Speedtrap: Internet-Scale IPv6 Alias Resolution

Matthew Luckie (CAIDA / UC San Diego)
Robert Beverly (NPS)
Billy Brinkmeyer (NPS)
kc claffy (CAIDA / UC San Diego)

Overview

- Alias resolution (IPv4 and IPv6)
- Speedtrap: IPv6
- Implementation and analysis
- Validation

Alias Resolution

- Router-level maps of the Internet are assembled using a set of hacks
- Traceroute returns interface IP addresses

- Which IP addresses belong to the same router?
  - Critical step: transforming an IP-interface graph to a router graph

Prior Work (IPv4)

- Analytical:
  - Graph Analysis (Rocketfuel, APAR, etc)
  - DNS (Rocketfuel)
- Fingerprinting
  - Common source address (Mercator)
  - Record Route (Discarte)
  - Pre-specified timestamps (Sherry IMC 2010)
  - IP-ID (Ally, Radargun, MIDAR)

IP-ID fingerprinting

- Ally (Spring et al., 2002)
  - Obtain sequence of IP-ID values from A and B which suggest a shared counter and therefore aliases

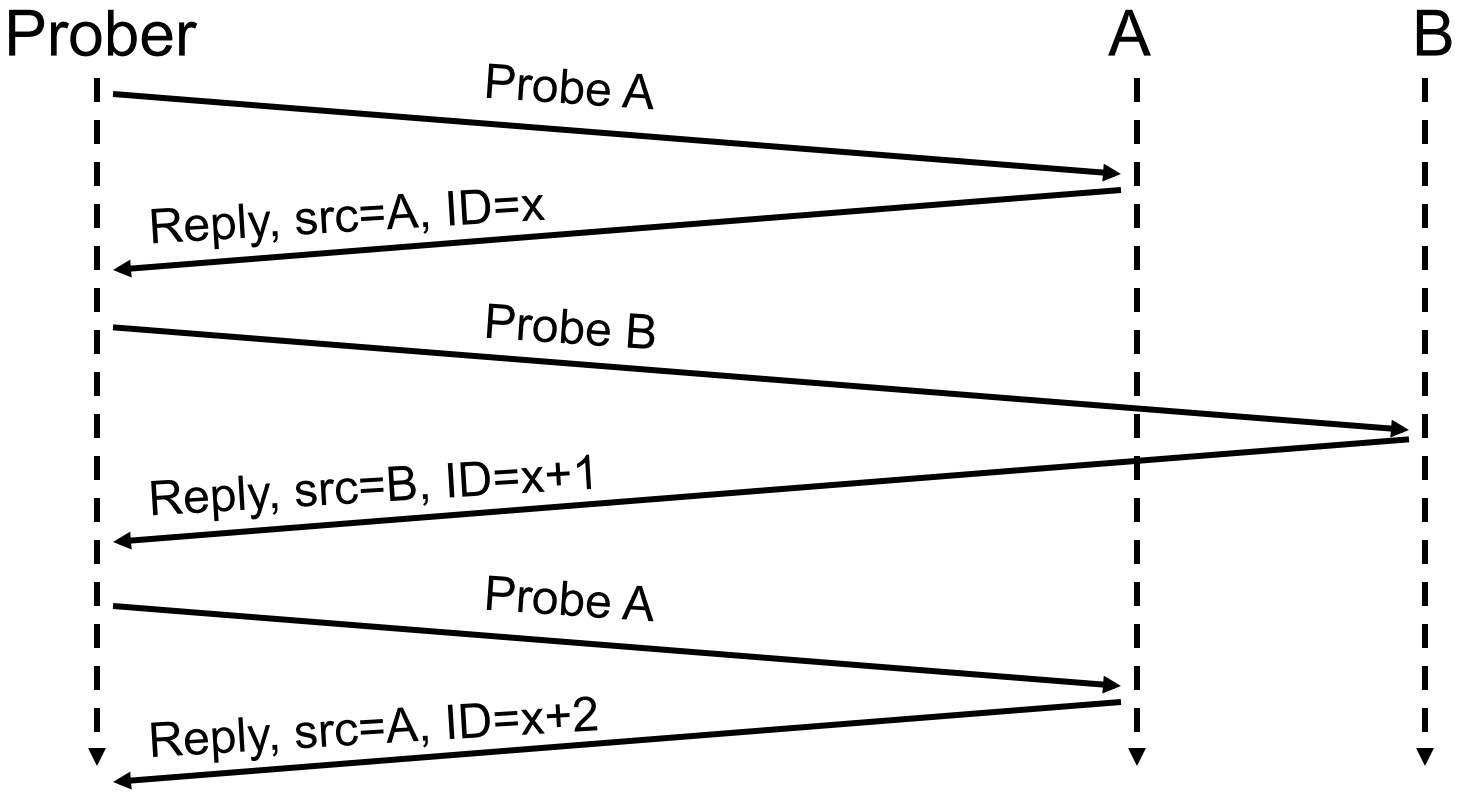

IP-ID fingerprinting: MIDAR Monotonic Bounds Test (MBT)

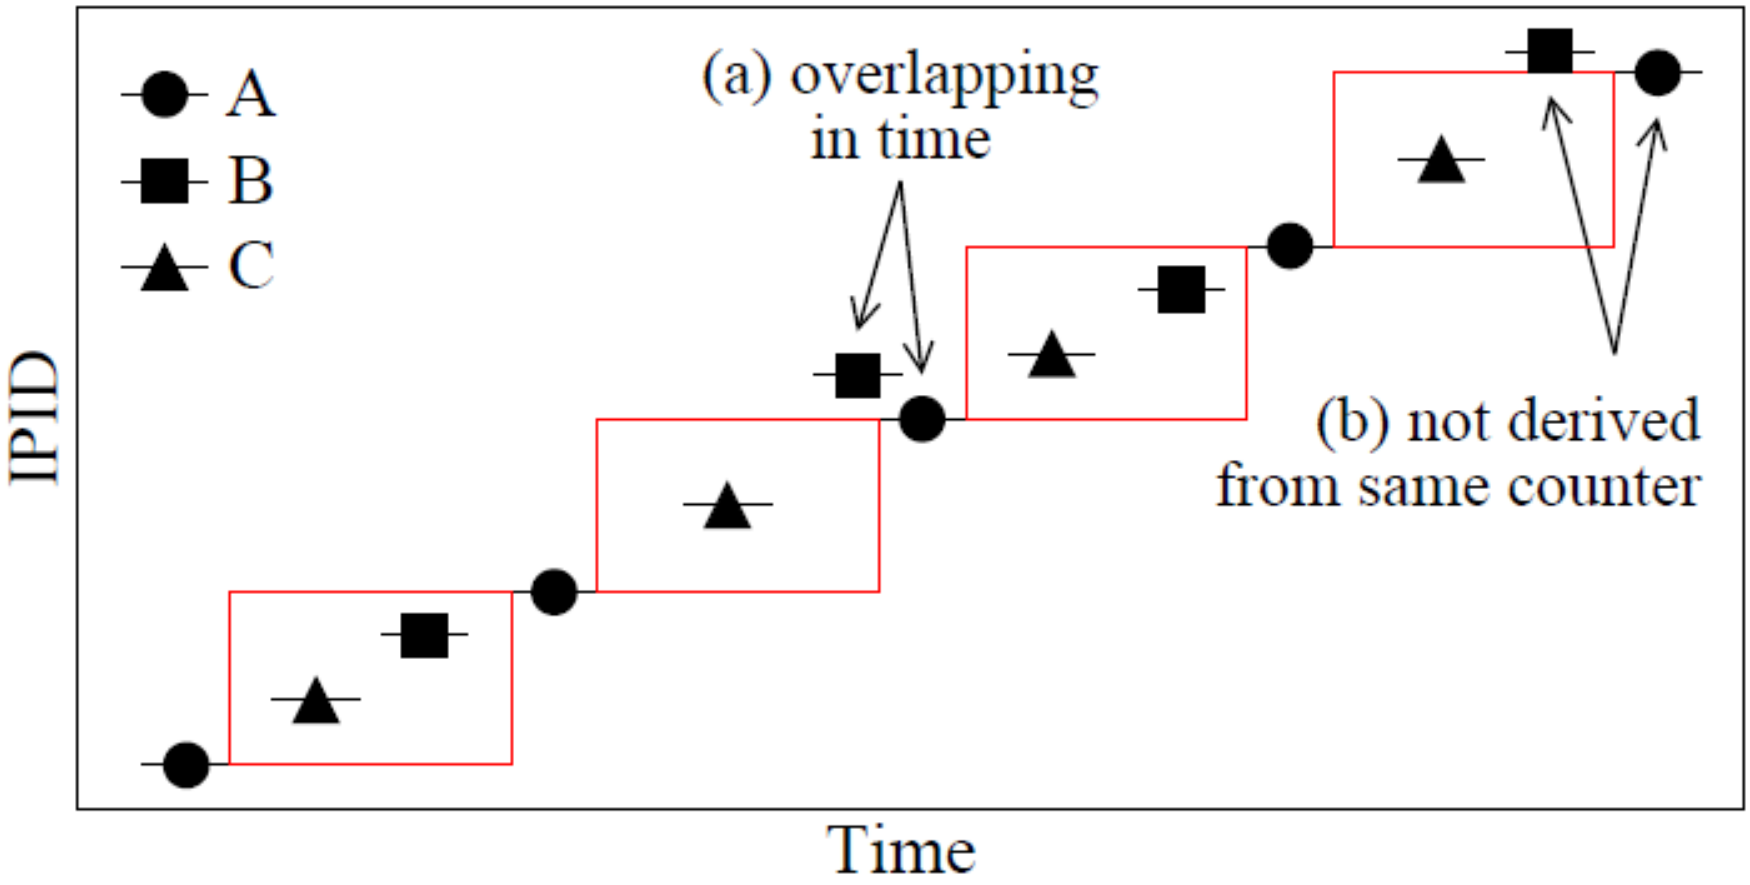

Prior Work (IPv6)

- Source routing
  - Waddington, et al. (2003): Atlas
  - Qian et al. (2010): Route positional method
- RFC 5095 (Dec 2007) deprecates source-routing functionality required
  - Denial of service through traffic amplification

- $O(N^2)$ probes, comparisons required

Too Big Trick (TBT)

- Induce router to send fragmented packets

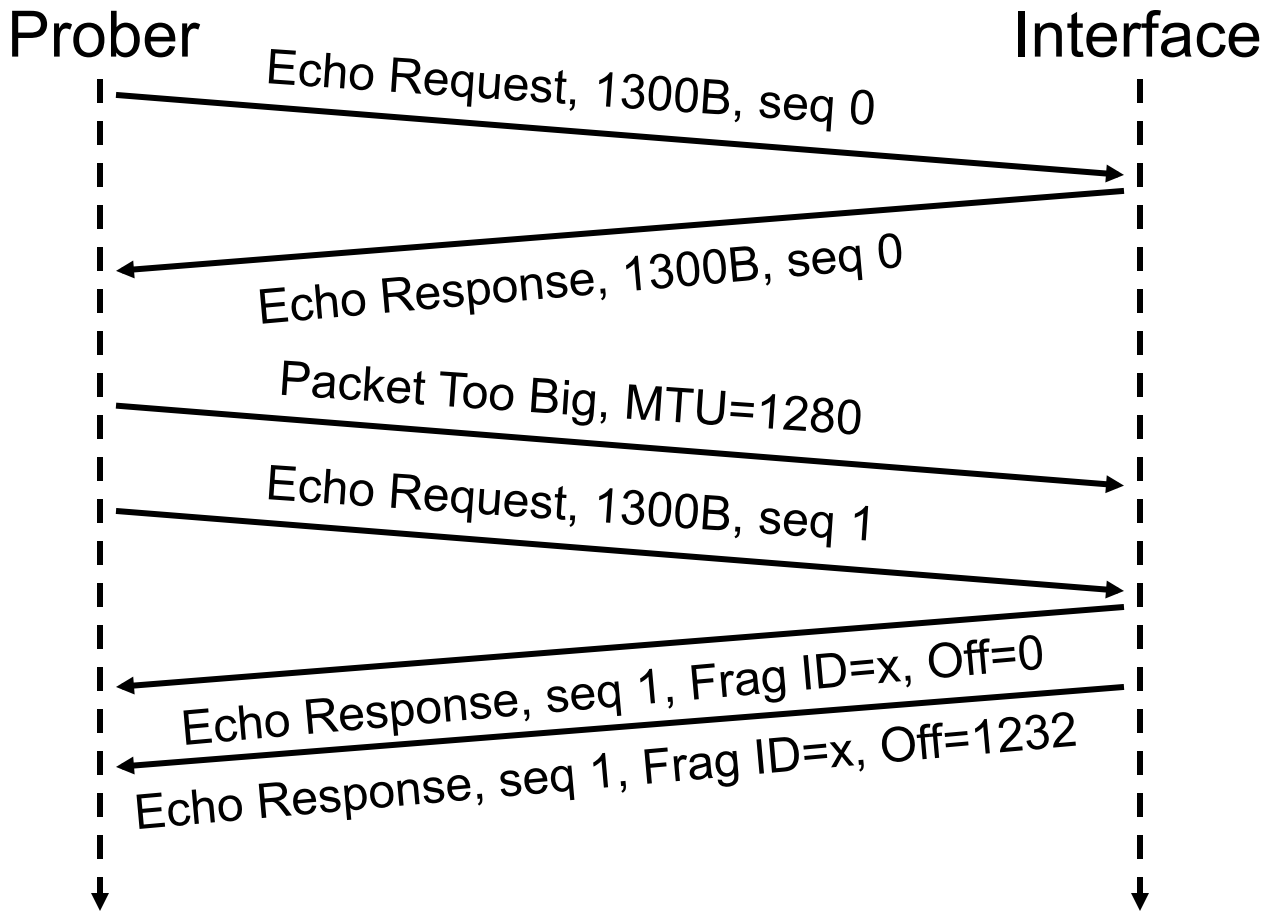

TBT Effectiveness

- March 2013 Interfaces observed by CAIDA's Archipelago
- 52,986 interfaces:
  - 32.1% sent fragments w/ incrementing IP-ID
  - 17.9% sent fragments w/ random IP-ID
  - 30.2% did not respond to echo probes
  - 19.8% did not send fragments

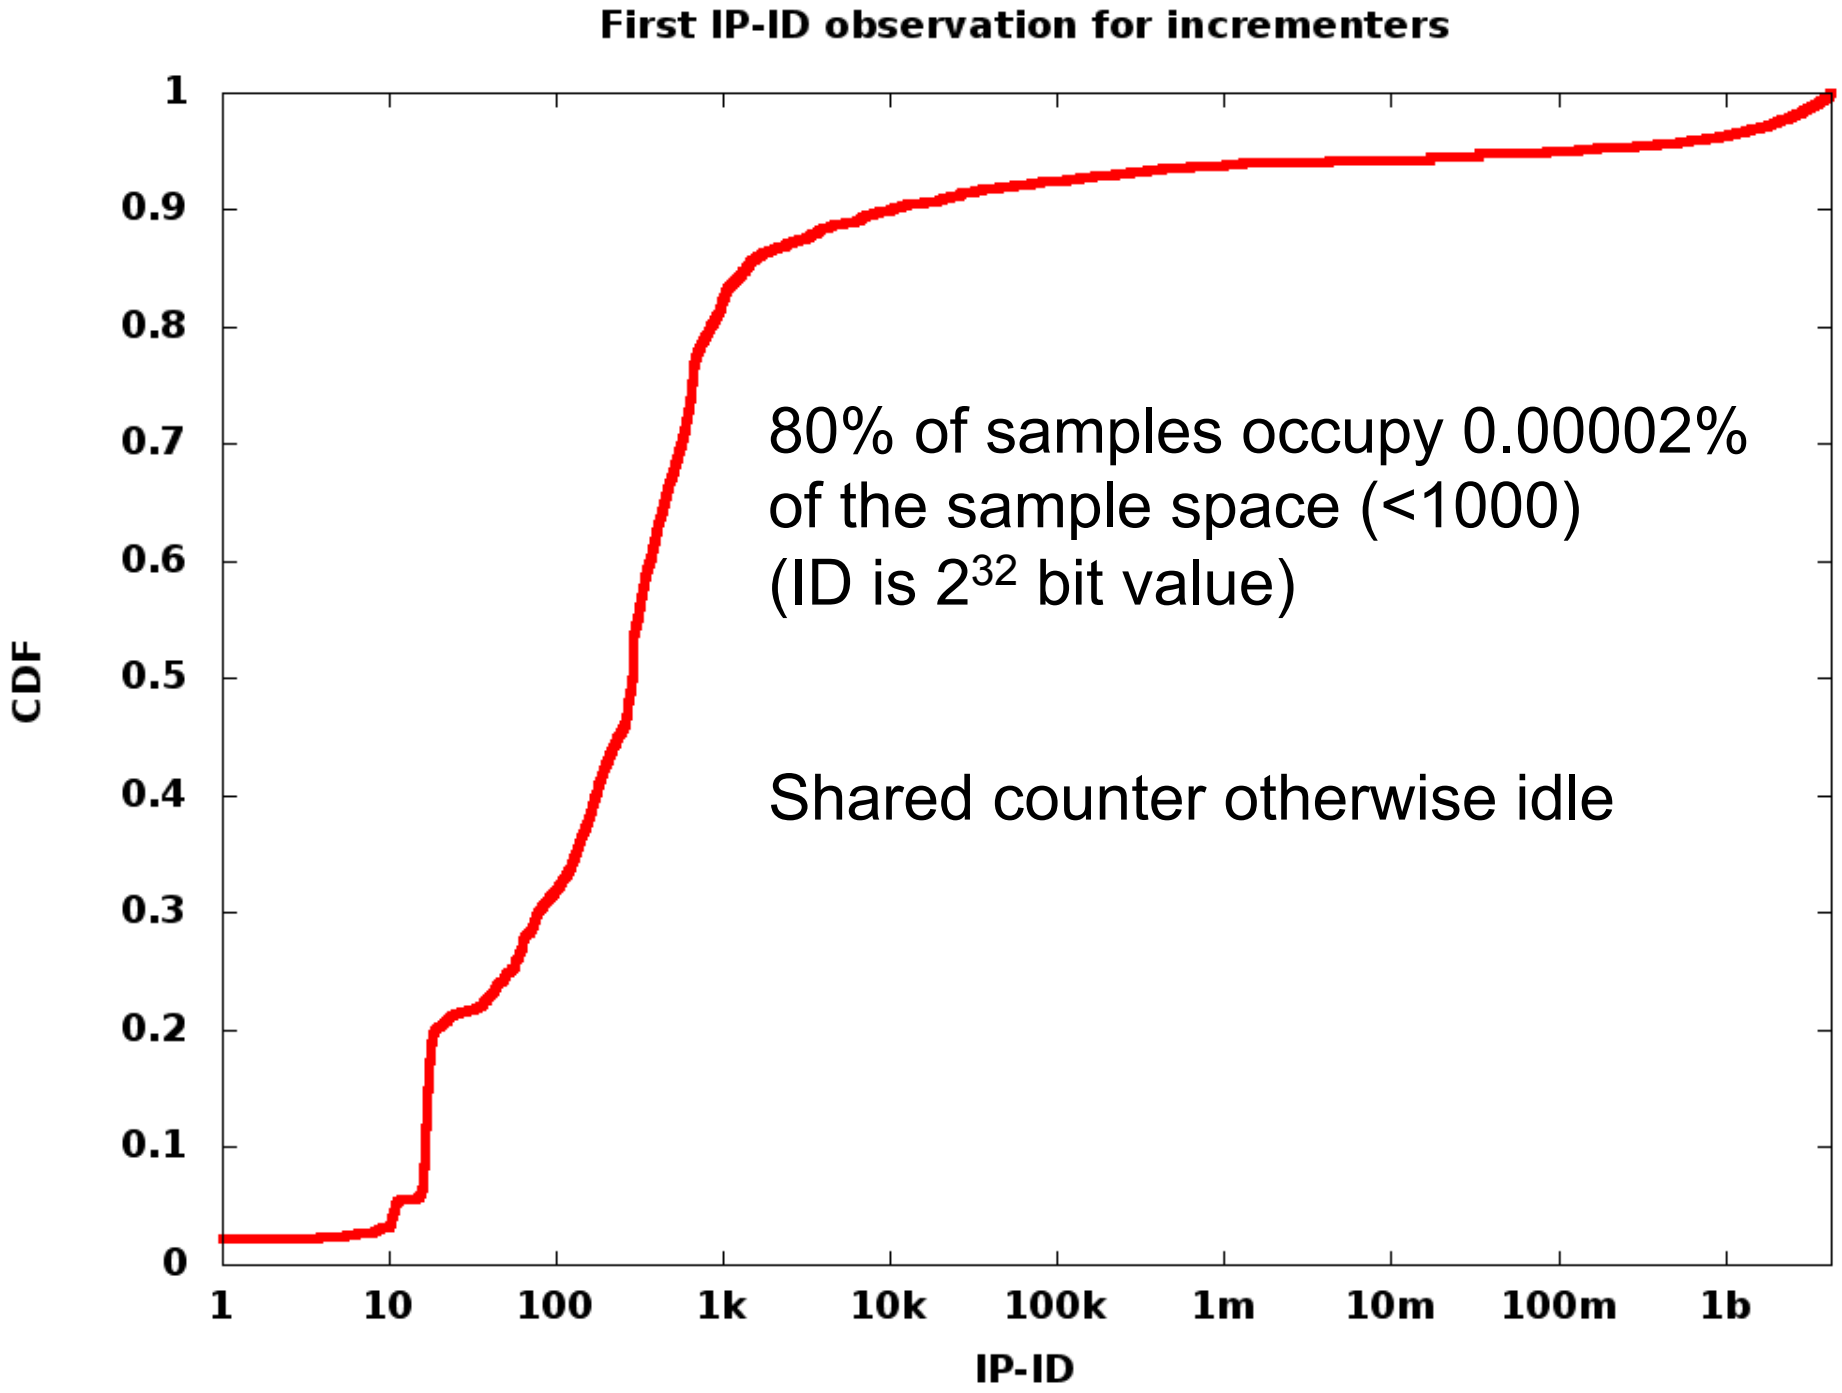
First IP-ID observation for incrementers
CDF
IP-ID
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
10
100
1k
10k
100k
1m
10m
100m
1b
80% of samples occupy 0.00002%
of the sample space (<1000)
(ID is 2^{32} bit value)
Shared counter otherwise idle

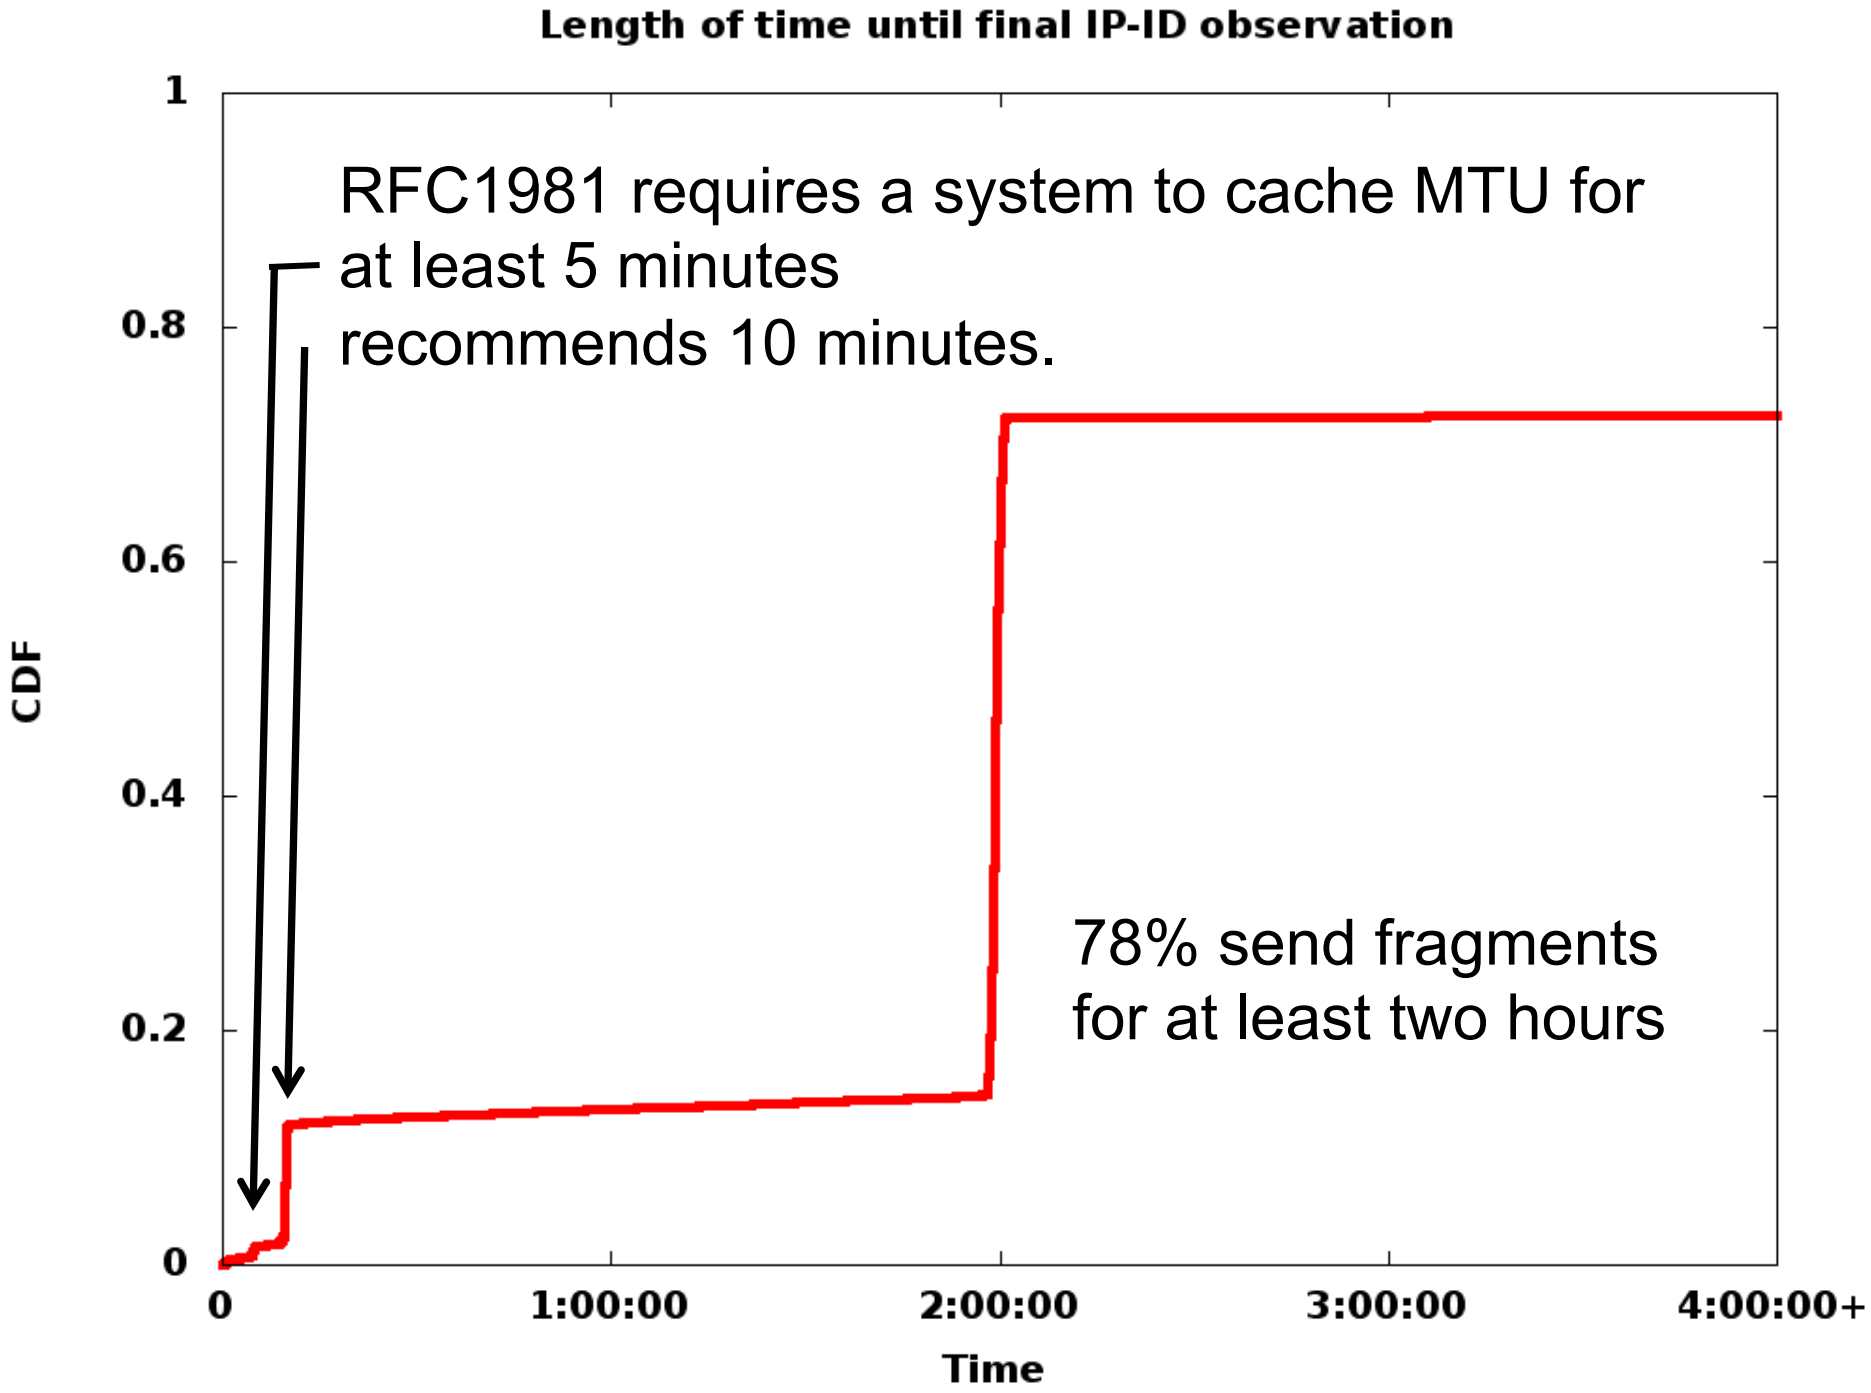
Length of time until final IP-ID observation
RFC1981 requires a system to cache MTU for at least 5 minutes
recommends 10 minutes.
78% send fragments for at least two hours
CDF
Time
1
0.8
0.6
0.4
0.2
0
0
1:00:00
2:00:00
3:00:00
4:00:00+

High-level Algorithm

1. Determine IPID behaviour of interfaces
2. Solicit non-overlapping sequence
3. Distill candidate routers
4. Pair-wise testing

Implementation

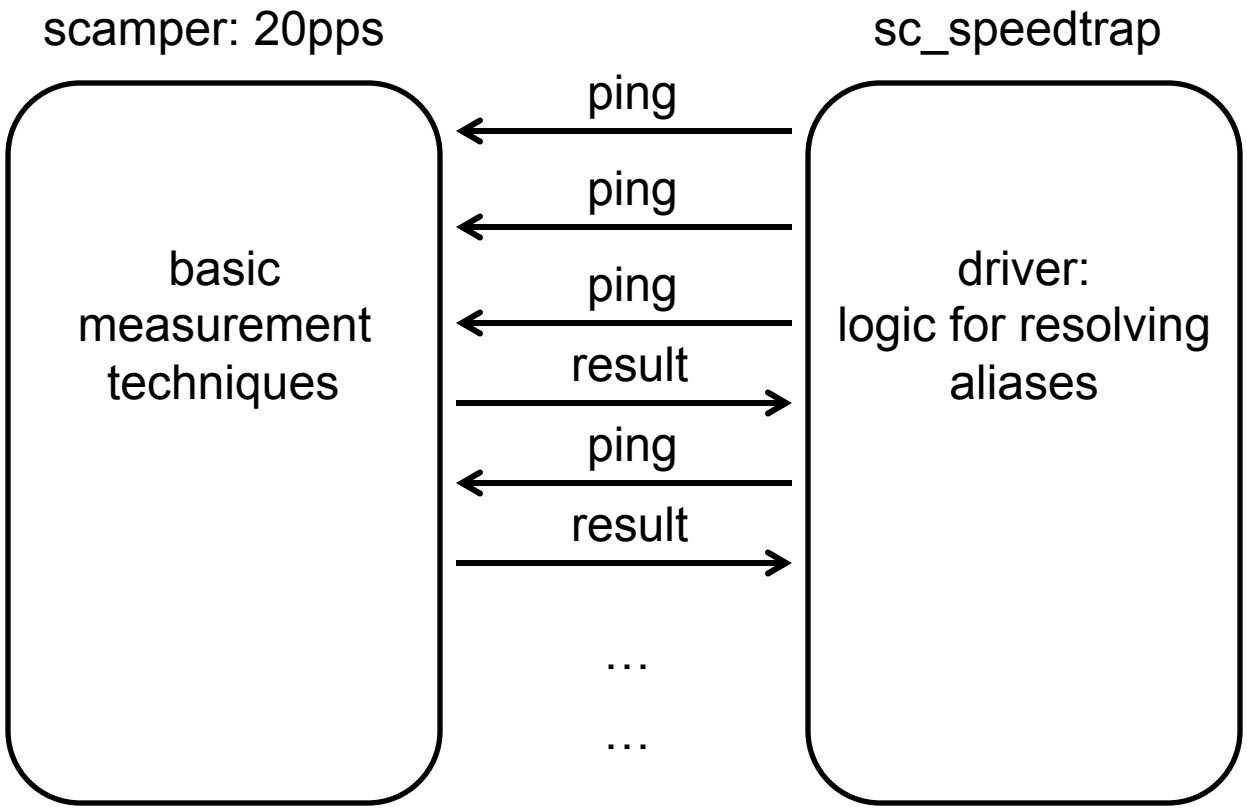

Step 2: non-overlapping sequence

- Stage 1: solicit fragmented response from all incrementers
- Stage 2: probe all interfaces which overlapped in time in stage 1
- Stage 3: solicit fragmented response from all incrementers

Step 2 stage 2: groups

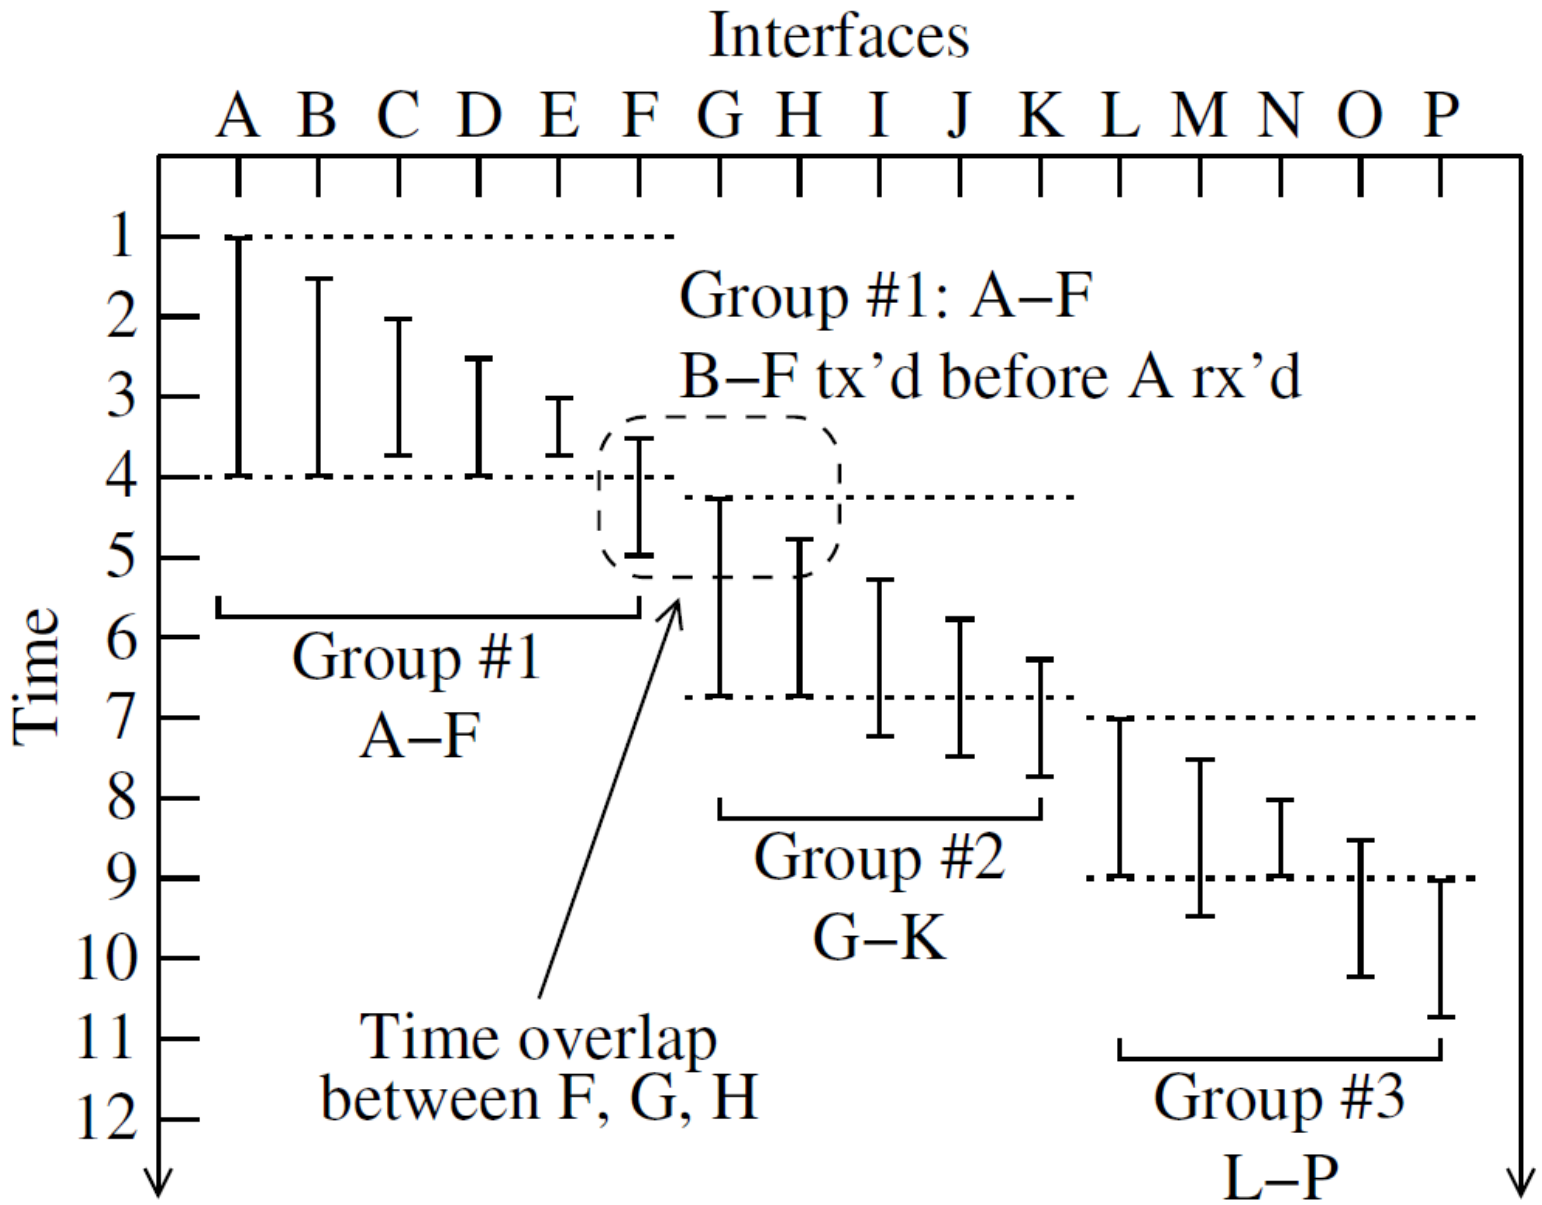

Step 3: distill candidate routers

- Infer candidate routers using a transitive closure (TC) of all interface-pairs that may share counter
- Try and cause divergence between A, B:
  - Probe A, B, A
  - One second apart
- TC groups can be probed in parallel

Step 4: pairwise testing

- Final pair-wise testing of candidate interfaces.
- In our data, 11,083 of 11,086 interfaces returned sequential IDs

Analysis: Packets and Time

Step		Packets	Time
1	IPID behaviour	317,814	5:35:44
2	Non-overlapping sequence	80,017	1:15:31
3	Distill candidate routers	34,659	1:15:43
4	Pair-wise testing	63,765	1:01:12
Total:		496,255	9:08:10

52,969 interfaces
20pps

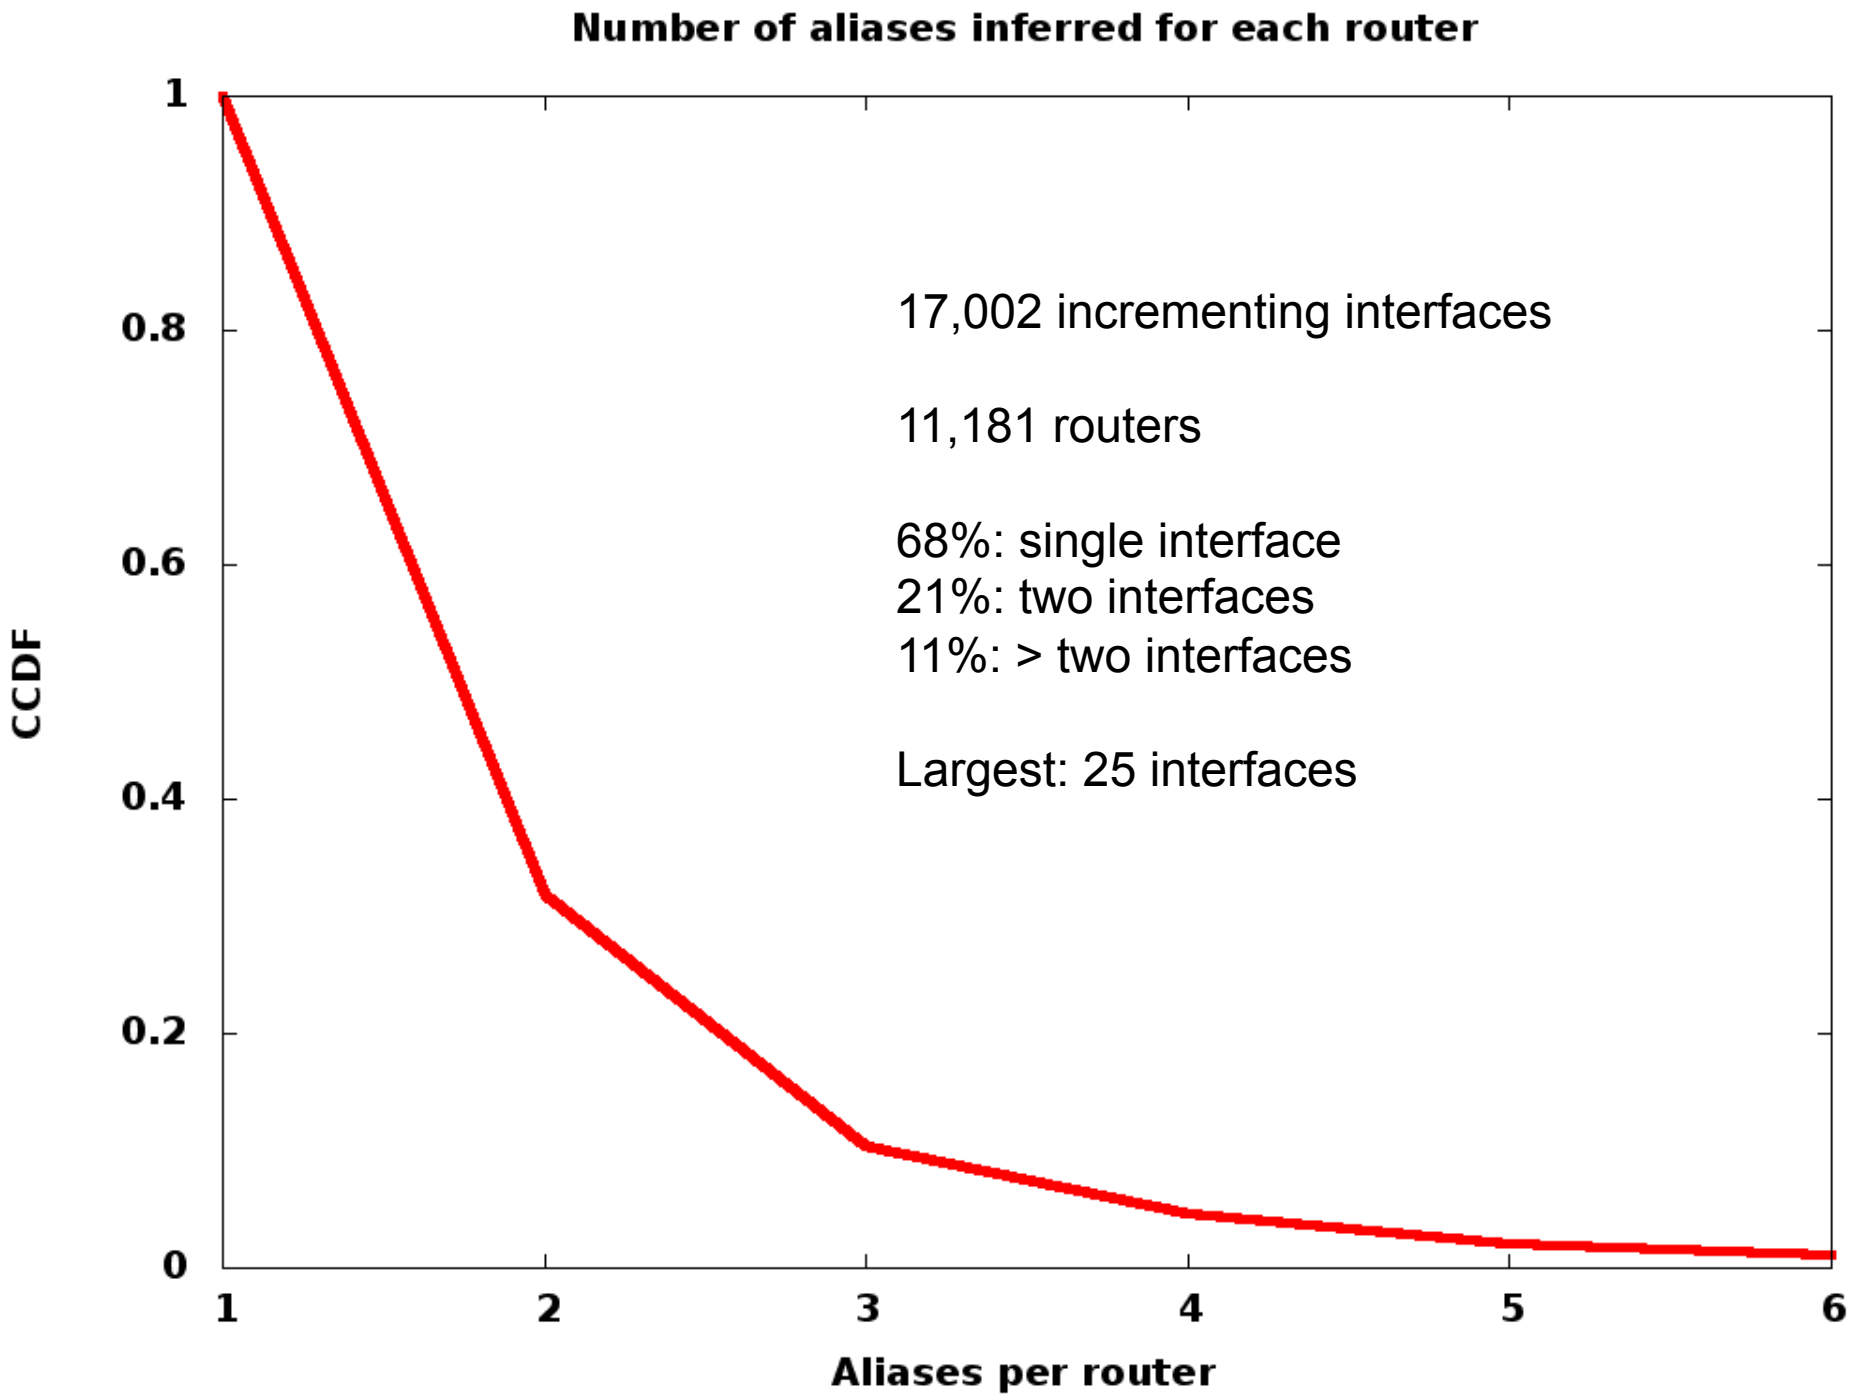
Number of aliases inferred for each router
CCDF
1
0.8
0.6
0.4
0.2
0
1
2
3
4
5
6
Aliases per router
17,002 incrementing interfaces
11,181 routers
68%: single interface
21%: two interfaces
11%: > two interfaces
Largest: 25 interfaces

Validation

- Four sources:
  - STP-1: RANCID — 70 routers
  - STP-2: DNS — 94 (observed)
  - AP: DNS — 267 (observed)
  - Tier-1: DNS — 239 (observed)
- Additional networks contacted, but DNS naming schemes were insufficient

Validation

Validation name	**STP-1**	**STP-2**	**AP**	**Tier1**
Data source	RANCID	DNS	DNS	DNS
Incrementing IPID	43	40	86	50
Random IPID		43	85	98
No Fragments		11	84	77
No Echo Replies			8	11
Mixed			4	3
Total Routers	70	94	267	239
Interfaces	151/750	85/279	138/1008	79/625
Correct assignments	150/151	85/85	137/138	79/79

451/453 correct assignments. 219 of 11,181 routers (2%)

Future Work

- Currently can resolve aliases for 1/3rd of routers
  - Investigate analytical approach to infer likely aliases
- Refine topology mapping process to focus on finding router subnets