New Methods for Topology Validation and Alias Resolution

Adam Bender
University of Maryland

Joint work with Rob Sherwood and Neil Spring

IPAM Internet MRA Workshop 1 - 9/26/08

Part 1:
Topology discovery and verification

Discarte: A Disjunctive Internet Cartographer
R. Sherwood, A. Bender, N. Spring
SIGCOMM 2008

Internet Topology

- Router-level maps of the Internet:
  - Are useful
    - Fault diagnosis, modeling, simulation of new protocols
  - Are difficult to obtain
    - Operators are reticent
    - No technological support for inference

Traceroute

- TTL (time to live)-limited probes
- Each probe returns information about a single IP address
- Sequential probes divulge links (usually)
- Prone to many errors, mostly due to incompleteness

Traceroute

```
 1  cat3750 (128.97.46.254)  3.202 ms
 2  ipam--murphy.backbone.ucla.net (169.232.49.229)  1.162 ms
 3  murphy--core-2-10ge.backbone.ucla.net (169.232.4.26)  1.311 ms
 4  core-2--border-1-10ge.backbone.ucla.net (169.232.4.103)  1.378 ms
 5  lax-hpr1--ucla-10ge.cenic.net (137.164.27.5)  3.763 ms
 6  nlr-packetnet--hpr-lax-hpr.cenic.net (137.164.26.131)  18.616 ms
 7  hous-losa-87.layer3.nlr.net (216.24.186.31)  33.291 ms
 8  atla-hous-70.layer3.nlr.net (216.24.186.9)  60.284 ms
 9  wash-atla-64.layer3.nlr.net (216.24.186.21)  70.235 ms
10  216.24.184.11 (216.24.184.11)  70.933 ms
11  xe-7-2-0-0.clpk-t640.maxgigapop.net (206.196.178.89)  71.524 ms
12  umd-i2-rtr.maxgigapop.net (206.196.177.126)  71.709 ms
13  gi3-5.ptx-fw-r1.net.umd.edu (129.2.0.233)  71.686 ms
14  gi5-8.css-core-r1.net.umd.edu (128.8.0.85)  71.941 ms
15  gi3-8.ptx-core-r1.net.umd.edu (129.2.0.1)  71.844 ms
16  gi4-1.ptx-dual-r2.net.umd.edu (129.2.0.114)  71.918 ms
17  128.8.239.70 (128.8.239.70)  71.901 ms
```

Problems with traceroute

- Only shows incoming interface on router
- Often blocked by network admins
- Paths can change mid-trace, implying false links

IP addresses are interfaces

- All routers have multiple interfaces
- Many-to-one mapping from interfaces to routers
- Two IPs on same router are aliases

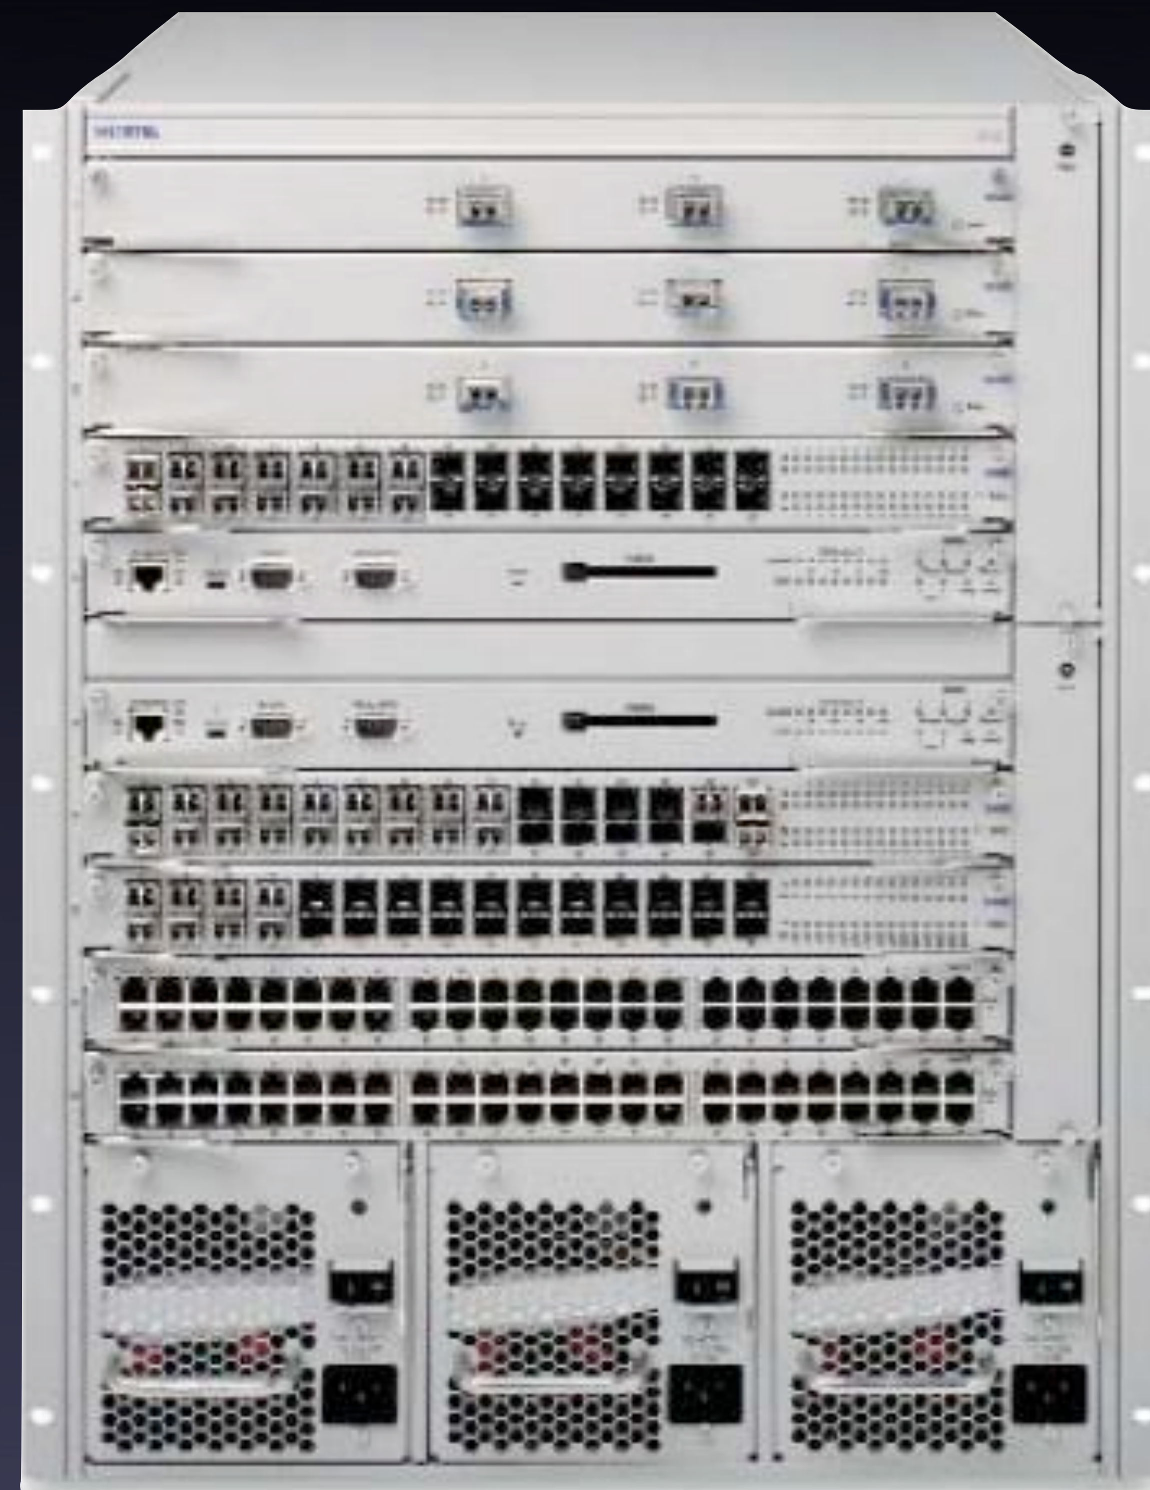

Alias resolution

- Process of mapping IP addresses (interfaces) to routers
- Necessary for accurate router count and degree distribution
  - False aliases collapse the network
  - Missed aliases inflate path diversity

Importance of alias resolution

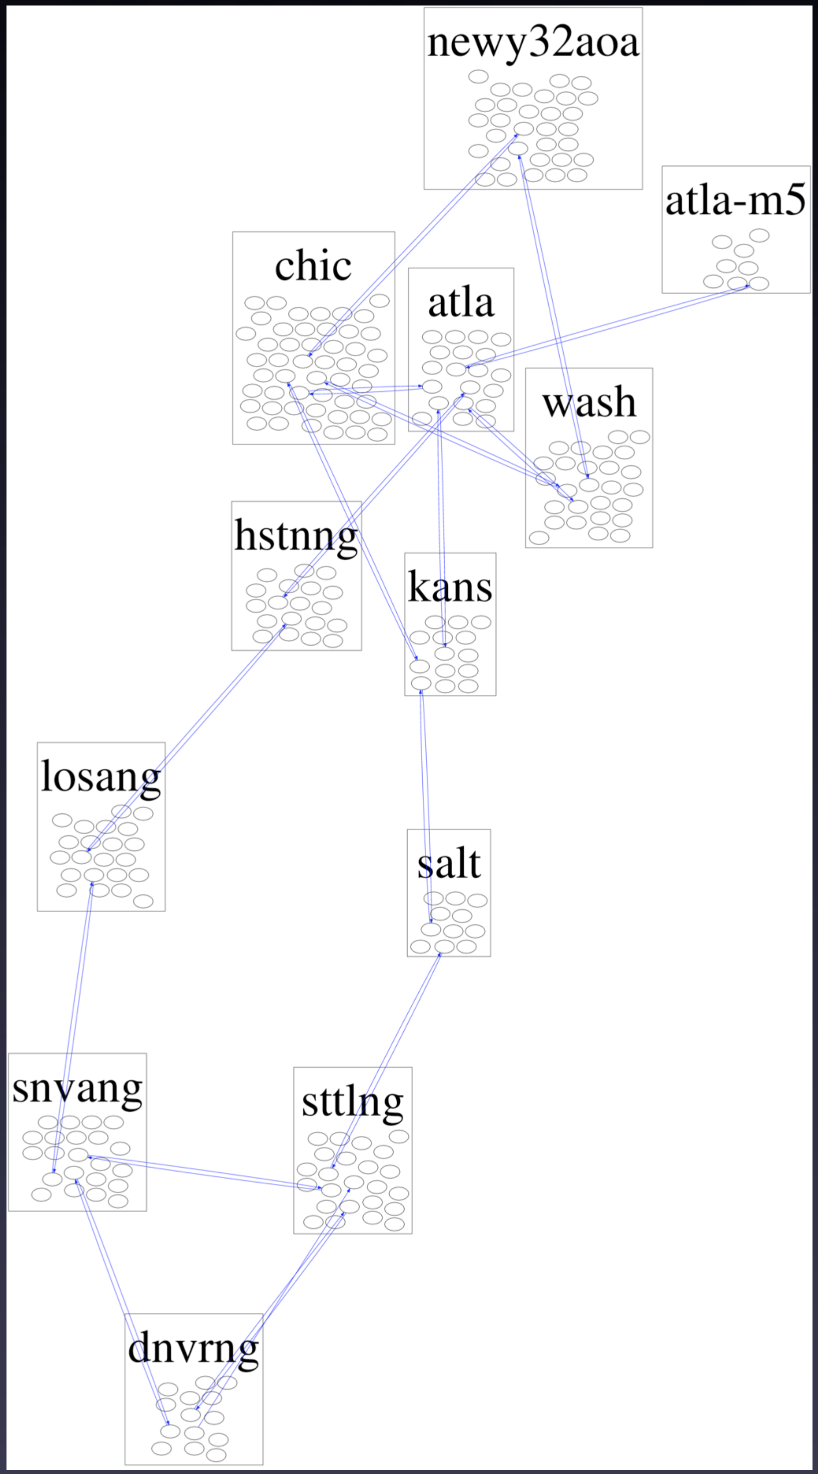

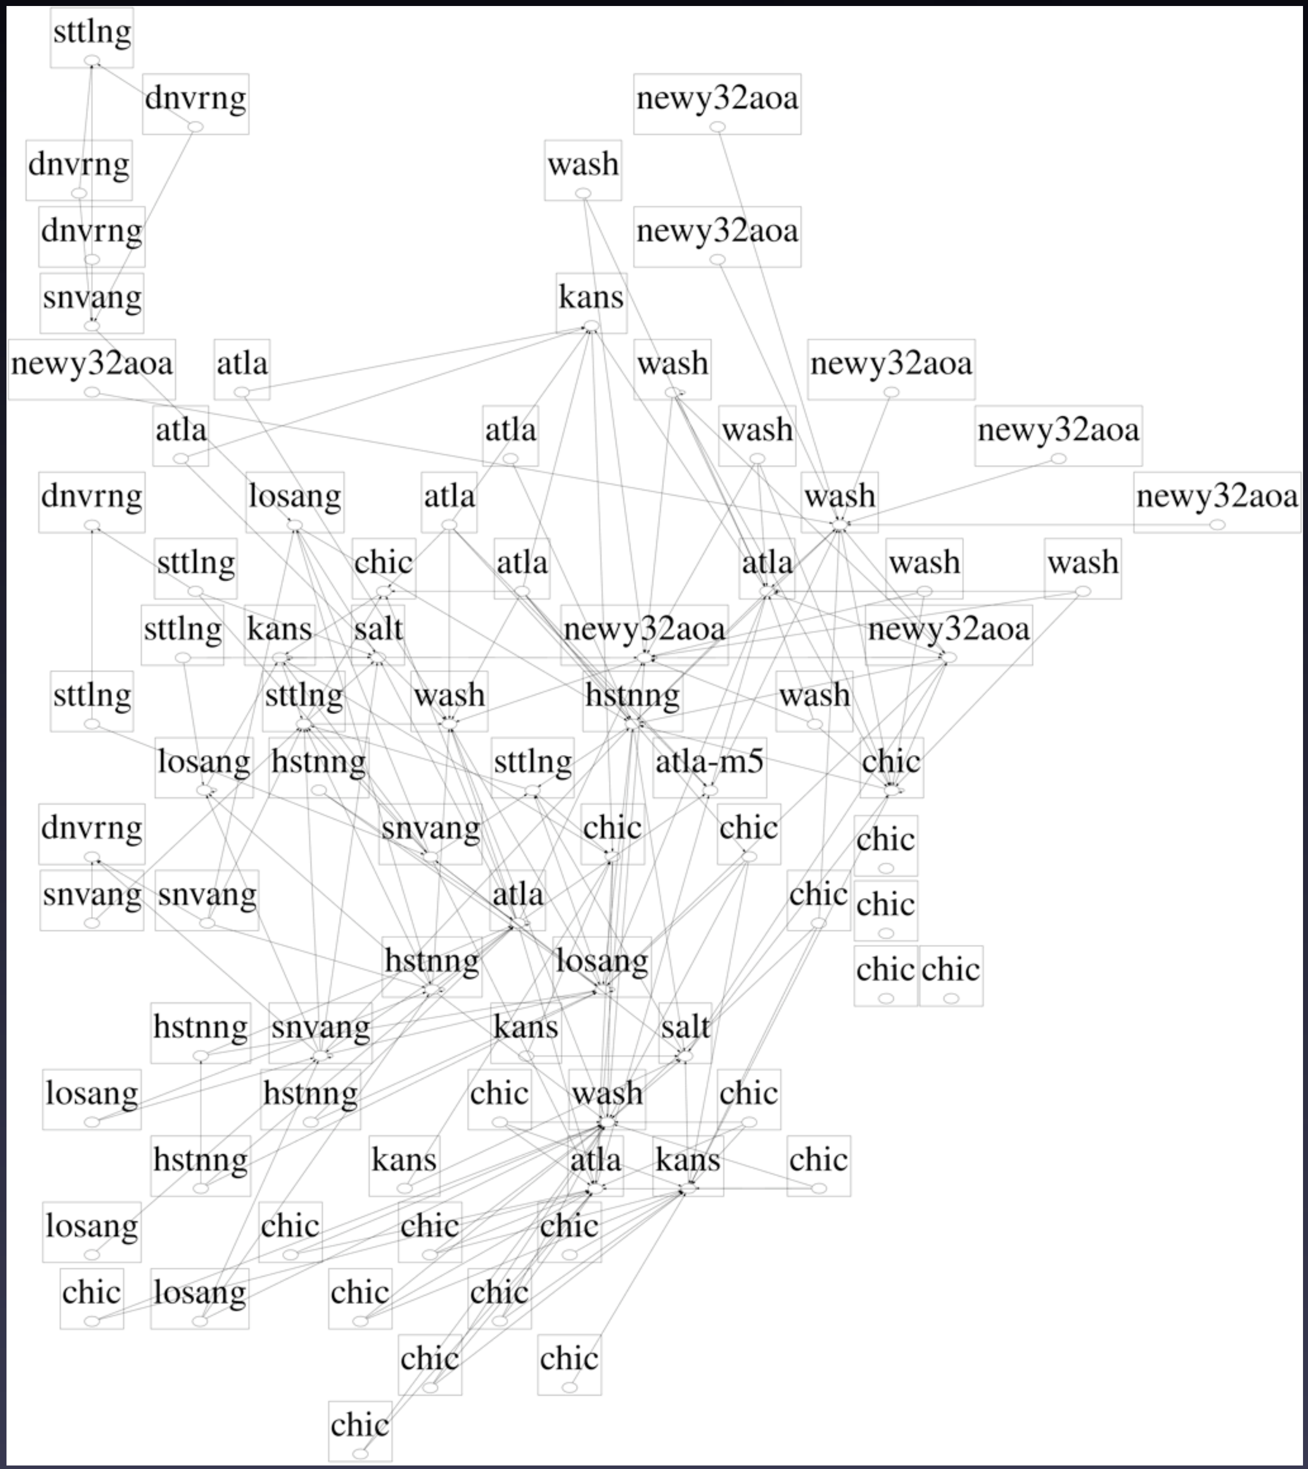

Another source of information

- Record route (RR) is another way to learn IP addresses on a path
- Will also allow us to do alias resolution
- Myth: routers drop RR packets
  - Reality: only 1% do

Record route

- IP option
- Stores first 9 IP addresses in header
- Recoverable from ICMP error response
- Unlike traceroute, stores outgoing interface

Address alignment

- Match addresses that are obtained with traceroute and those obtained with record route
- Allows alias resolution without direct probing

Address alignment

Traceroute:

Address alignment

Traceroute:

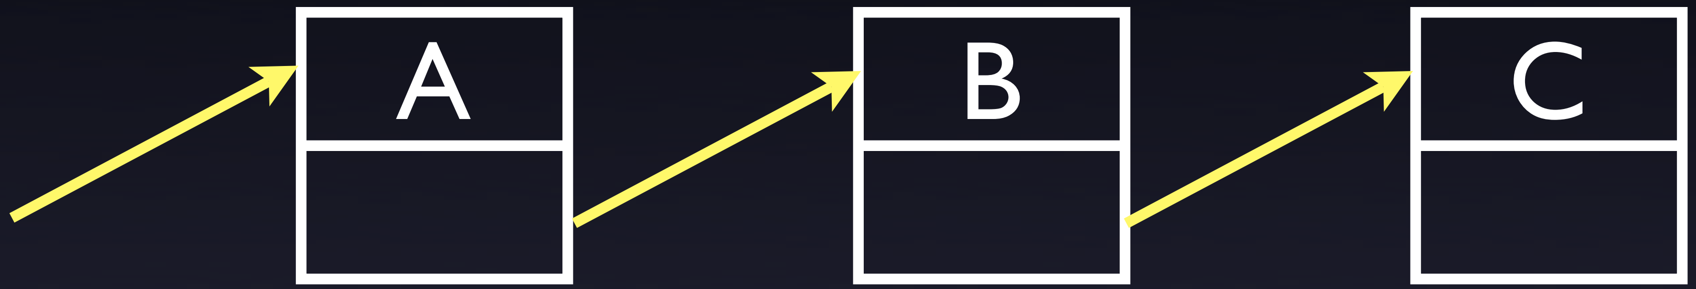

Record route:

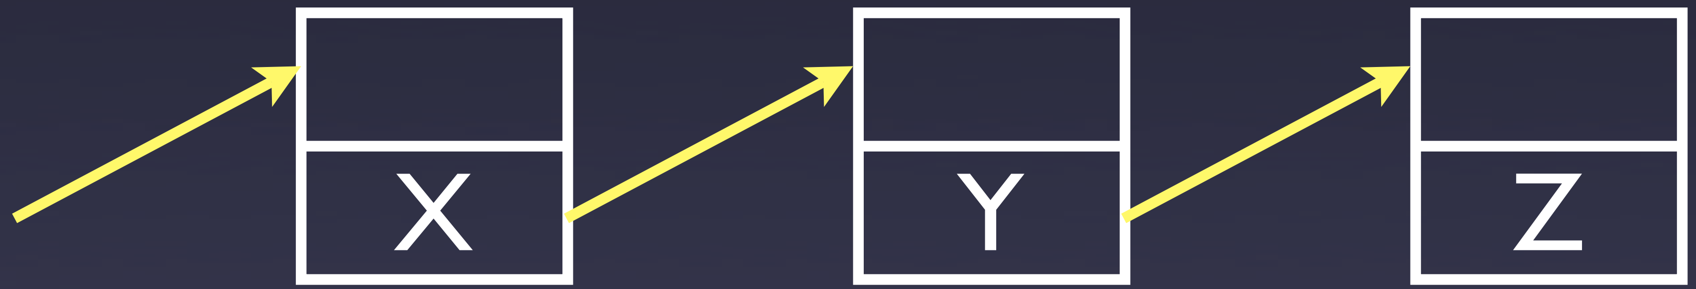

Address alignment

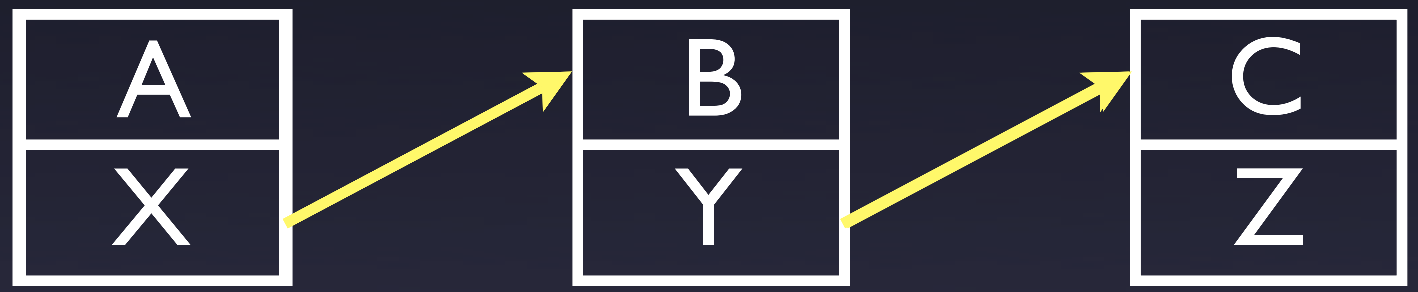

Alignment frustrations

- At most 9 addresses stored in a probe
  - Probe from multiple sources (PlanetLab)
- Record route is understandardized
  - 7 different implementations
  - Some don’t decrement TTL; some don’t write address into RR field

Disjunctive logic programming

- Technique from AI community
- Goal is to model the implementation type of each router
  - Once that is done, alignment is easier
- Given set of facts and inference rules, outputs model that agrees with most facts

Facts

- Facts are of the form “IP 128.8.126.92, IP 128.8.129.7, delta = 1”
  - IP address pairs are 1 TTL apart
  - delta is change in size of RR field

Inference rules

- Inference rules dictate which implementations are possible, for each value of delta
- Aided by observed standard practices
  - No self-loops, interfaces at ends of a link share first 31 bits

Results

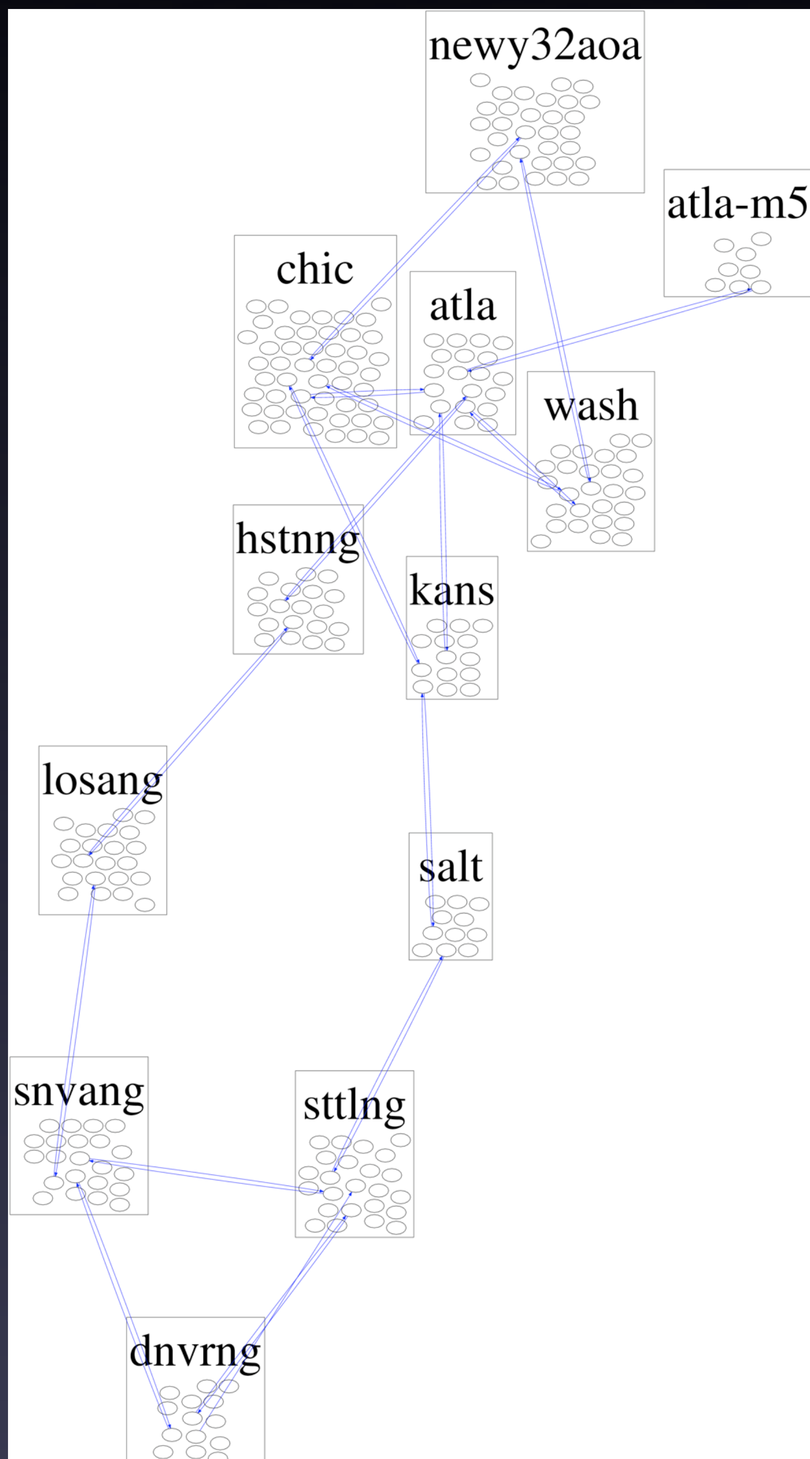
newy32aoa
atla-m5
chic
atla
wash
hstnng
kans
losang
salt
snvang
sttlng
dnvrng

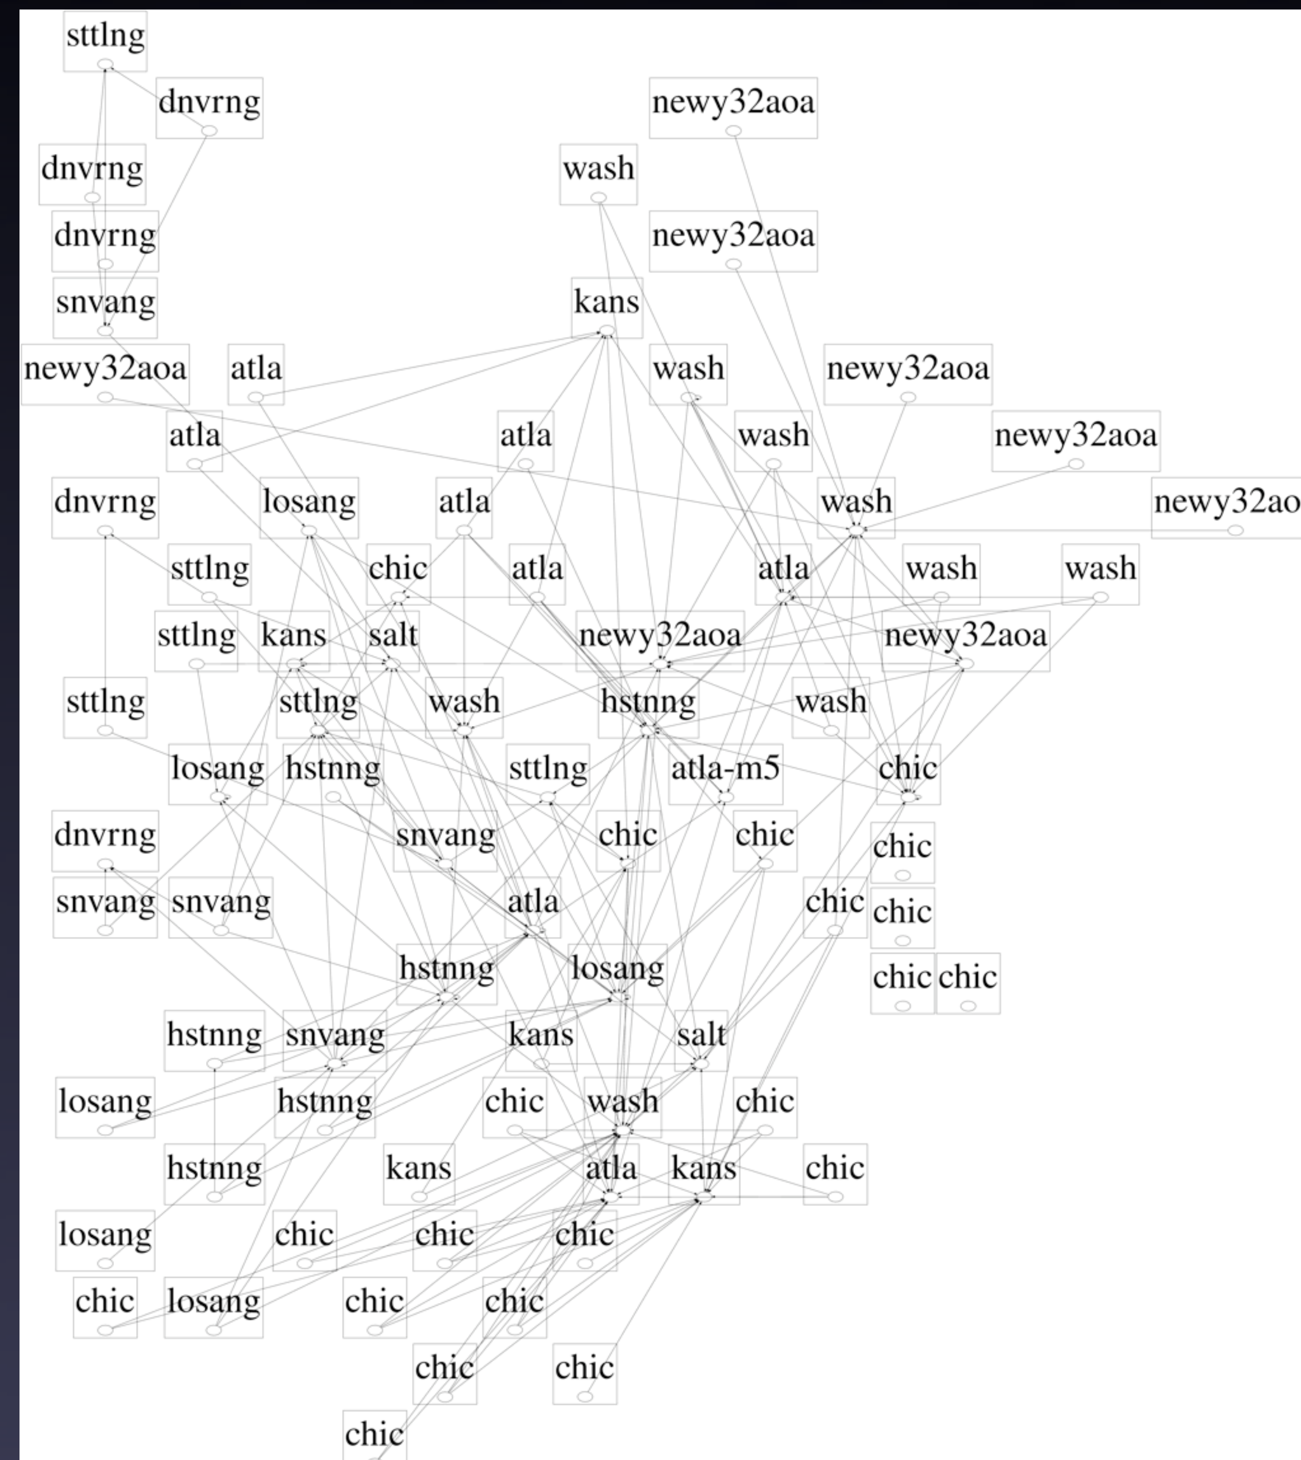
sttlng
dnvrng
newy32aoa
dnvrng
wash
dnvrng
newy32aoa
snvang
kans
newy32aoa
atla
wash
newy32aoa
atla
atla
wash
newy32aoa
dnvrng
losang
atla
wash
newy32aoa
sttlng
chic
atla
atla
wash
wash
sttlng
kans
salt
newy32aoa
newy32aoa
sttlng
sttlng
wash
hstnng
wash
losang
hstnng
sttlng
atla-m5
chic
dnvrng
snvang
chic
chic
chic
snvang
snvang
atla
chic
chic
hstnng
losang
chic
chic
hstnng
snvang
kans
salt
losang
hstnng
chic
wash
chic
hstnng
kans
atla
kans
chic
losang
chic
chic
chic
chic
losang
chic
chic
chic
chic
chic

Results

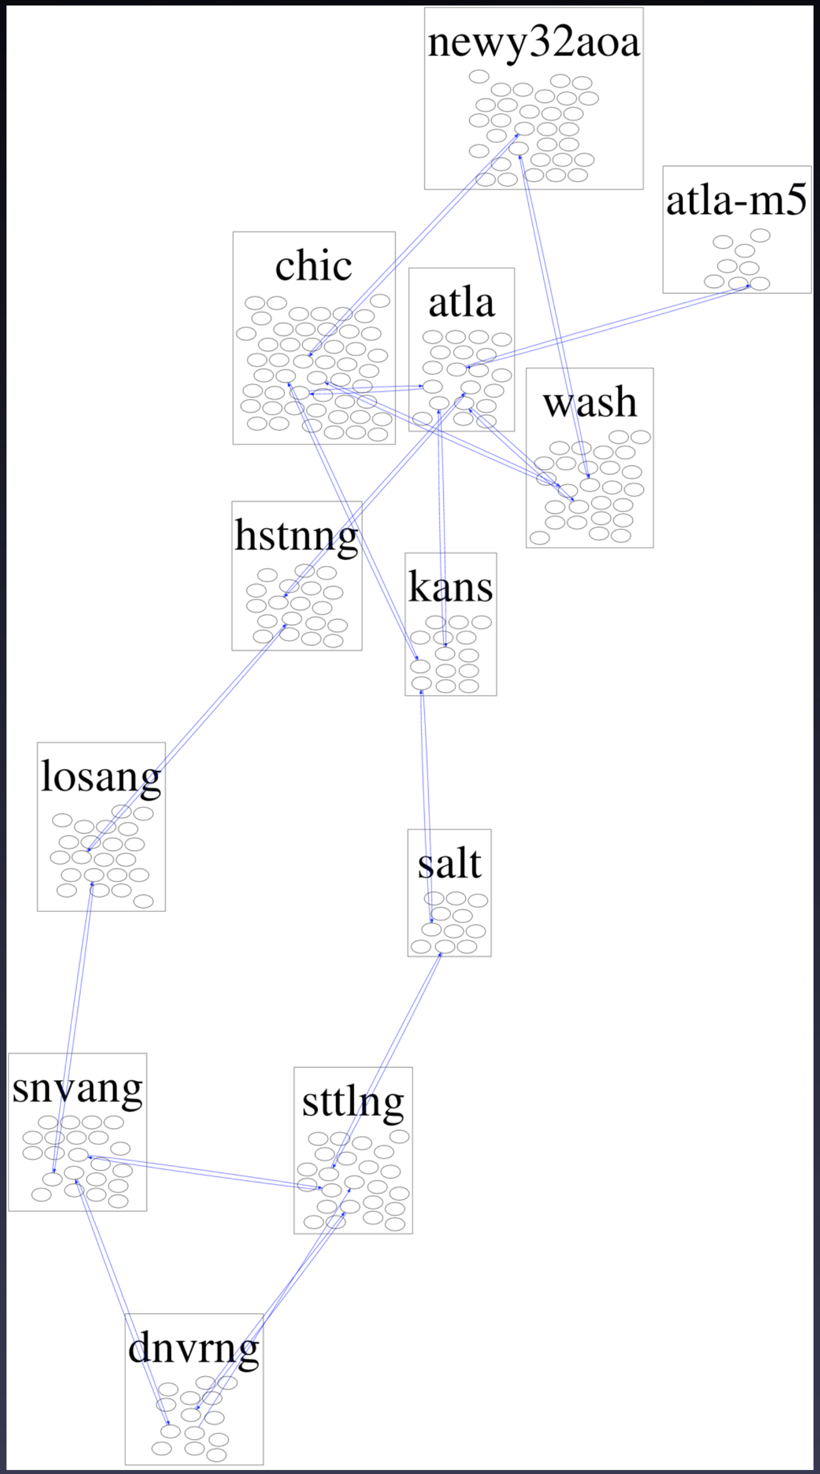
newy32aoa
atla-m5
chic
atla
wash
hstnng
kans
losang
salt
snvang
sttlng
dnvrng

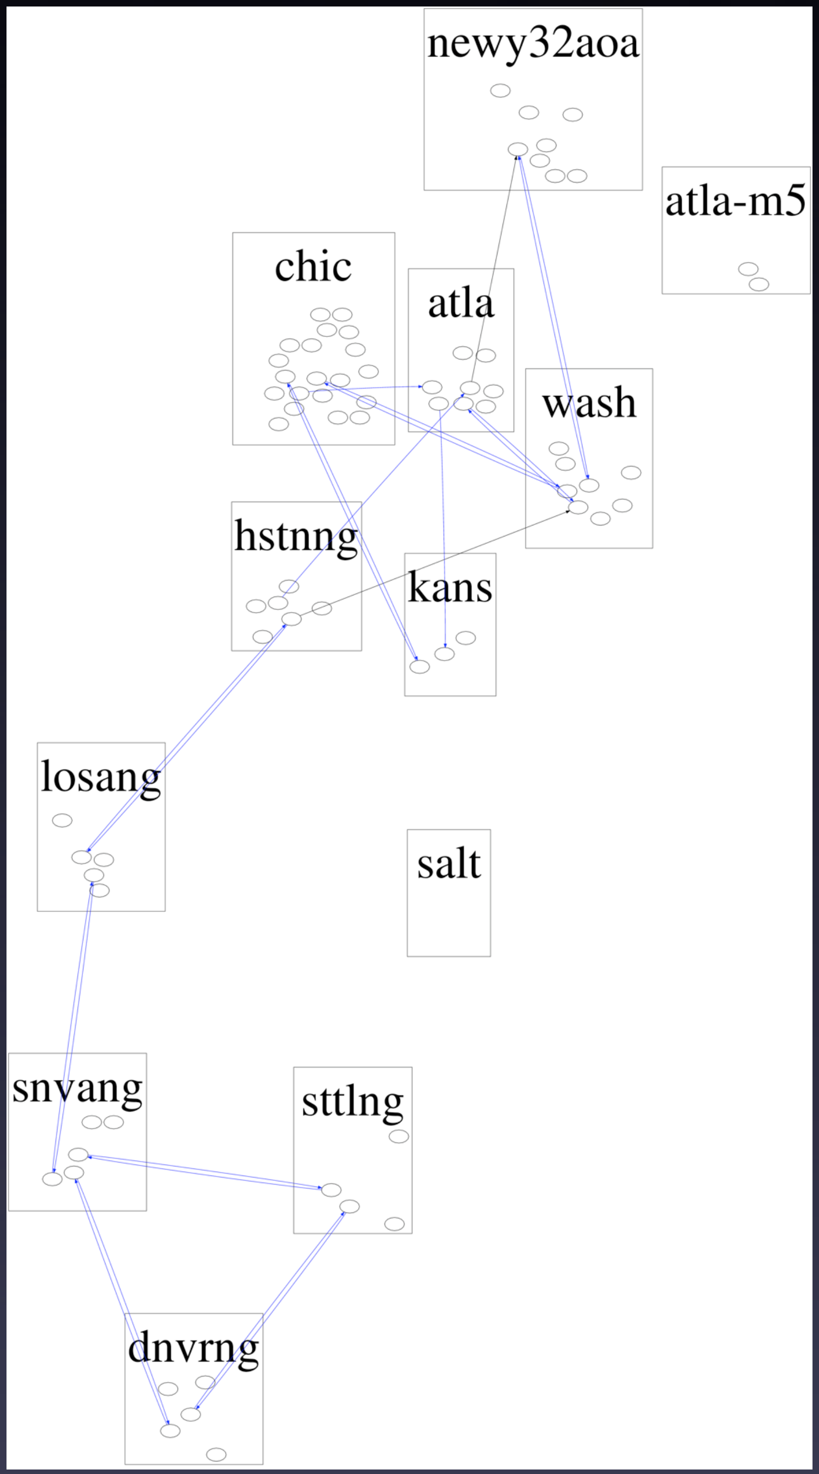
newy32aoa
atla-m5
chic
atla
wash
hstnng
kans
losang
salt
snvang
sttlng
dnvrng

Part 2:
Accurate and scalable alias resolution

Fixing Ally's Growing Pains with Velocity Modeling
A. Bender, R. Sherwood, N. Spring
IMC 2008

Alias resolution methods

- Source IP matching (Mercator)
- Common DNS naming practices
  - Does not require active probing
- Record route alignment
- Ally

IP header


Version
Length
TOS
Total length
IP identifier
Flags
Offset
TTL
Protocol
Checksum
Source address
Destination address
Options / Data...

IP header

<table>
<tr><td>Version</td><td>Length</td><td>TOS</td><td colspan="2">Total length</td></tr>
<tr><td colspan="3">IP identifier</td><td>Flags</td><td>Offset</td></tr>
<tr><td colspan="2">TTL</td><td>Protocol</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="5">Source address</td></tr>
<tr><td colspan="5">Destination address</td></tr>
<tr><td colspan="5">Options / Data...</td></tr>
</table>

The IP identifier

- Used for reassembling fragmented packets
  - Fragmented packets with same IP ID are reassembled at destination
  - Values must be unique
- Often implemented as a counter
  - Max. value is 2^{16}-1, then wraps to 0
- Some OSes (BSD) use random values

Ally

- Rocketfuel tool for alias resolution
- Sends probes to 2 addresses
- If IP IDs are “close”, re-probe after a slight delay
- If both pairs of probes are “close”, output alias
- Else output non-alias

Problems with Ally

- False negatives
  - For random and "fast" IP IDs
- False positives
  - If routers happen to have similar IP ID
- To resolve aliases in a list of *n* IP address, requires $O(n^2)$ probes
  - Each pair requires fresh probes
- Rate-limiting - incomplete results

Idea: model routers over time

- For routers that implement their IP ID as a counter, how fast does that counter increment? Can we compare counters off-line?
- Send series of probes over time, look at results

IP IDs can be modeled linearly

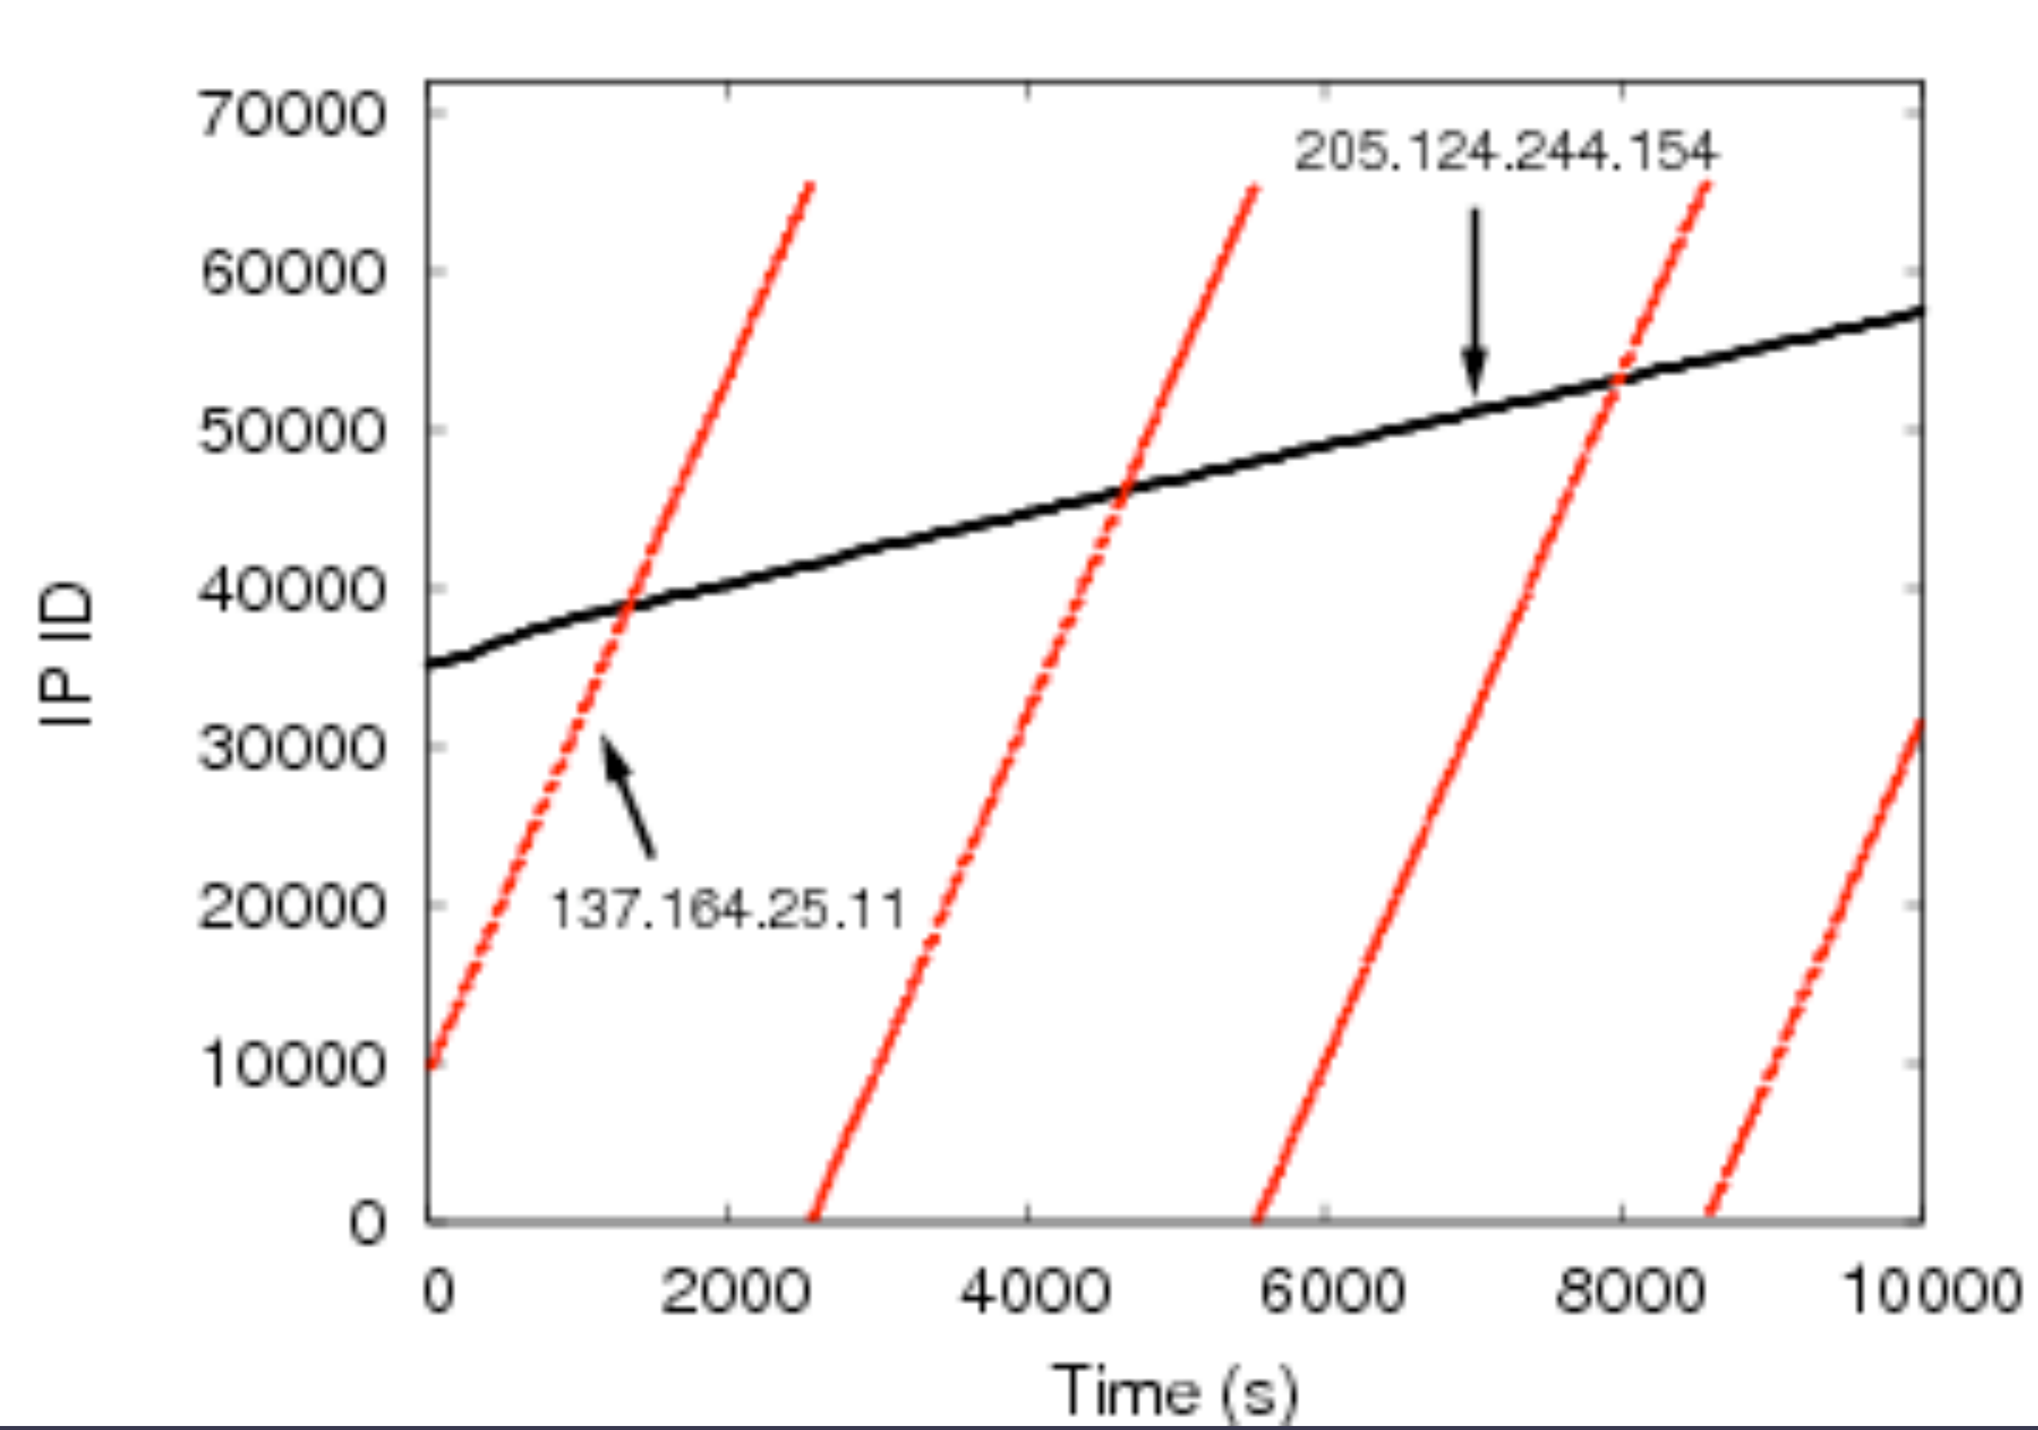

Data set

- 9,056 IP addresses
- Intra-PlanetLab routers, found with Discarte
- Sent 30 probes to each address, 34 seconds apart
- Goal is to model slope of each IP (if possible)

Problems with modeling

- 37% unresponsive to TCP and UDP
- “Uwrapping” (time, IP ID) samples to get linear sequence
- If probes are insufficiently frequent, hard to distinguish random IP IDs from fast IP IDs
- Unexpected observations: reflecting probe IP ID, IP ID always 0

Modeling slopes

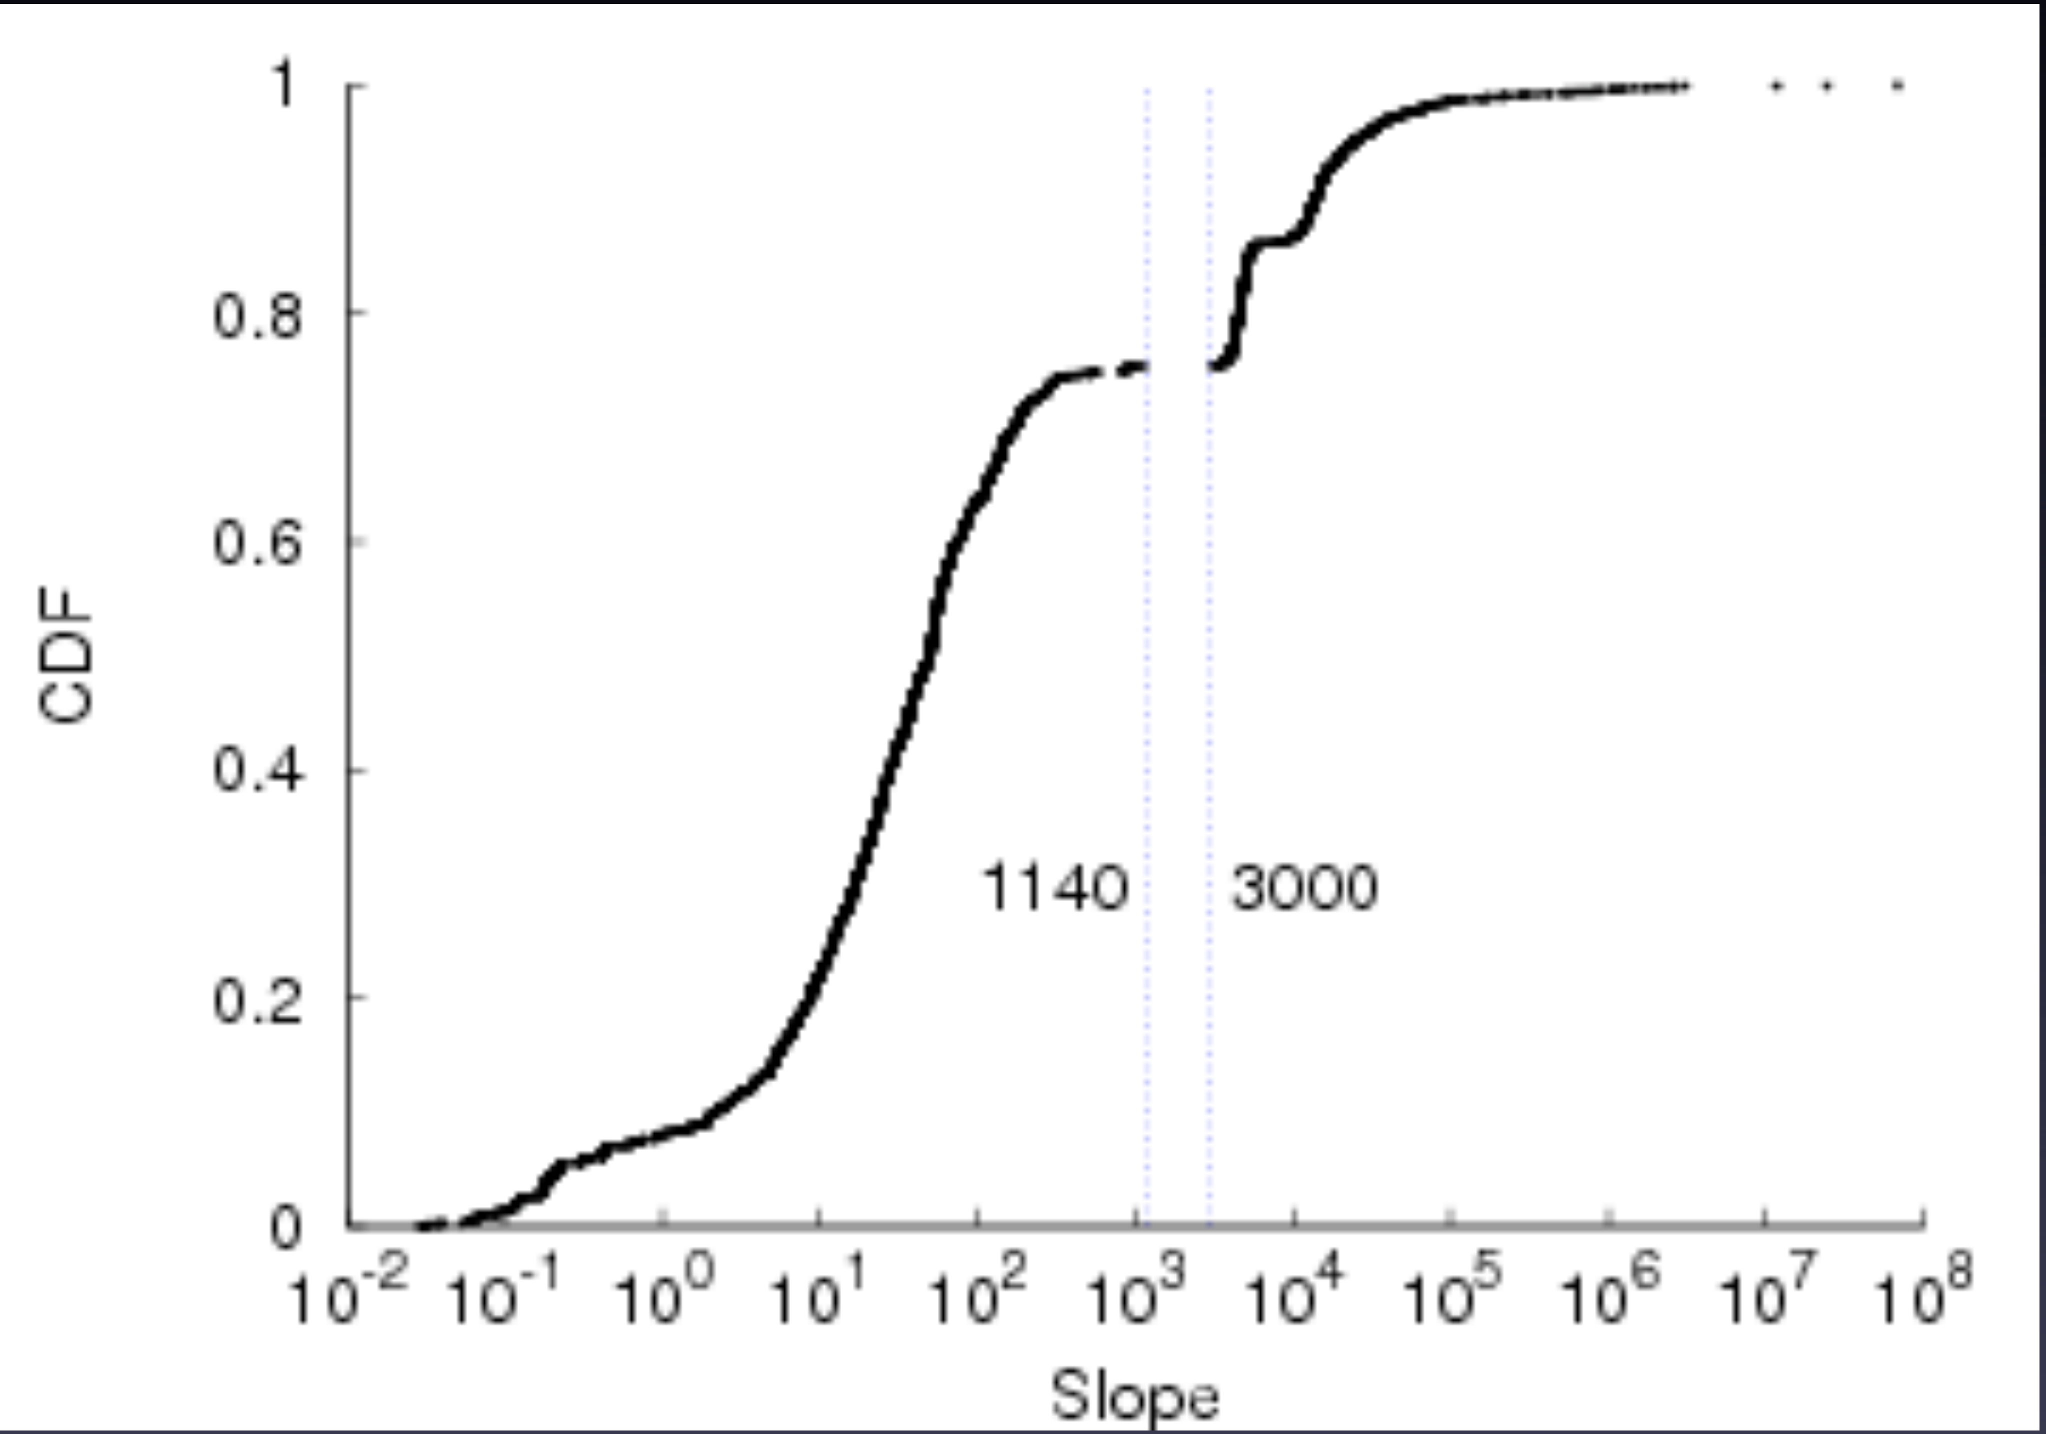

Inferring aliases

- Given two time series, determine if they are drawn from the same (linear) sample
  - If both are random, can't say anything
- For each sample at time *t*, estimate the IP ID of the other router at time *t*

Algorithm

- for (t, id) in samples$_A$
  1. find closest (t_1, id_1) and (t_2, id_2) from samples$_B$ such that $t_1 < t < t_2$
  2. est = $(id_2 - id_1) \dfrac{t-t_1}{t_2-t_1} + id_1$
  3. sum += |id-est|
- If no such t_1, t_2, estimate using inferred slope
- Divide sum by number of samples

Validation

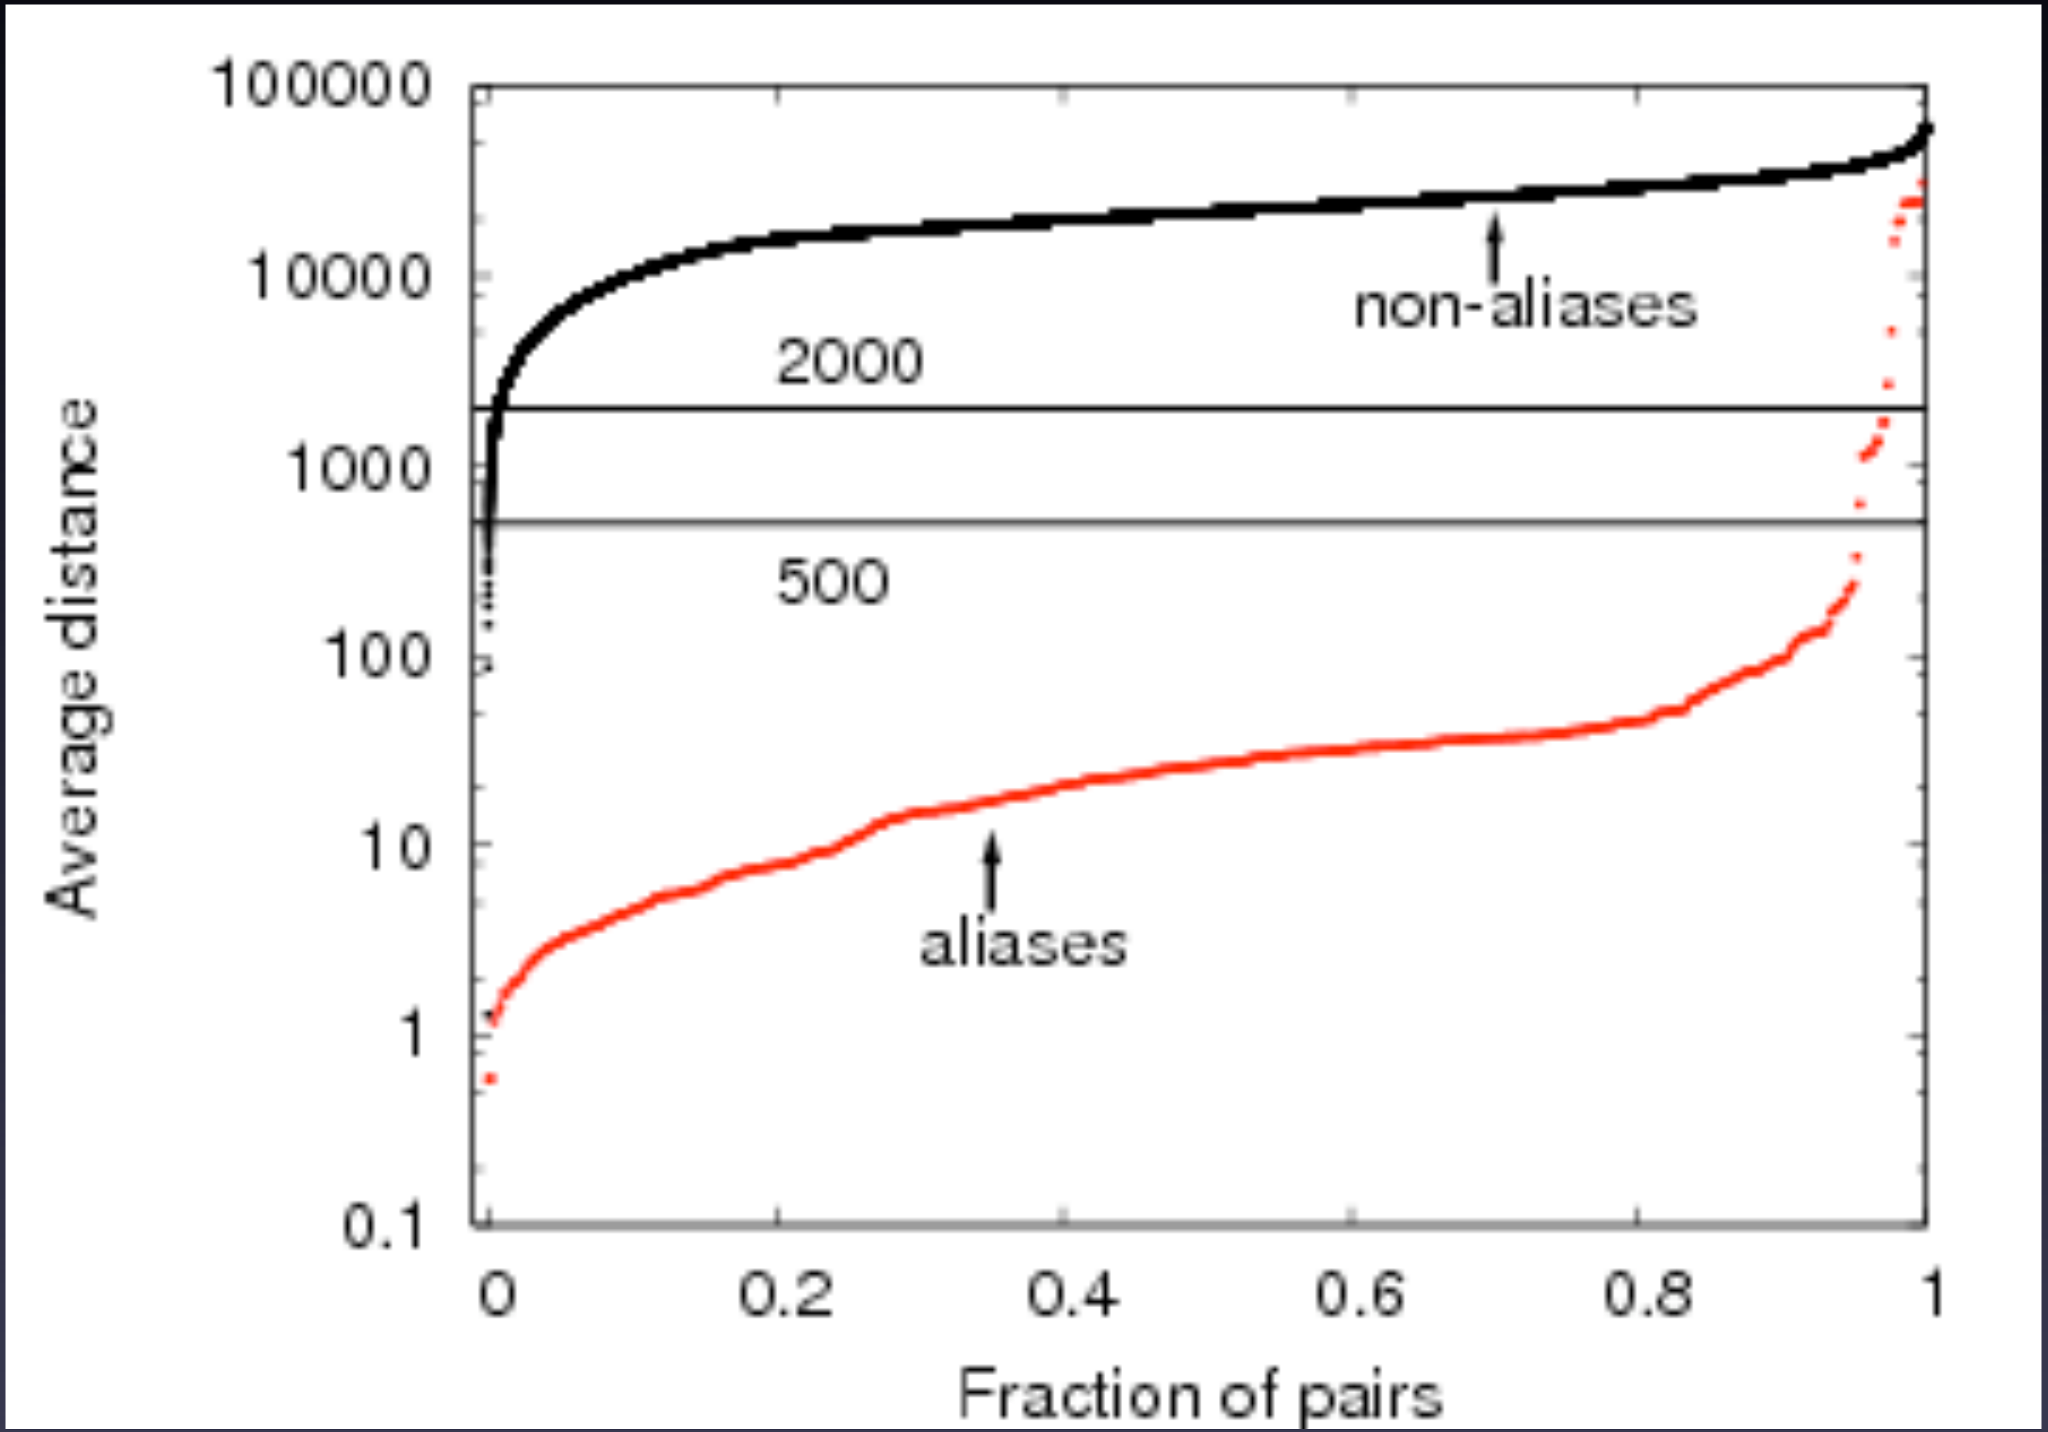
100000
10000
1000
100
10
1
0.1
Average distance
0
0.2
0.4
0.6
0.8
1
Fraction of pairs
non-aliases
2000
500
aliases

Scalability

- Given a list of *n* IPs, group into aliases
- Two rounds: probing and resolving
- Probing requires ~30 probes per IP
  - Ally requires O(*n*) probes per IP
- Resolving is done off-line
- Host with 10Mb/s connection can probe 500,000 hosts in 17 minutes

Routing update?

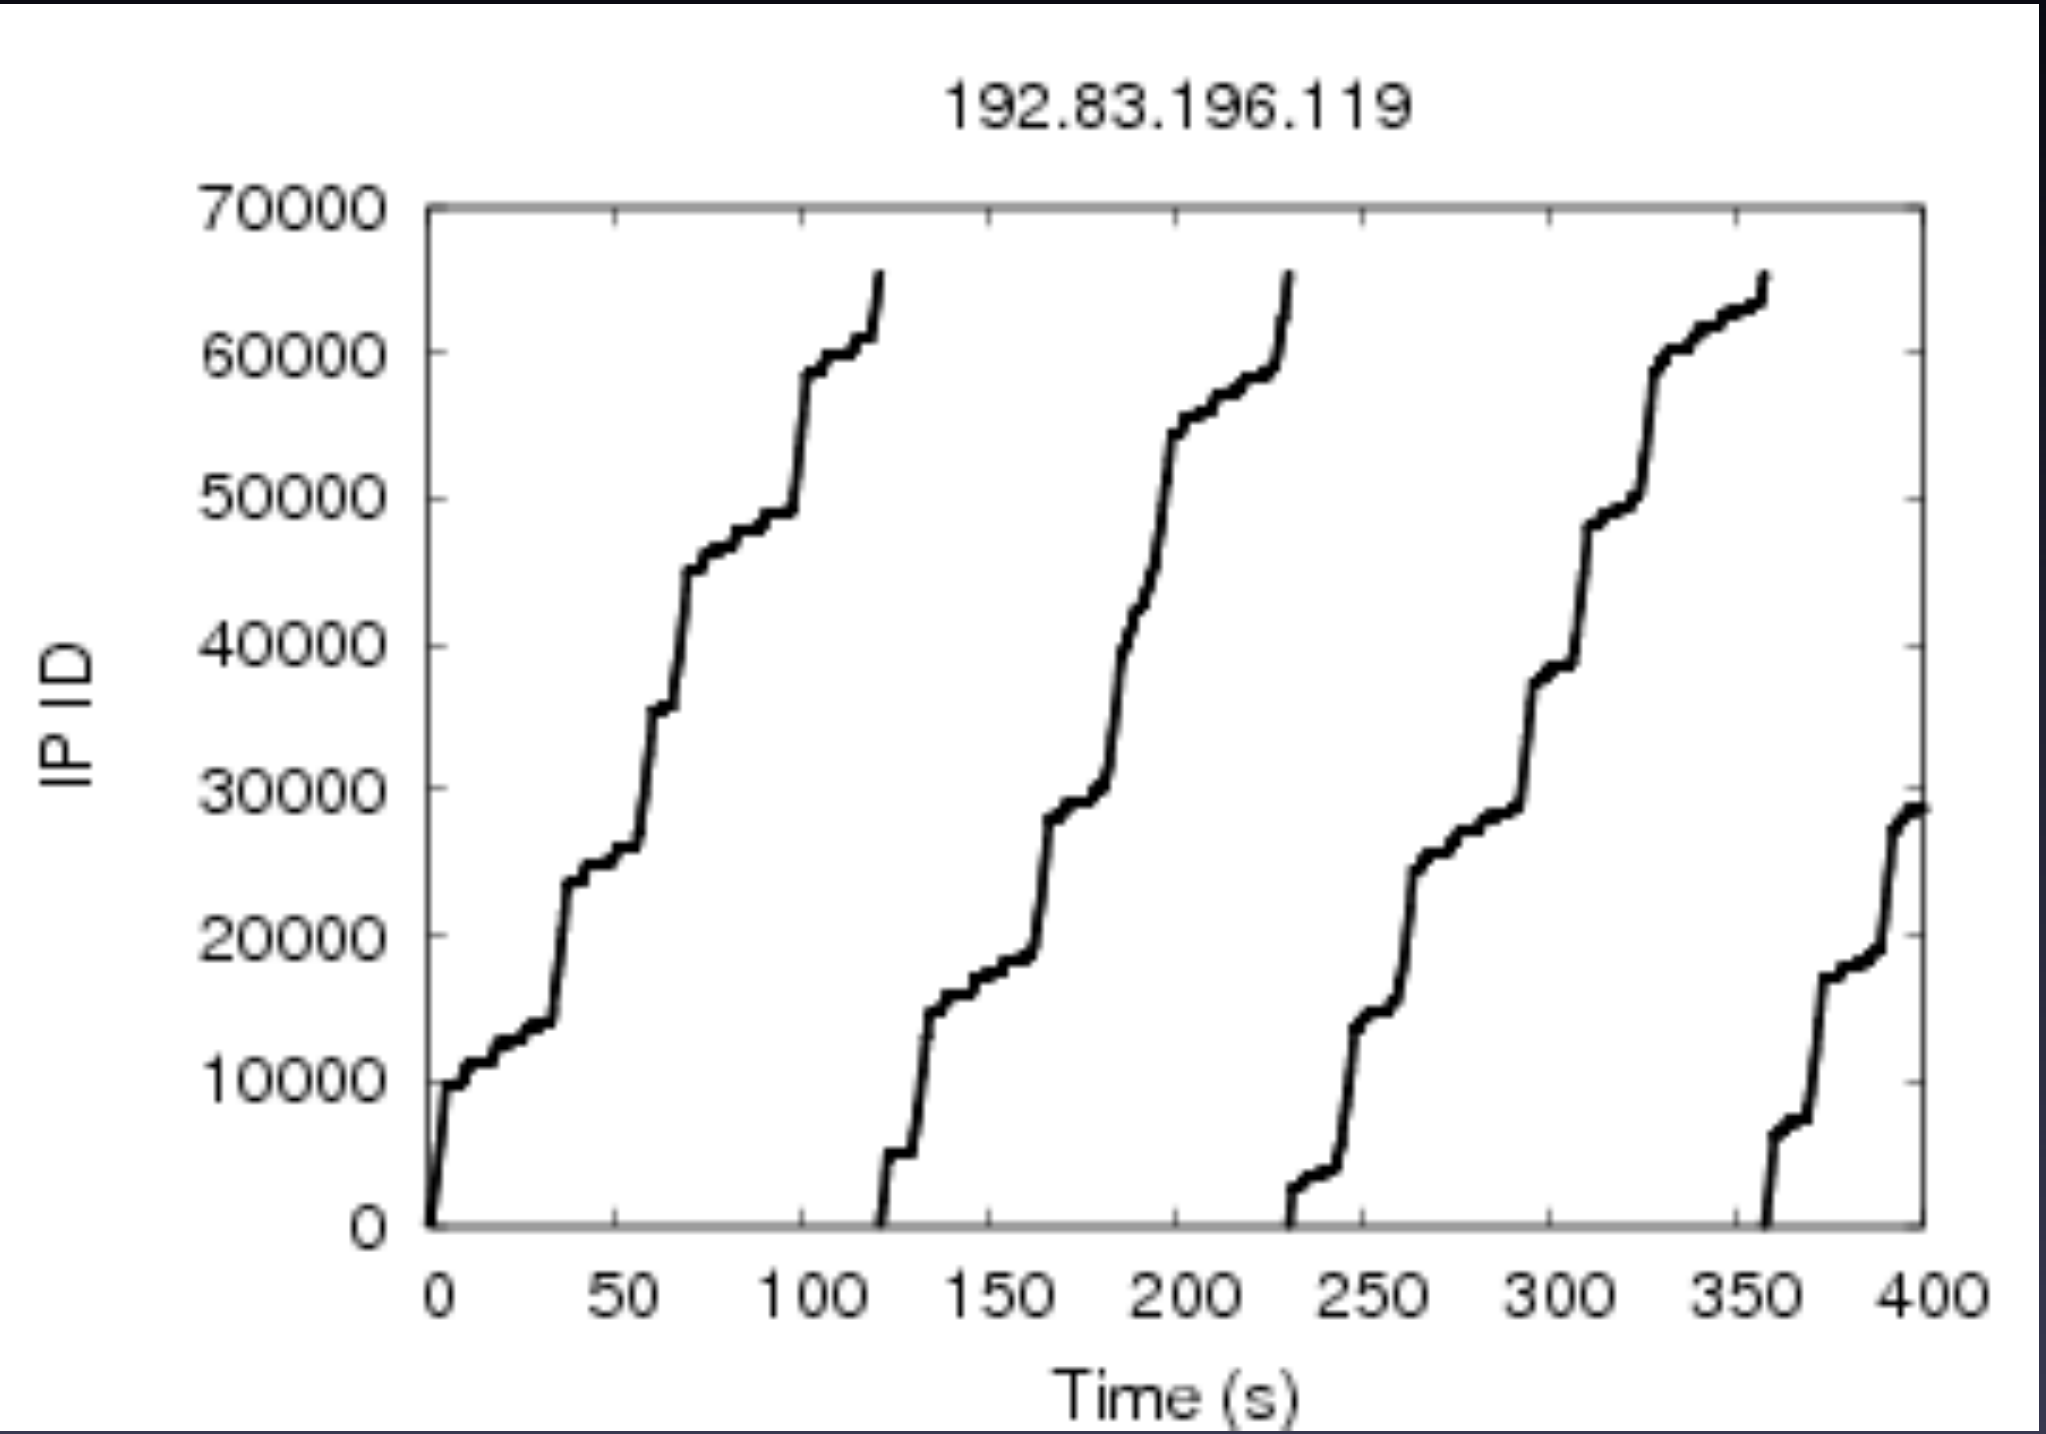

??

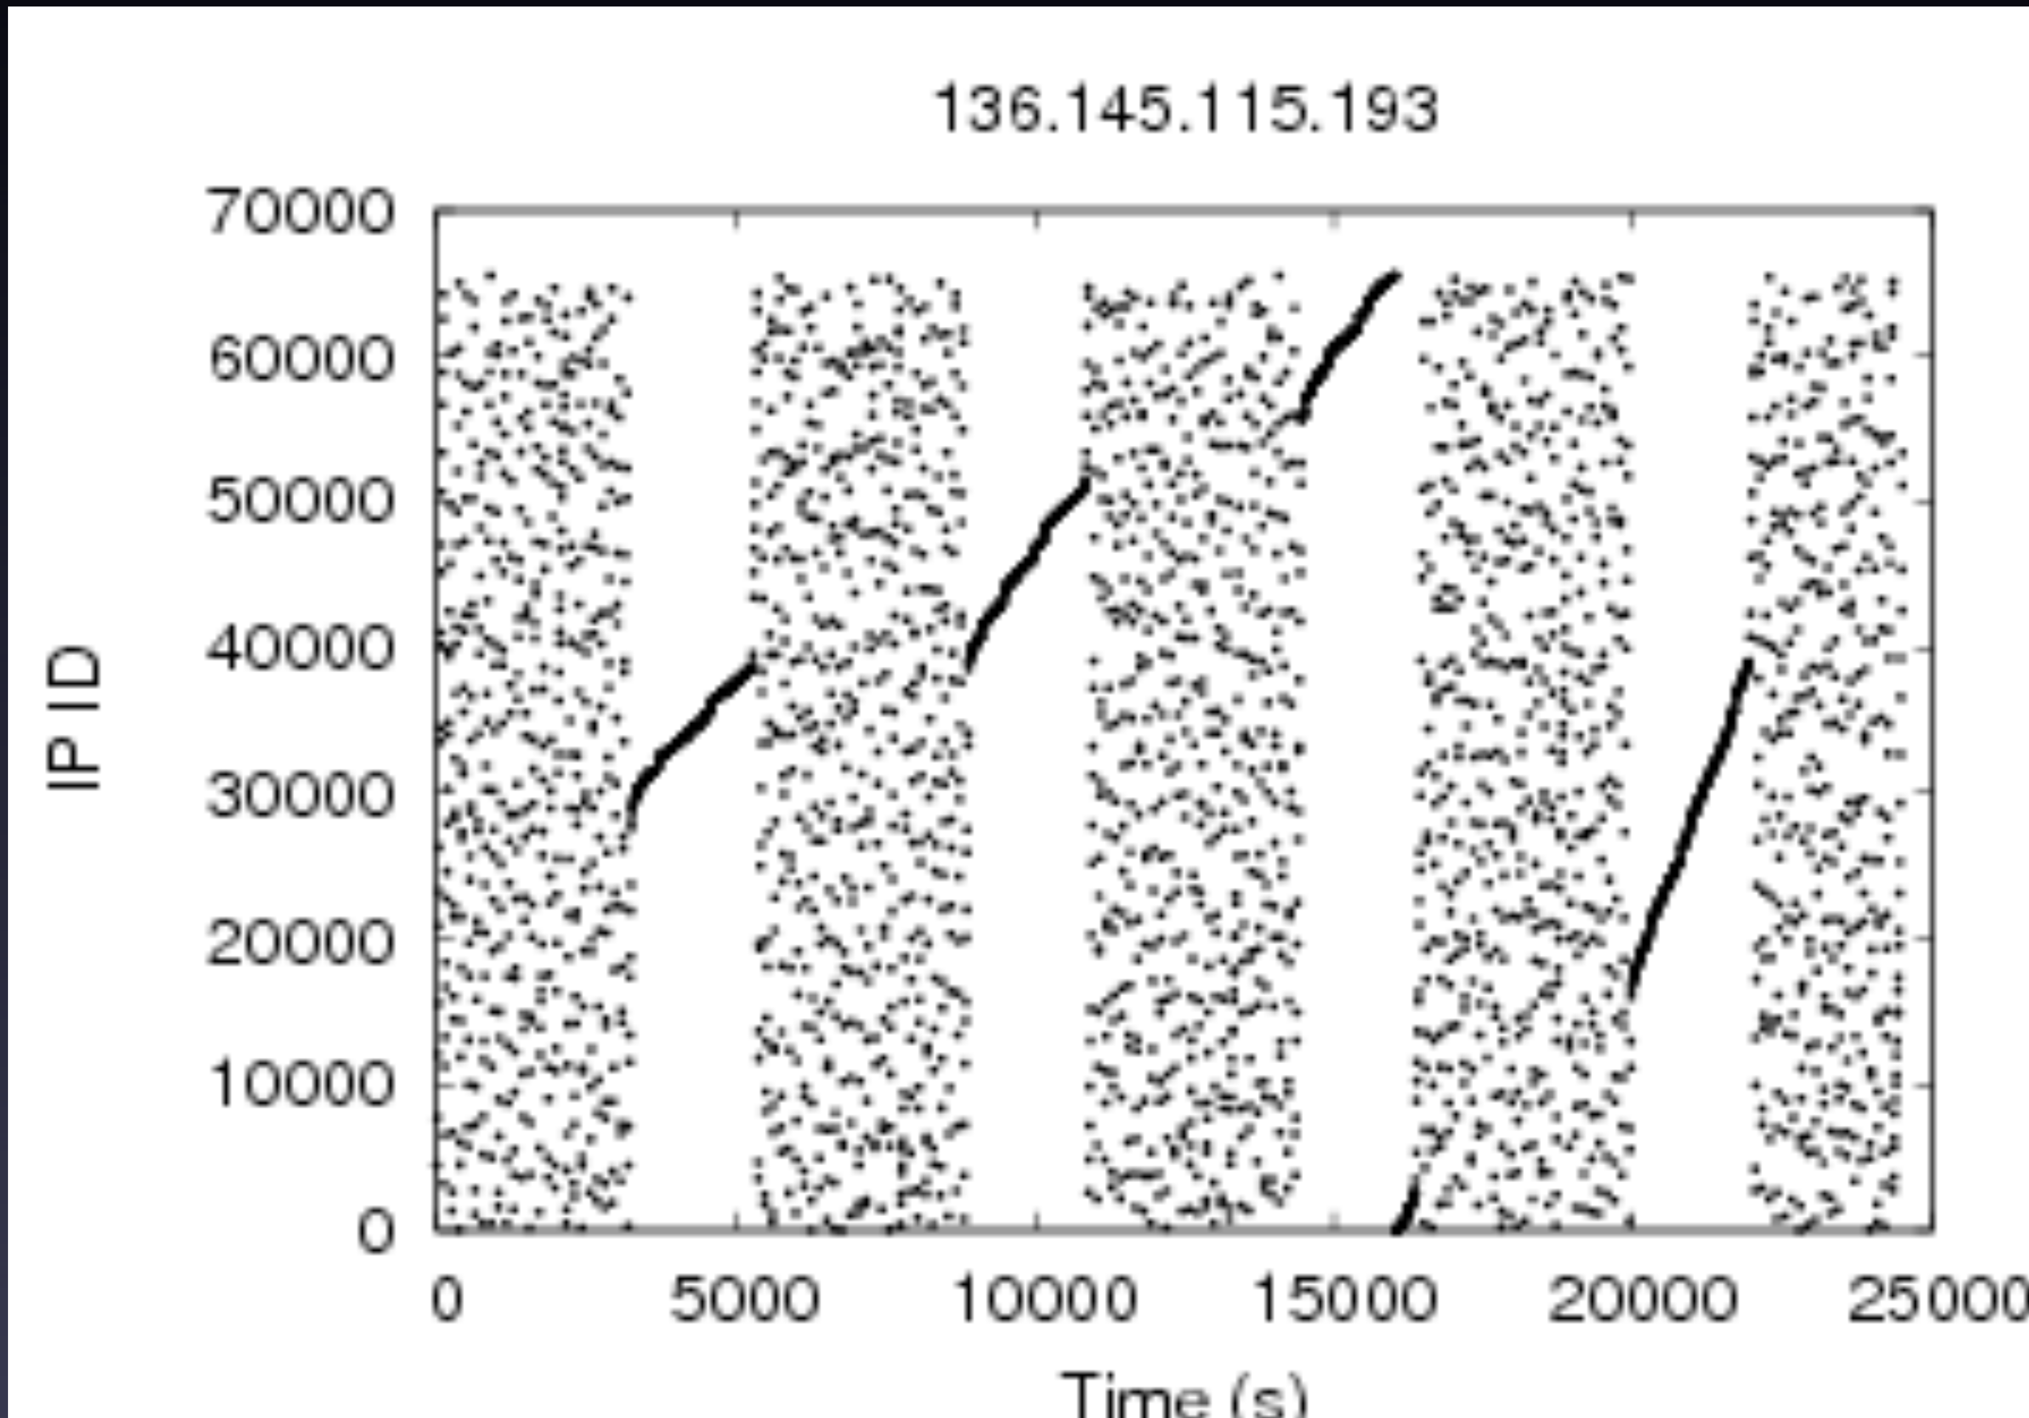

136.145.115.193
70000
60000
50000
40000
30000
20000
10000
0
IP ID
0
5000
10000
15000
20000
25000
Time (s)

Future Work

- Determine if TCP, UDP, and ICMP probes can be mixed
- Other uses of velocity modeling
  - Observing routing updates
    - Link and anomaly detection