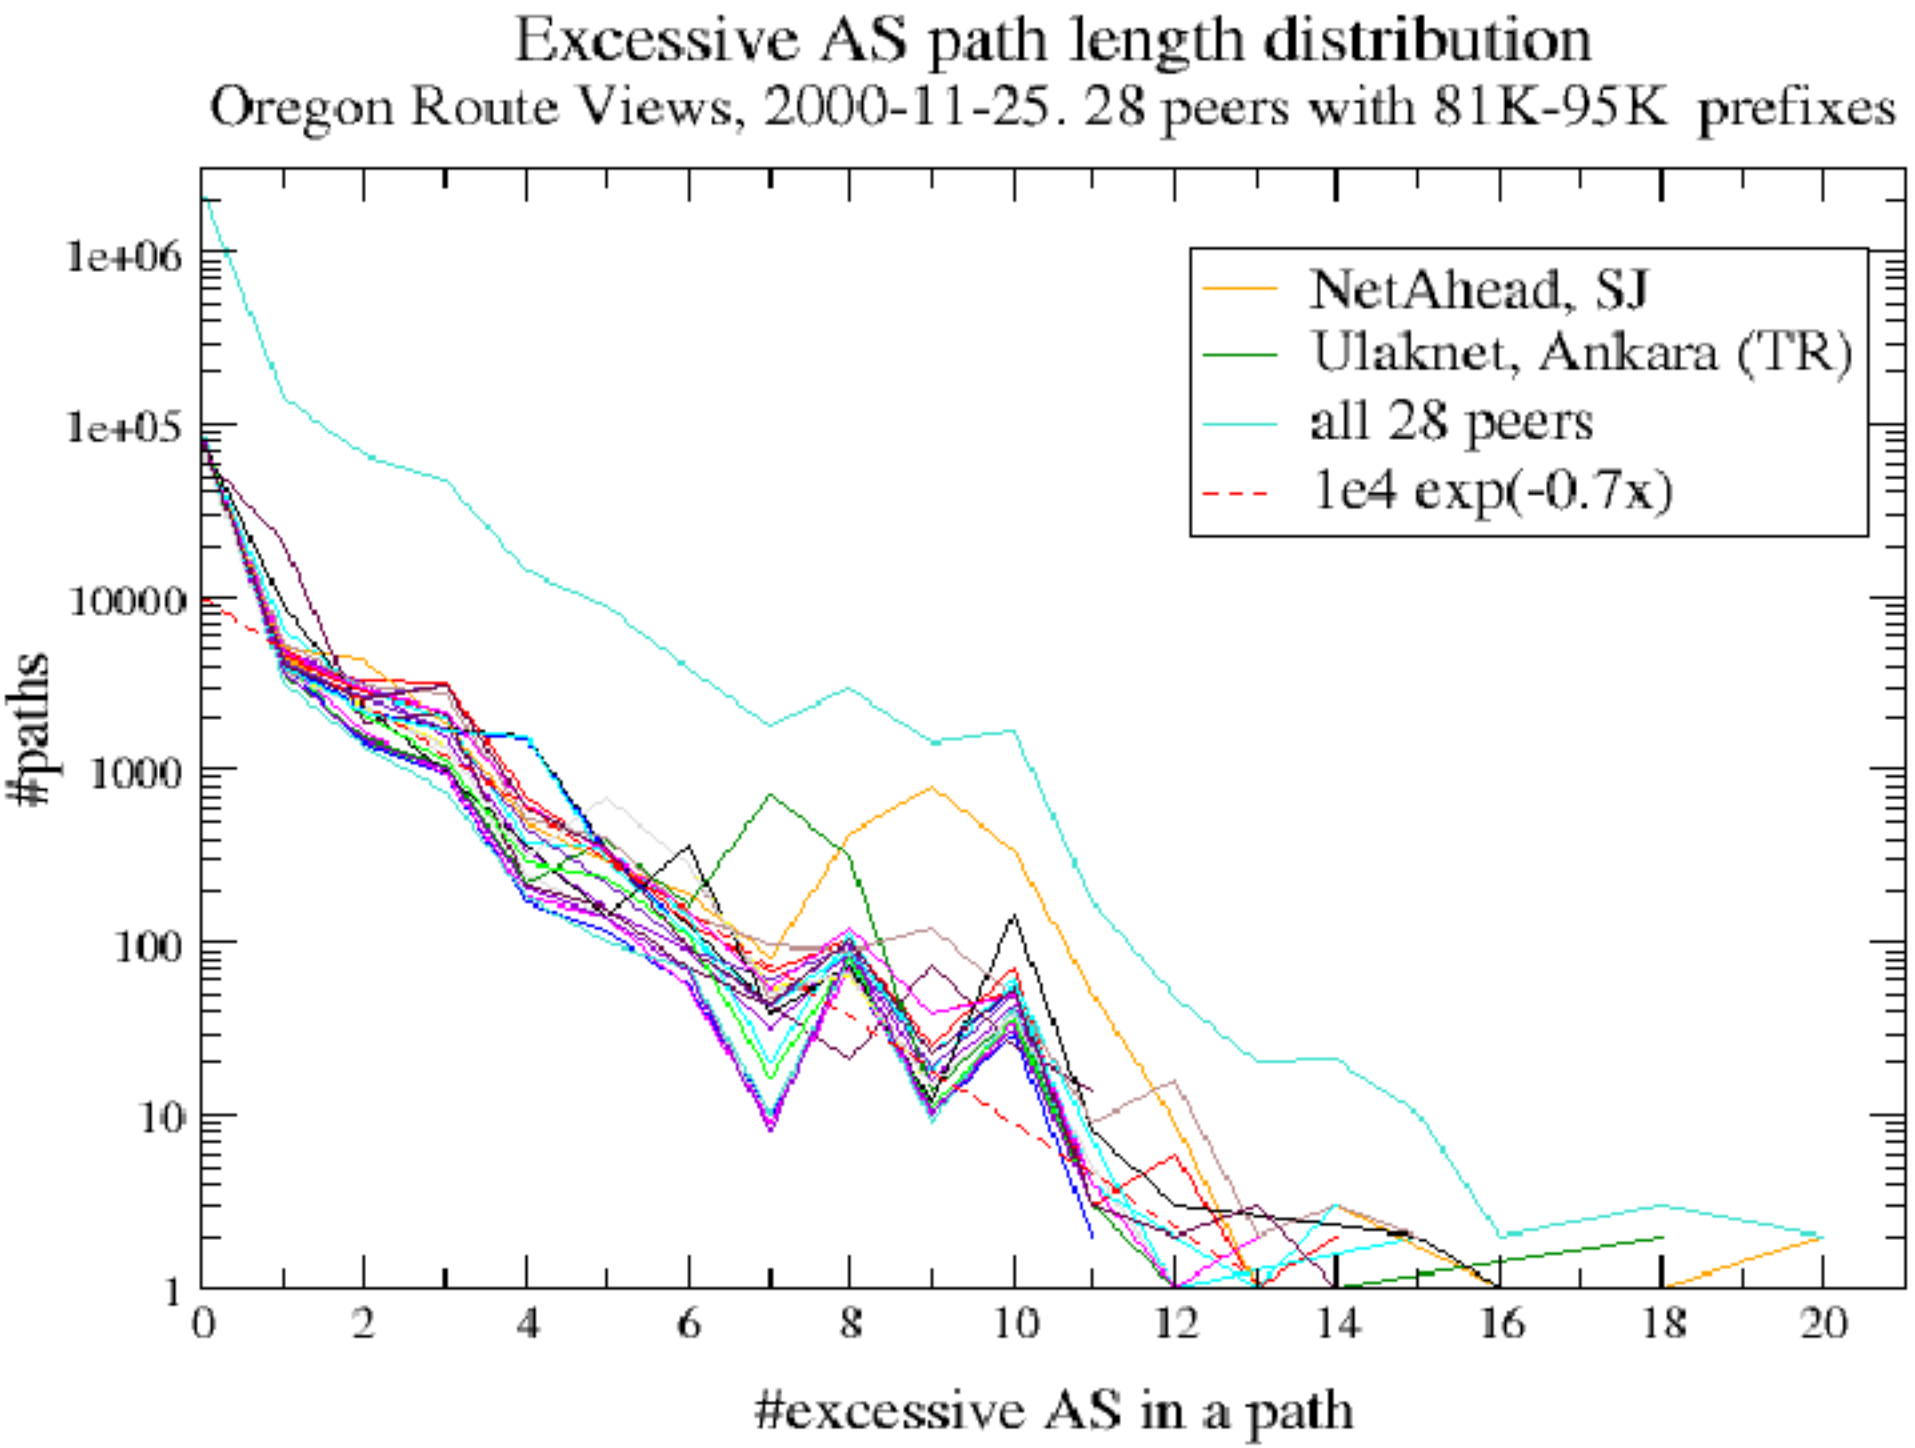
Excessive AS path length distribution
Oregon Route Views, 2000-11-25. 28 peers with 81K-95K prefixes
NetAhead, SJ
Ulaknet, Ankara (TR)
all 28 peers
1e4 exp(-0.7x)
#paths
1e+06
1e+05
10000
1000
100
10
1
0
2
4
6
8
10
12
14
16
18
20
#excessive AS in a path
Wed Dec 6 01:27:44 2000

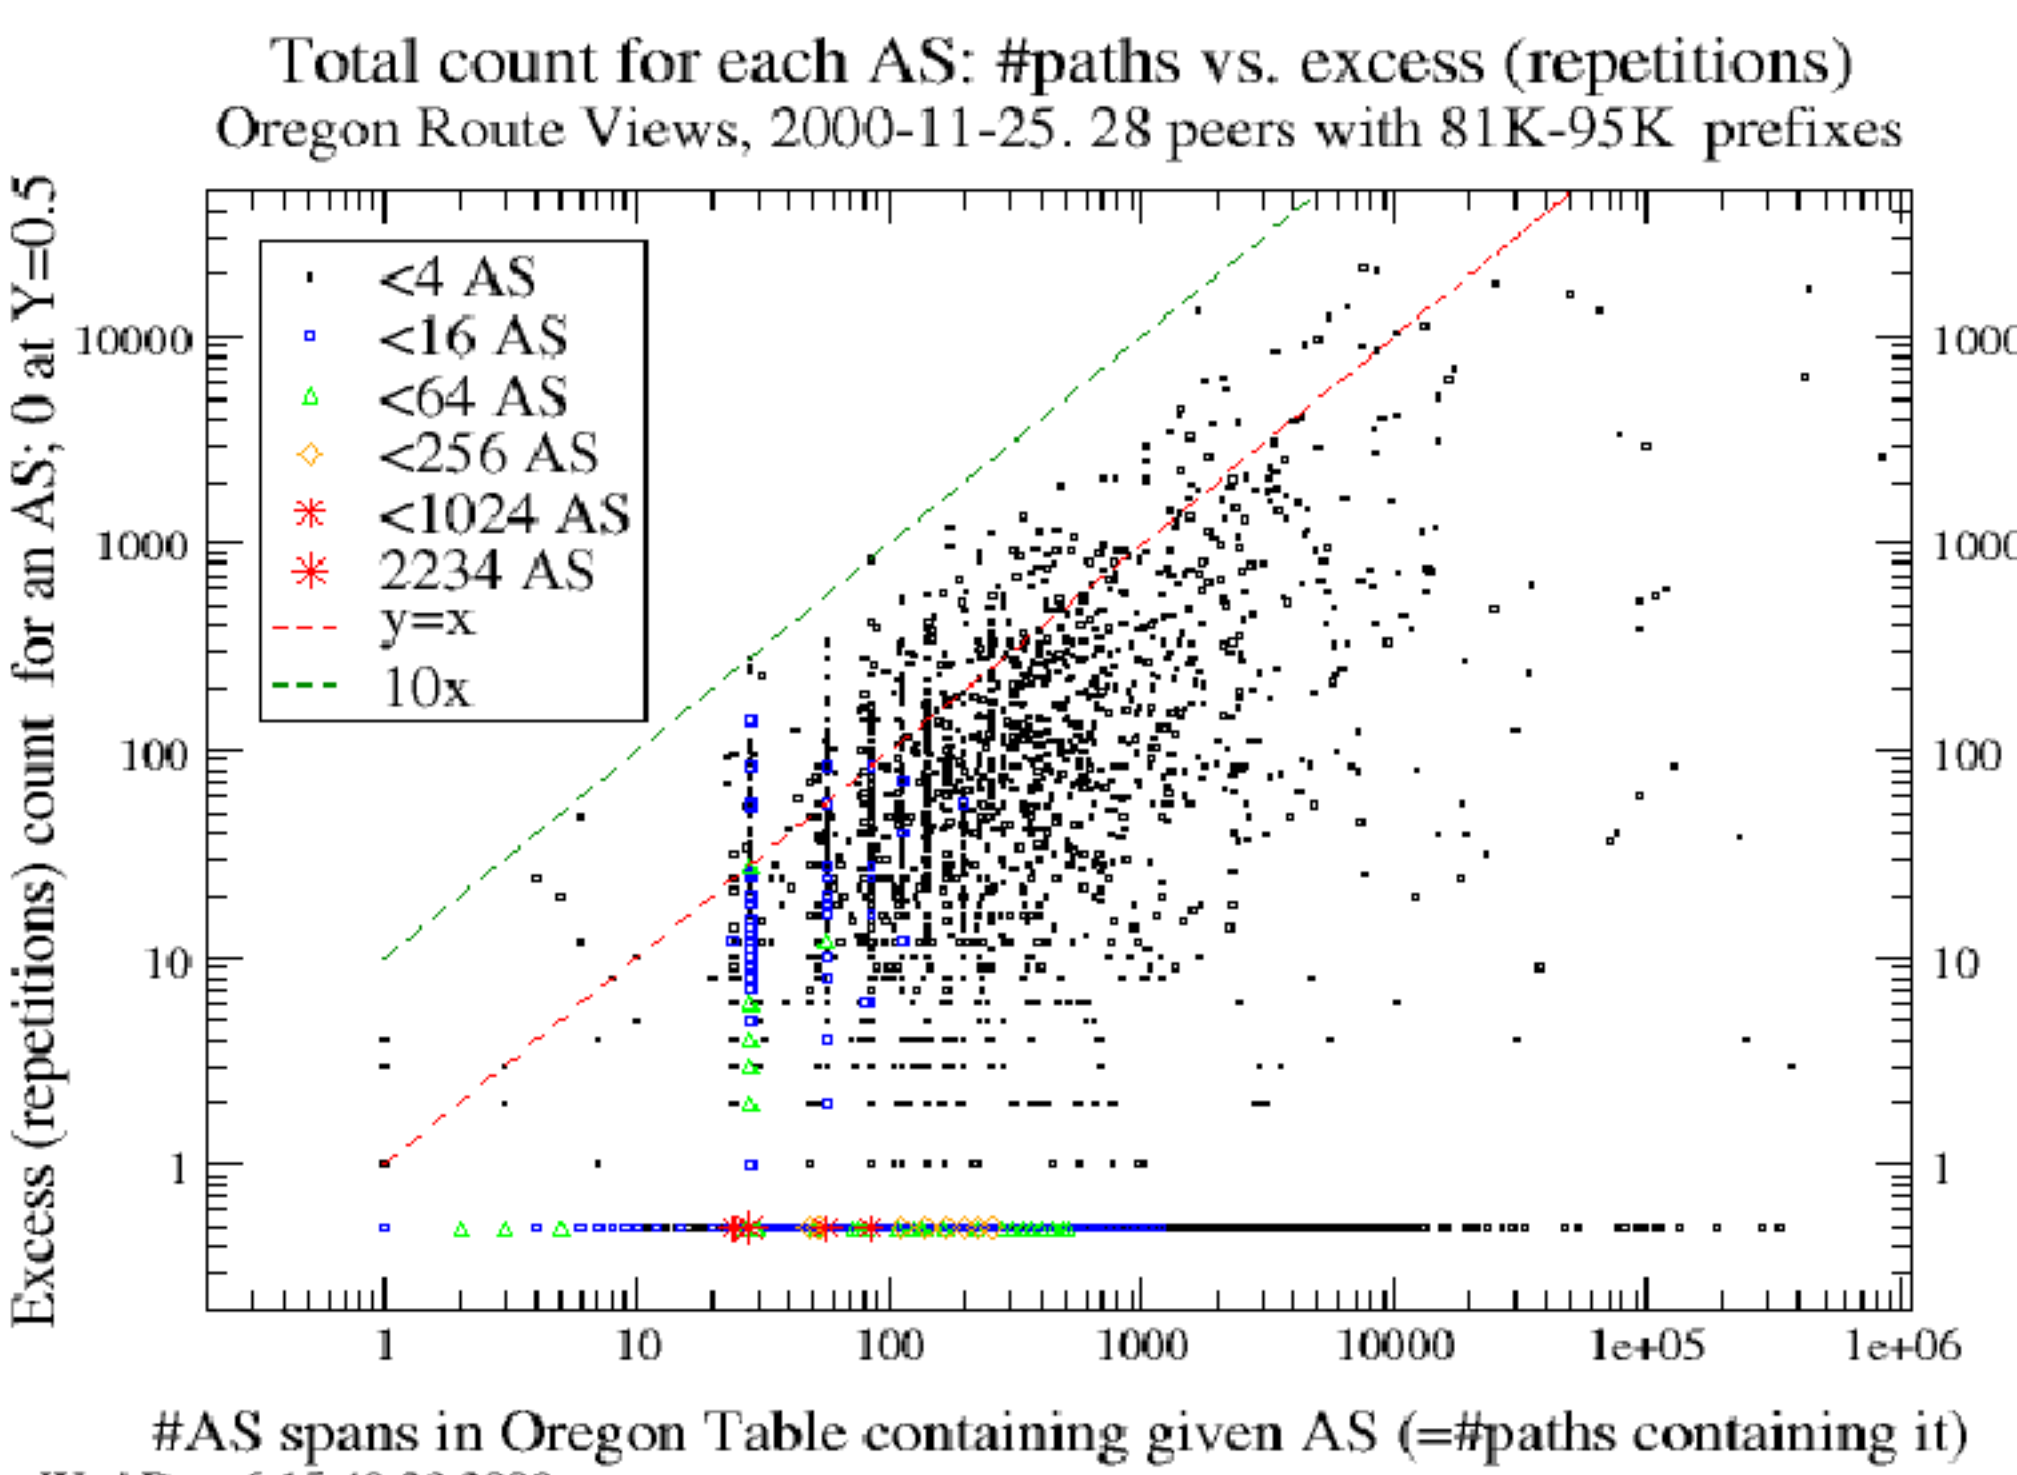
Total count for each AS: #paths vs. excess (repetitions)
Oregon Route Views, 2000-11-25. 28 peers with 81K-95K prefixes
Excess (repetitions) count for an AS; 0 at Y=0.5
<4 AS
<16 AS
<64 AS
<256 AS
<1024 AS
2234 AS
y=x
10x
10000
1000
100
10
1
1
10
100
1000
10000
1e+05
1e+06
#AS spans in Oregon Table containing given AS (=#paths containing it)
Wed Dec 6 15:48:23 2000

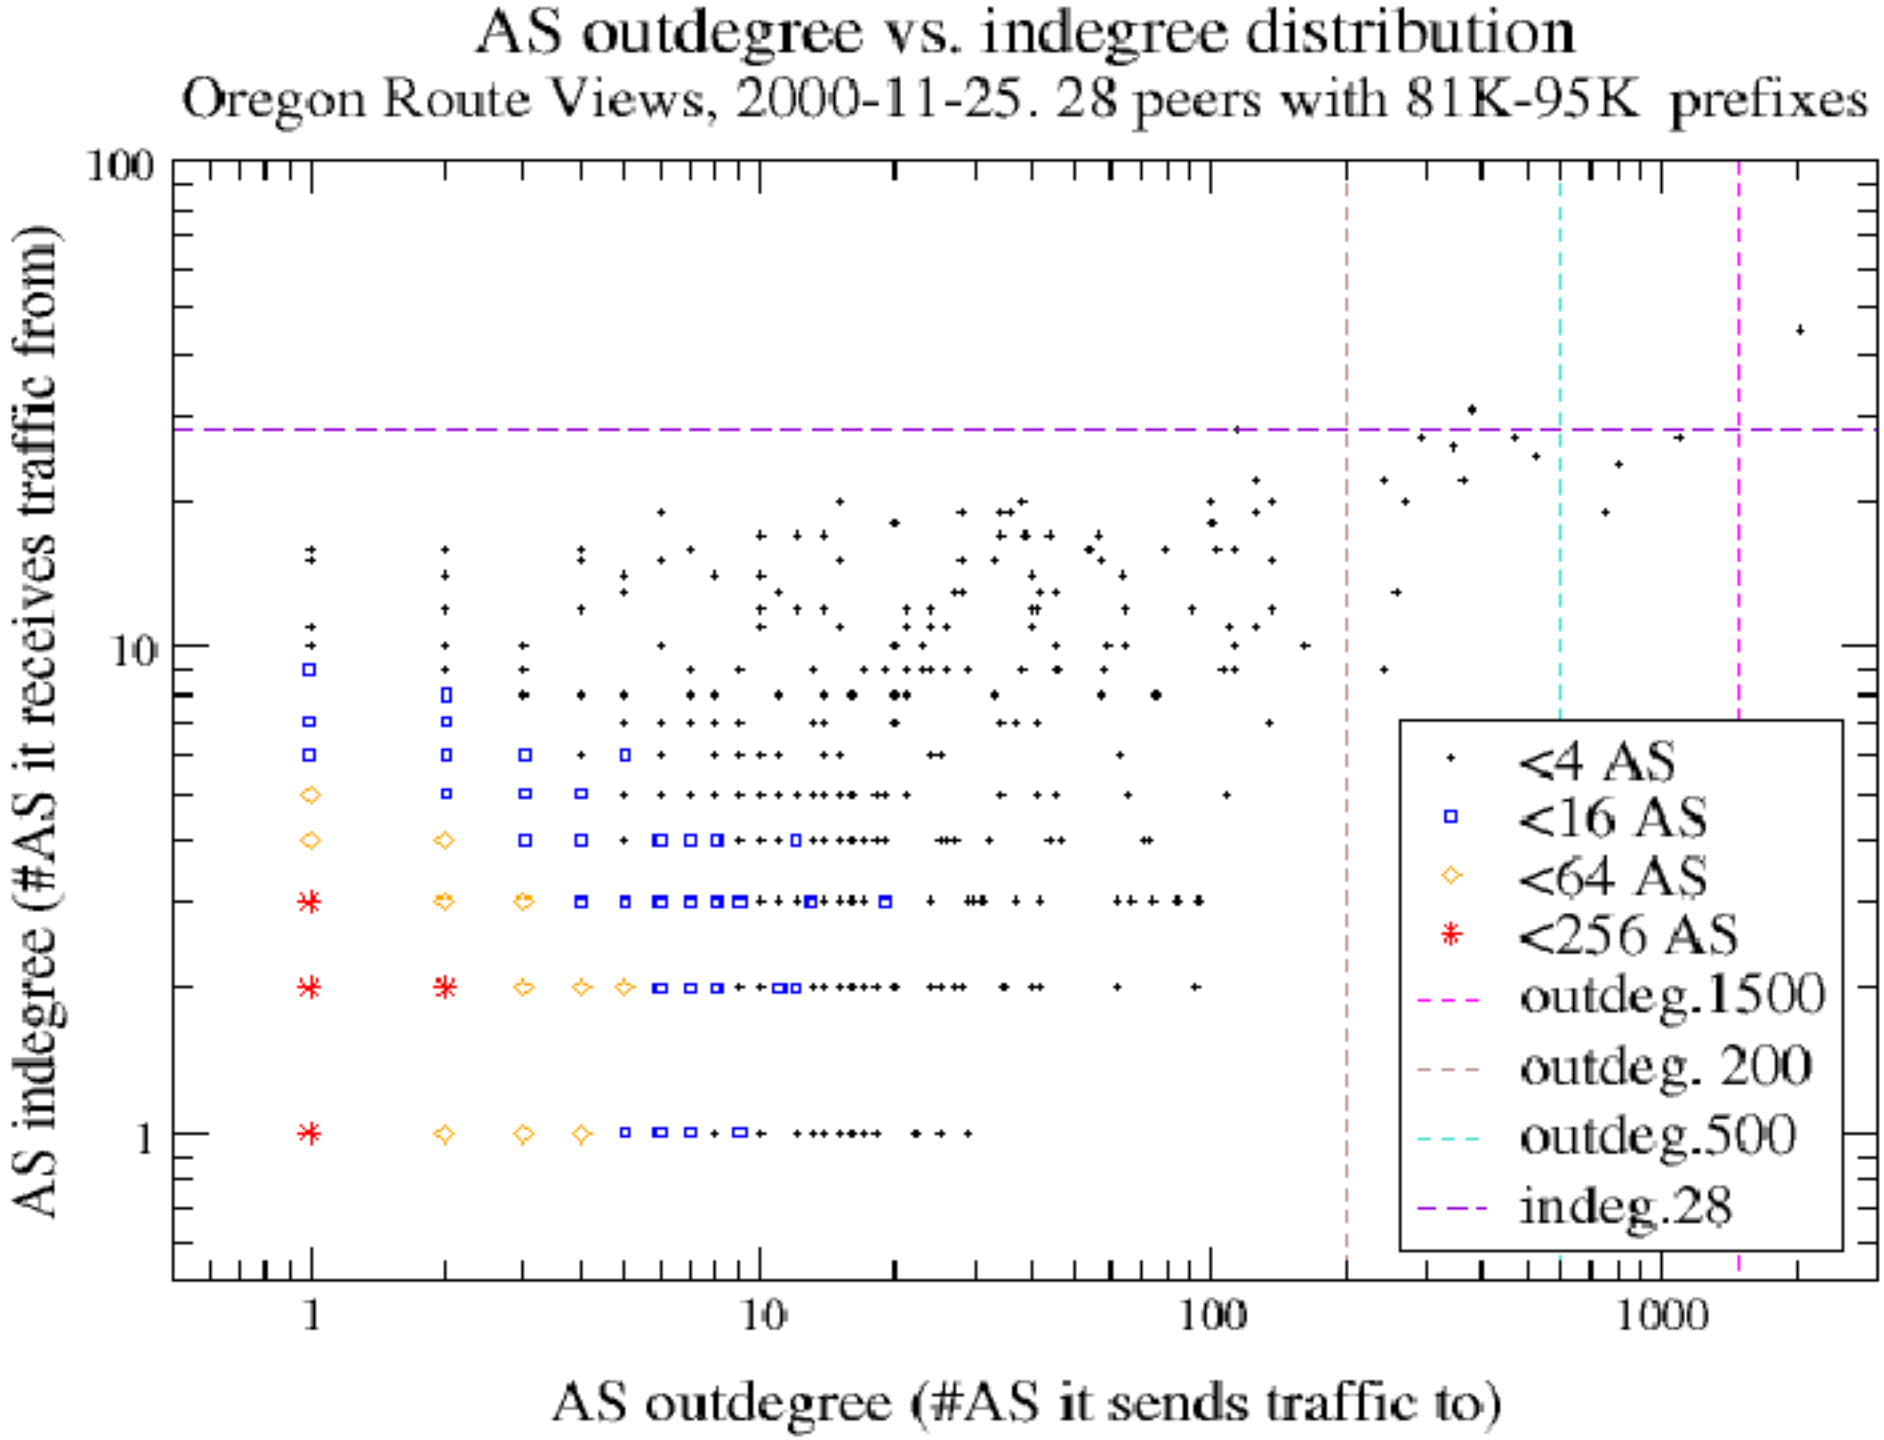
AS outdegree vs. indegree distribution
Oregon Route Views, 2000-11-25. 28 peers with 81K-95K prefixes
AS indegree (#AS it receives traffic from)
AS outdegree (#AS it sends traffic to)
100
10
1
1
10
100
1000
<4 AS
<16 AS
<64 AS
<256 AS
outdeg.1500
outdeg. 200
outdeg.500
indeg.28

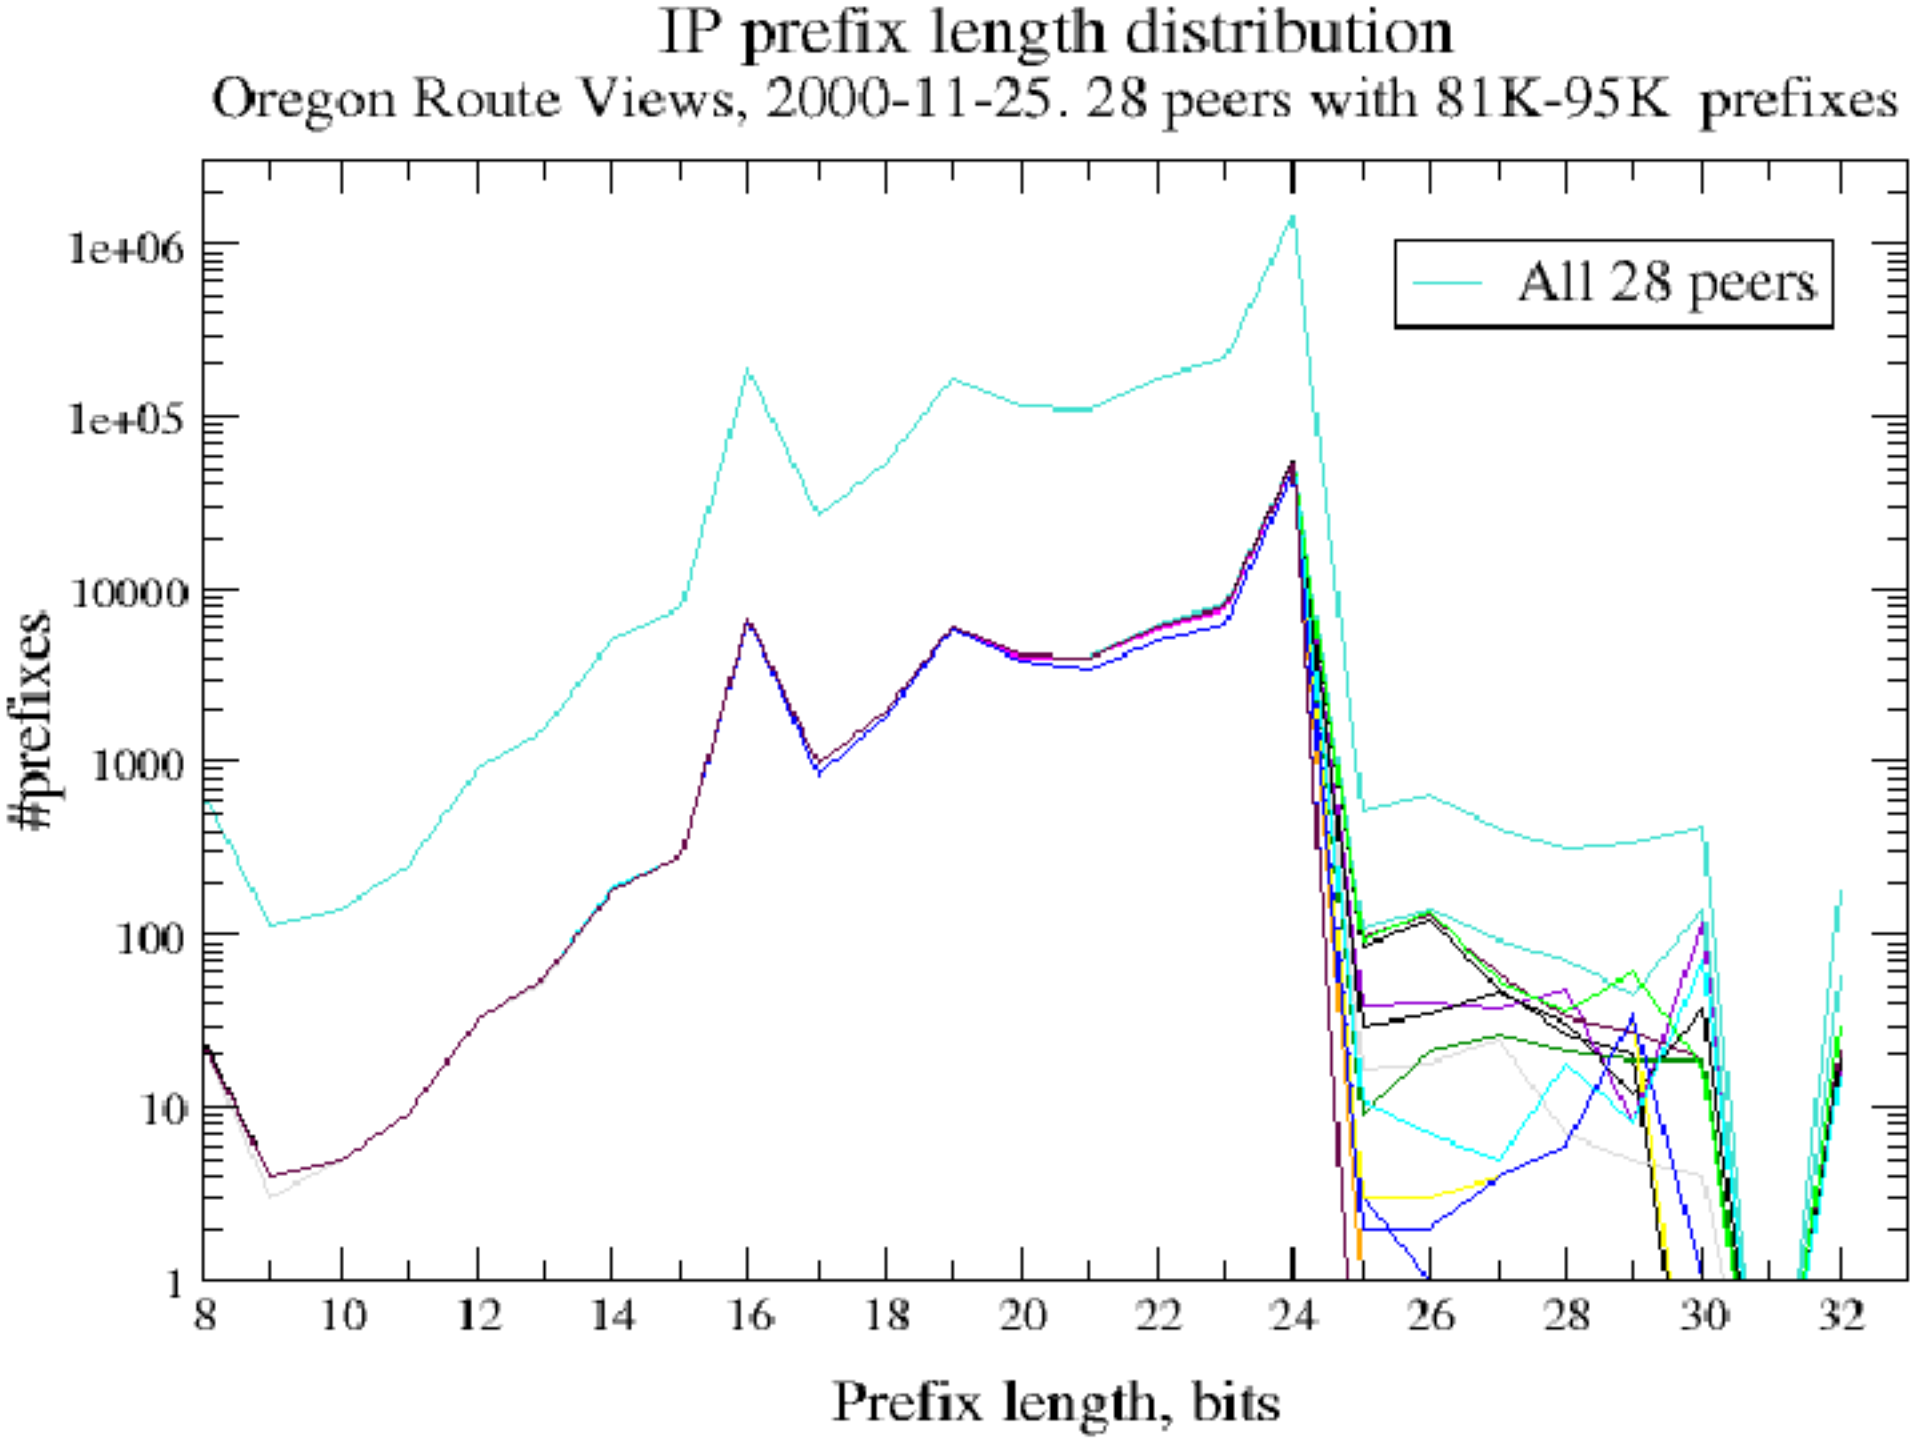
IP prefix length distribution
Oregon Route Views, 2000-11-25. 28 peers with 81K-95K prefixes
All 28 peers
#prefixes
1e+06
1e+05
10000
1000
100
10
1
8
10
12
14
16
18
20
22
24
26
28
30
32
Prefix length, bits
Thu Dec 7 14:49:04 2000

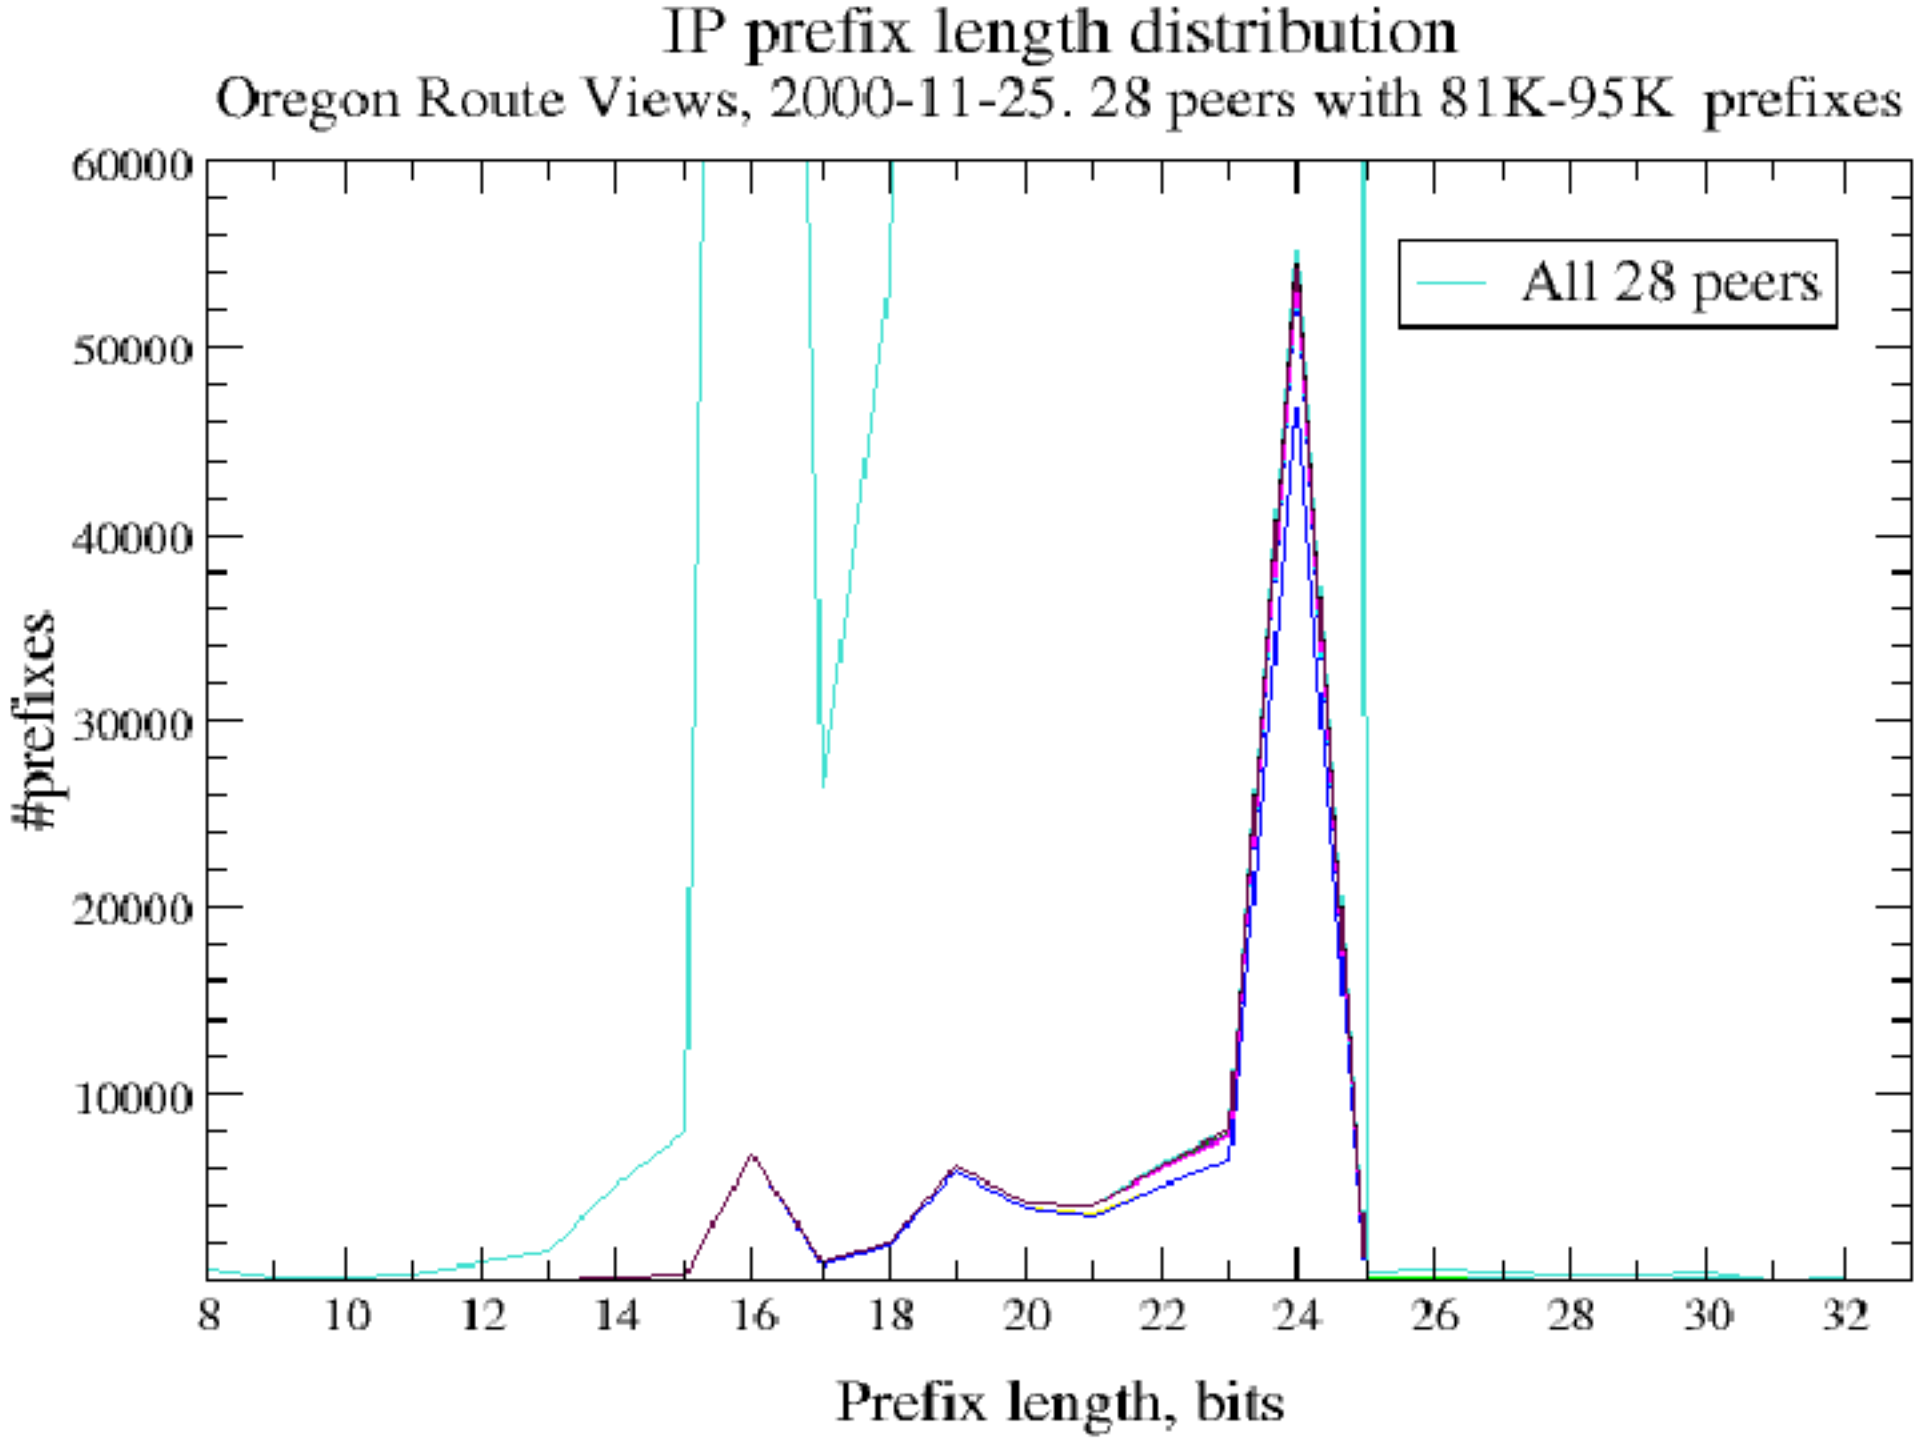
IP prefix length distribution
Oregon Route Views, 2000-11-25. 28 peers with 81K-95K prefixes
All 28 peers
#prefixes
60000
50000
40000
30000
20000
10000
8
10
12
14
16
18
20
22
24
26
28
30
32
Prefix length, bits
Thu Dec 7 14:45:11 2000

Prefix distribution

- most frequent: /16, /19 and /24
- Almost 60% are /24
- Almost no prefixes over /24
- Some providers filter prefixes over /24, e.g. Ulaknet (TR)

Atoms: new unit of routing/connectivity analysis

Atoms are:

- equivalence class of prefixes that share same set of AS paths as seen e.g. by peers in Oregon Route Views BGP tables
- unit of granularity in between prefixes and ASes
- we propose as new way to analyze routing system

Atom size in prefixes

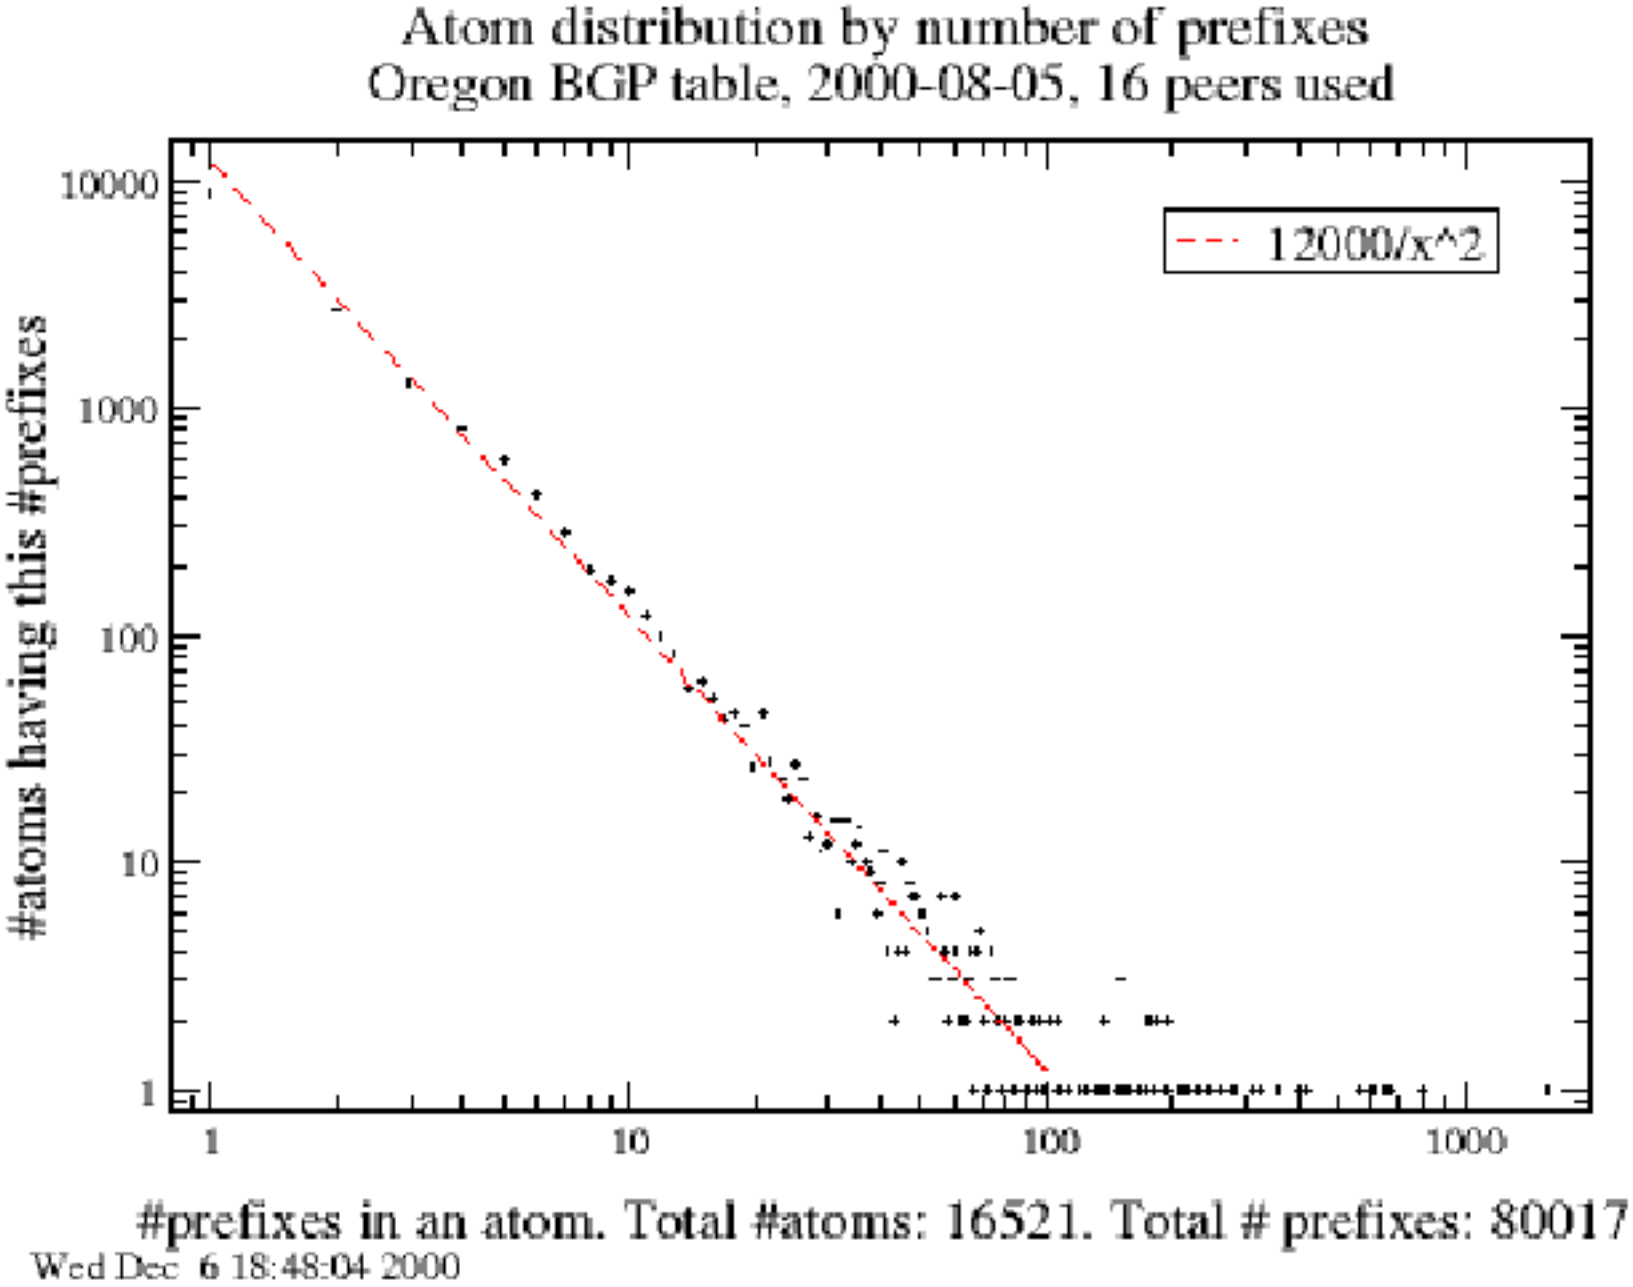

Number of atoms per AS

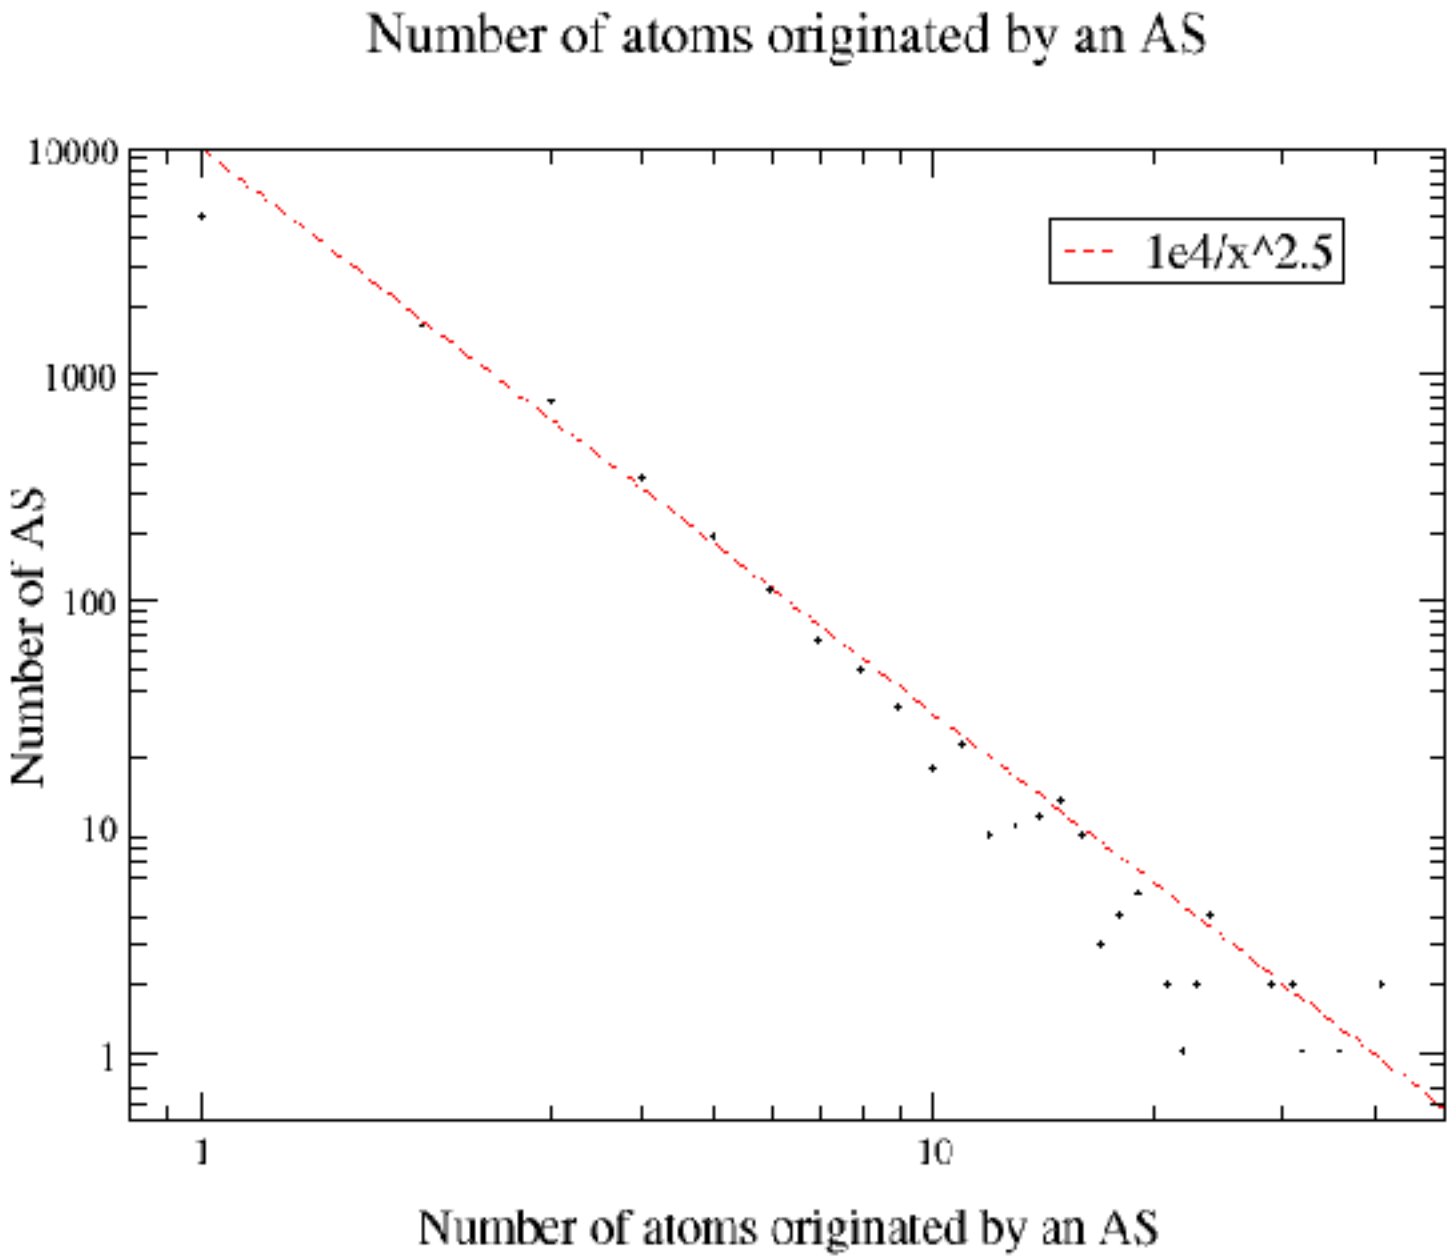

atoms	#AS	percent
1	4943	59.84
2	1647	19.94
3	750	9.08
4	347	4.20
5	191	2.31
6	111	1.34
7	67	0.81
8	49	0.59
9	34	0.41
10	18	0.22
11	23	0.28
12	9	0.11
13	10	0.12
14	11	0.13
15	13	0.16
16	9	0.11
17	3	0.04
18	4	0.05
19	5	0.06
21	2	0.02
22	1	0.01
23	2	0.02
...		
Total	8261	

http://www.caida.org/outreach/isma/0012/talks/kcandre/

Andre Broido
kc claffy
UCSD/SDSC/CAIDA
kc@caida.org
www.caida.org