DNS root traffic: how do anycast instances behave?

Nevil Brownlee, CAIDA

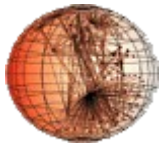

nevil@caida.org

Outline

- **DNS-OARC introduction**
- **Previous DNS-OARC data collection**
- **Analysis topics**
- **Some preliminary results**
- **Recommendations for future collections**

DNS-OARC

- **DNS Operations, Analysis and Research Centre**
- **Run by ISC, for membership info see https://oarc.isc.org/docs/dns-oarc-overview.html**
- **Collects trace files of DNS packets at root server instances**
- **CAIDA has been working on analysing that trace data ...**

Previous OARC data collection

- **Participants**
  - **C-root (cogent): 7 instances**
  - **E-root (nasa) : 4 instances**
  - **F-root (isc) : 71 instances**
  - **K-root (ripe) : 31 instances**
- **Dates (UTC, 48-hour period each)**
  - **2005: 09/22~23, 09/28~29**
  - **2006: 01/10~11 (most completed)**

Previous OARC data collection (2)

- **Disk usage (.gzip) – 320 GB altogether**

root server	total	per instance/day
C (cogent)	60GB	1 ~ 4.5GB
E (nasa)	32GB	1 ~ 3.5GB
F (isc)	210GB	local: 10s ~ 100s MB global: 20 ~ 25GB
K (ripe)	58GB	local: 10s ~ 100s MB global: 1 ~ 10GB

Previous OARC data collection (3)

- **Composition of traffic**

root server	time	packets	bytes	avg. pkt size (bytes)
C (cogent)	2005/09/22~23	UDP: 796M TCP: 31M (3.8%)	UDP: 59G TCP: 1.7G (2.8%)	UDP: 74 TCP: 54
F (isc)	2005/09/28~29	UDP: 1.4G TCP: 36M (2.5%)	UDP: 104G TCP: 1.5G (1.5%)	UDP: 74 TCP: 42
	2006/01/10~11	UDP: 1.8G TCP: 65M (3.3%)	UDP: 214G TCP: 2.8G (1.3%)	UDP: 113 TCP: 43
K (ripe)	2006/01/10~11	UDP: 1.6G TCP: 50M (3.0%)	UDP: 167G TCP: 2.1G (1.3%)	UDP: 104 TCP: 43

Analysis topics

- **Time-of-day usage difference**
- **Distribution of queries across anycast instances**
- **Distribution of response sizes and types**
- **Distribution of queries by gTLD and ccTLD**
- **Growth in and impact of DNSSEC**
- **Fraction of TCP request/response which are genuine DNS requests, not bogus**

Some preliminary results

- **Time-of-day usage difference**

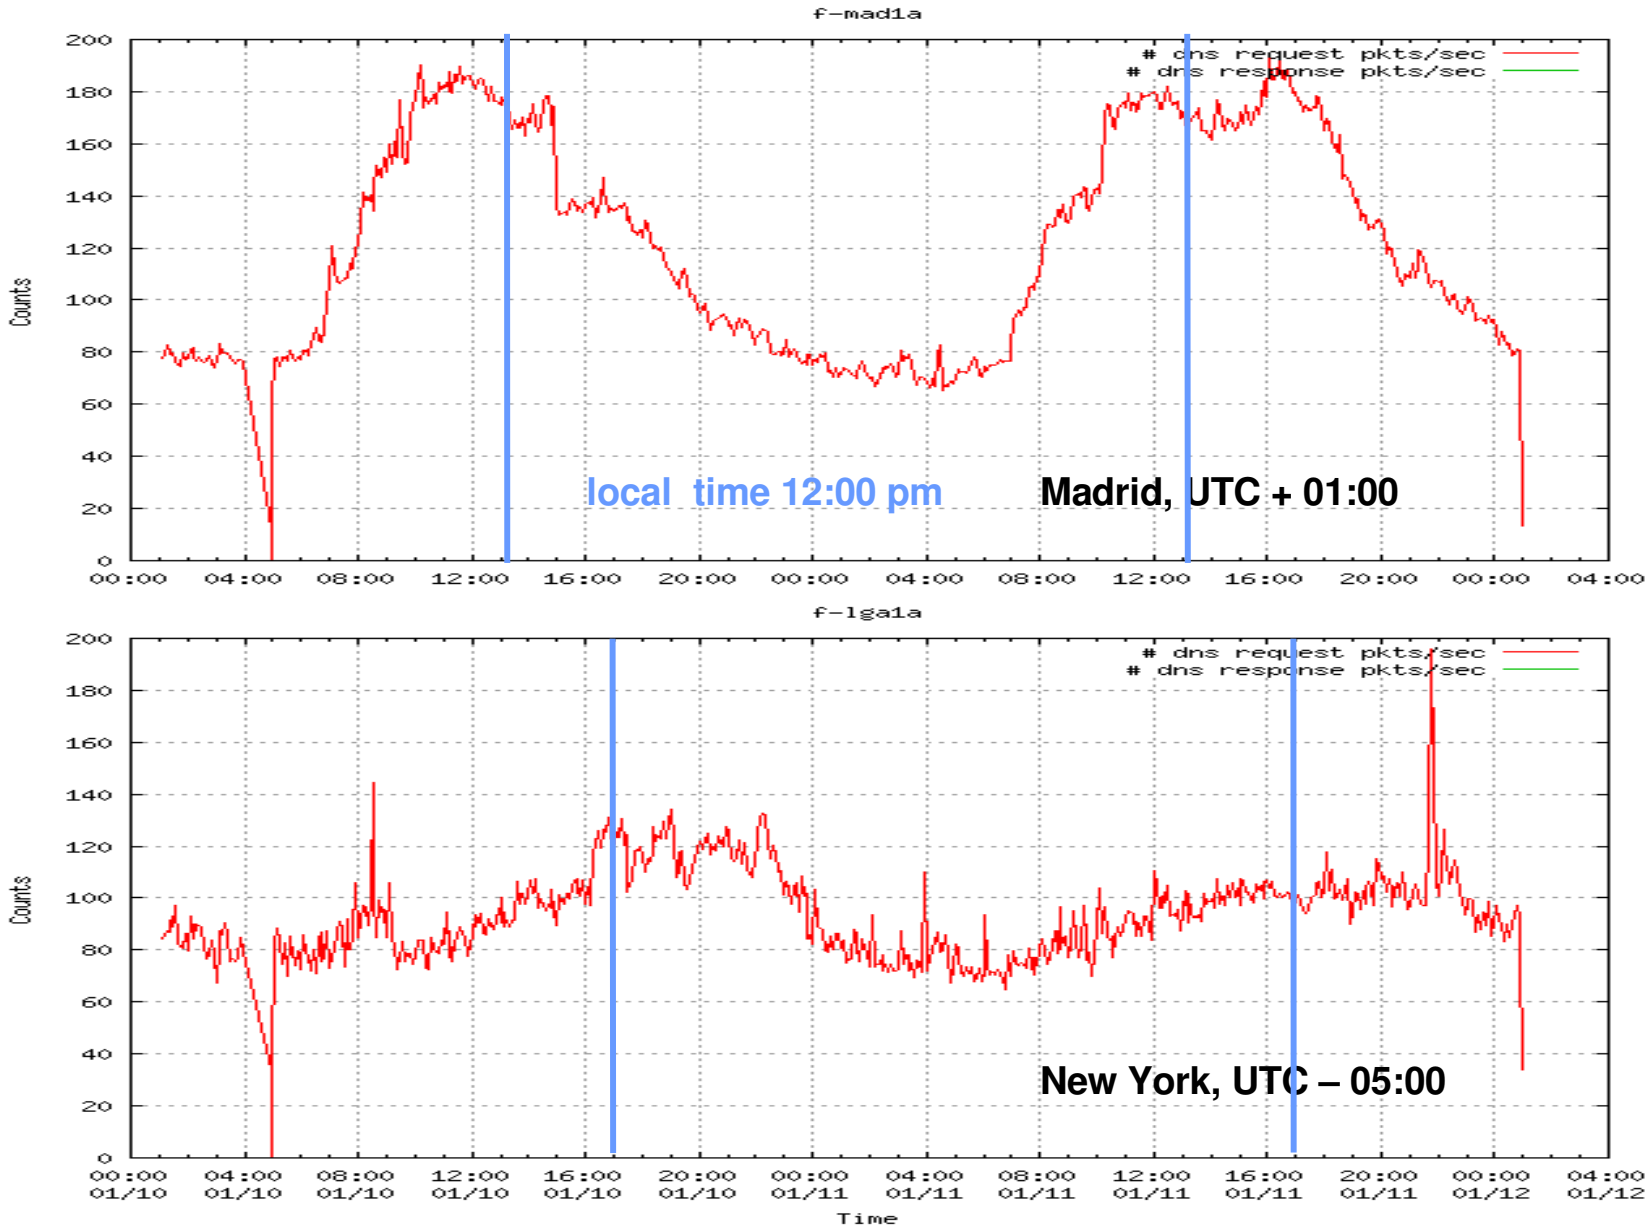

- **Time-of-day usage difference**

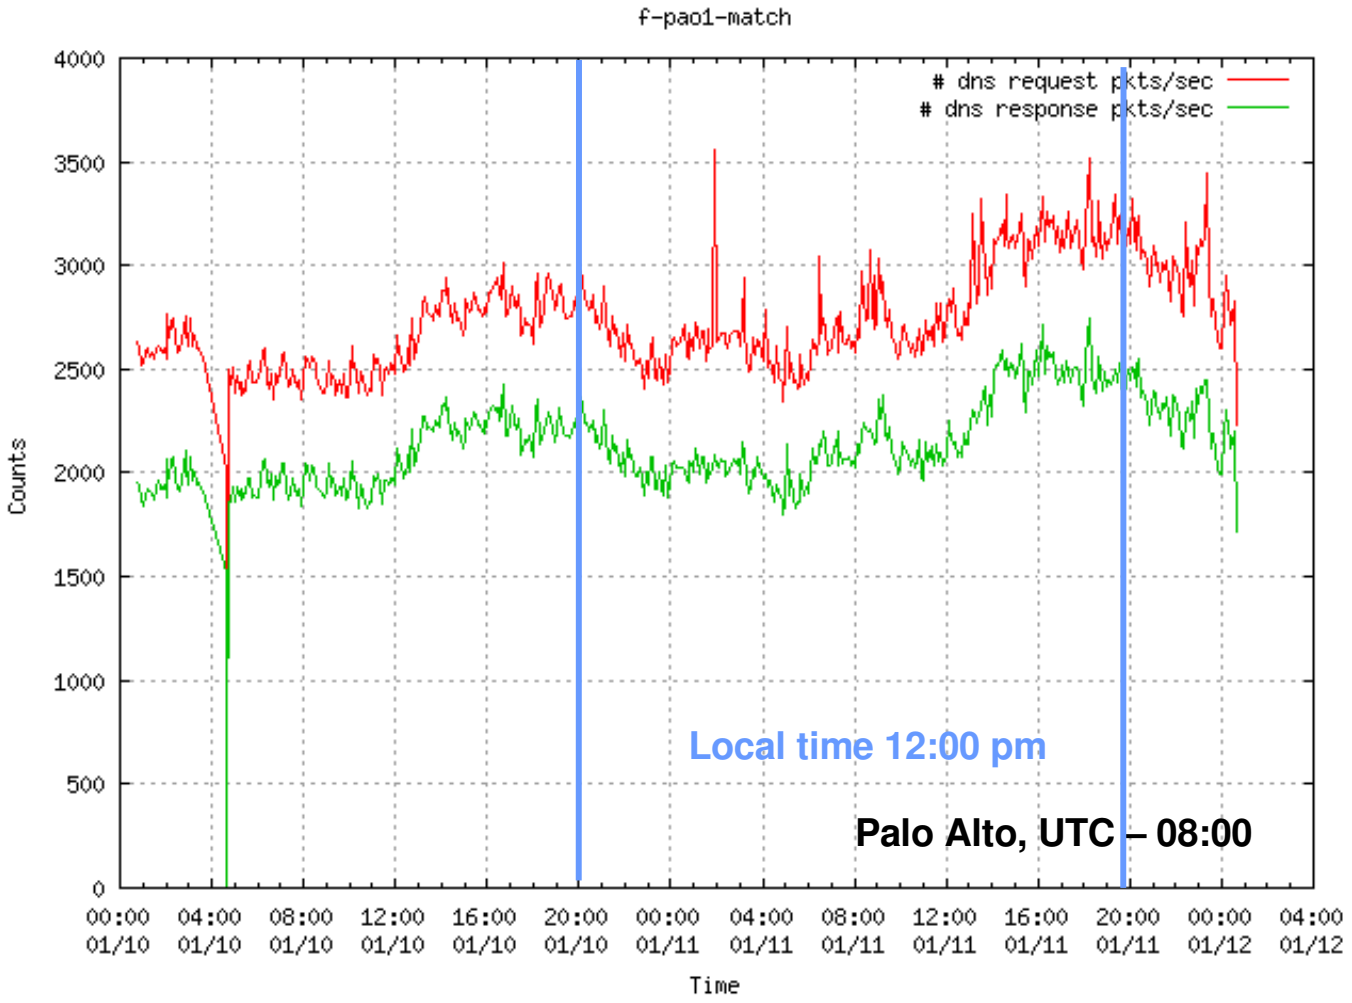

Even the global instance shows slightly diurnal pattern!

(though we can not just match it with the local time)

- **DNS anycast analysis**

Datasets

- Date: 2006/01/10~11
- Member: Cogent (C root), RIPE (K root), ISC (F root)
- Geographic: Netacuity database for geographic mapping
- Topological: RouteViews BGP tables for ASs and prefixes (Jan 10, 2006)

Scope

- Observed ASs: 19,237 (RouteViews: 21,883)
- Observed prefixes: 104,832 (RouteViews: 192,316)
- Observed IPs: 2,554,419

- **DNS anycast analysis – geographical distribution of clients**

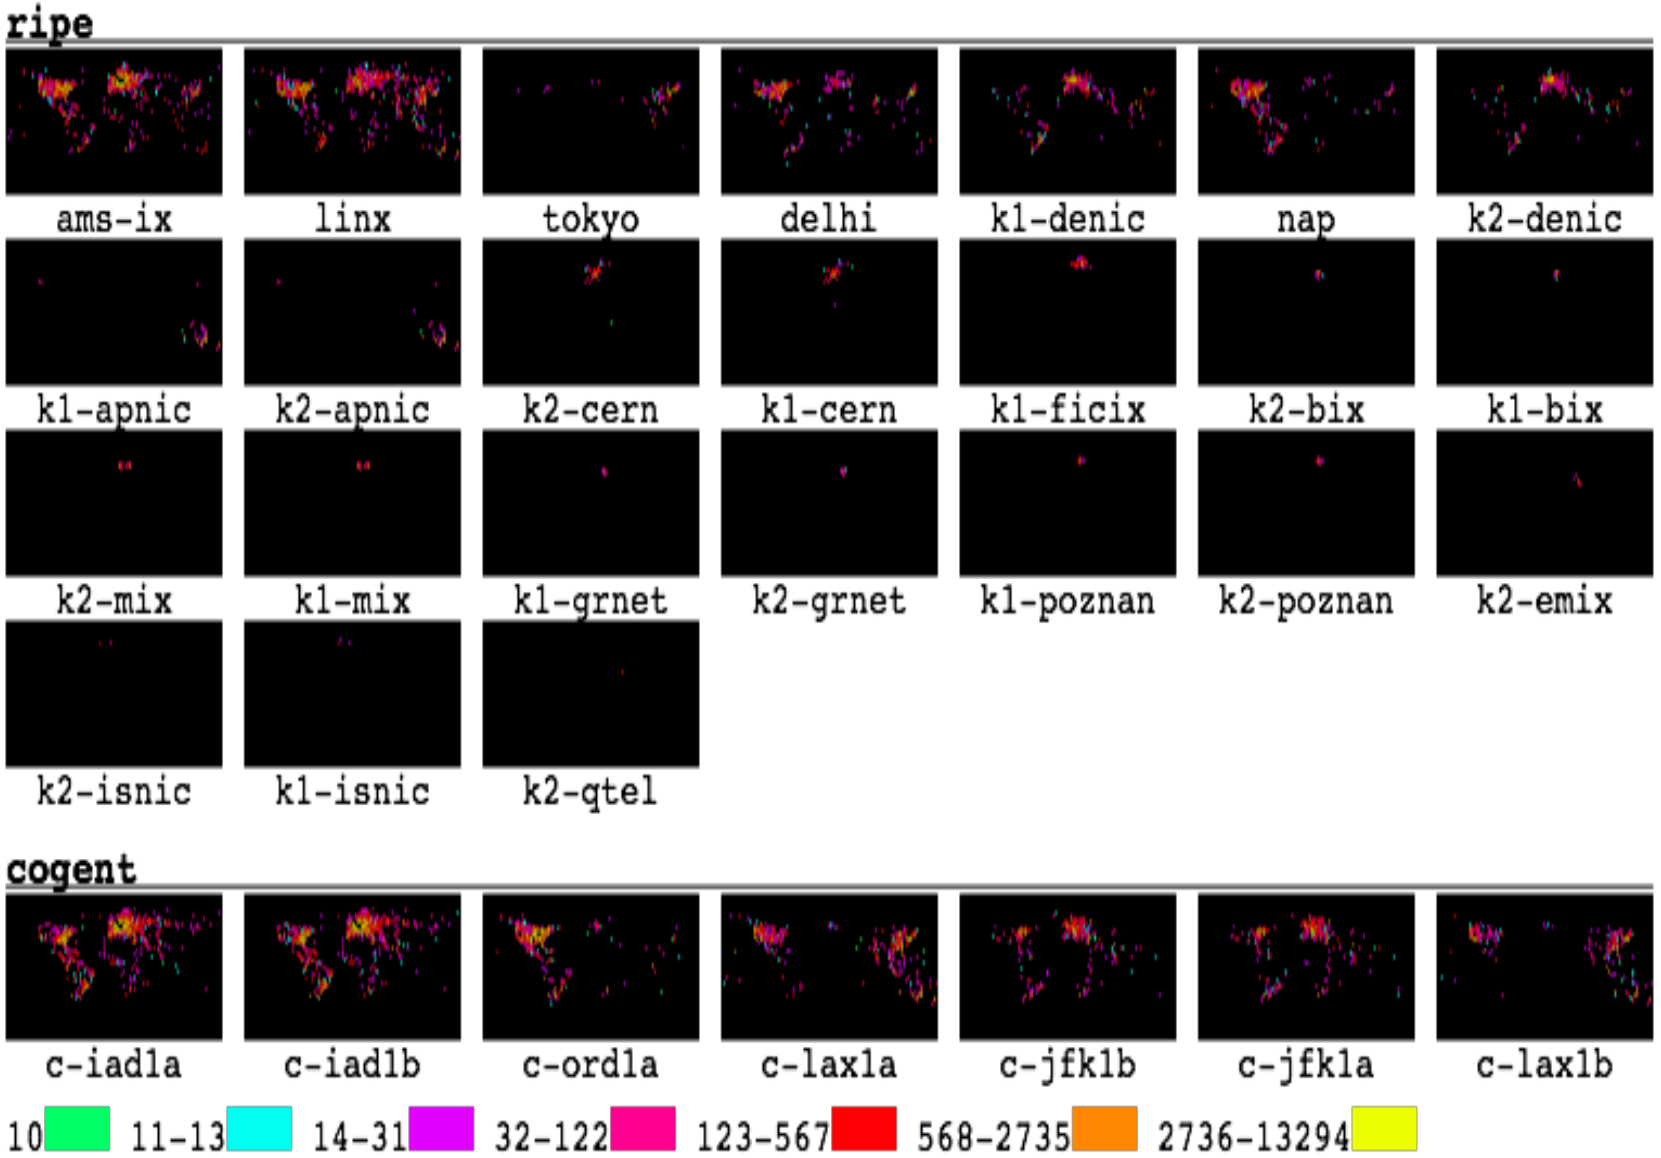

- **DNS anycast analysis – geographical distribution of clients**

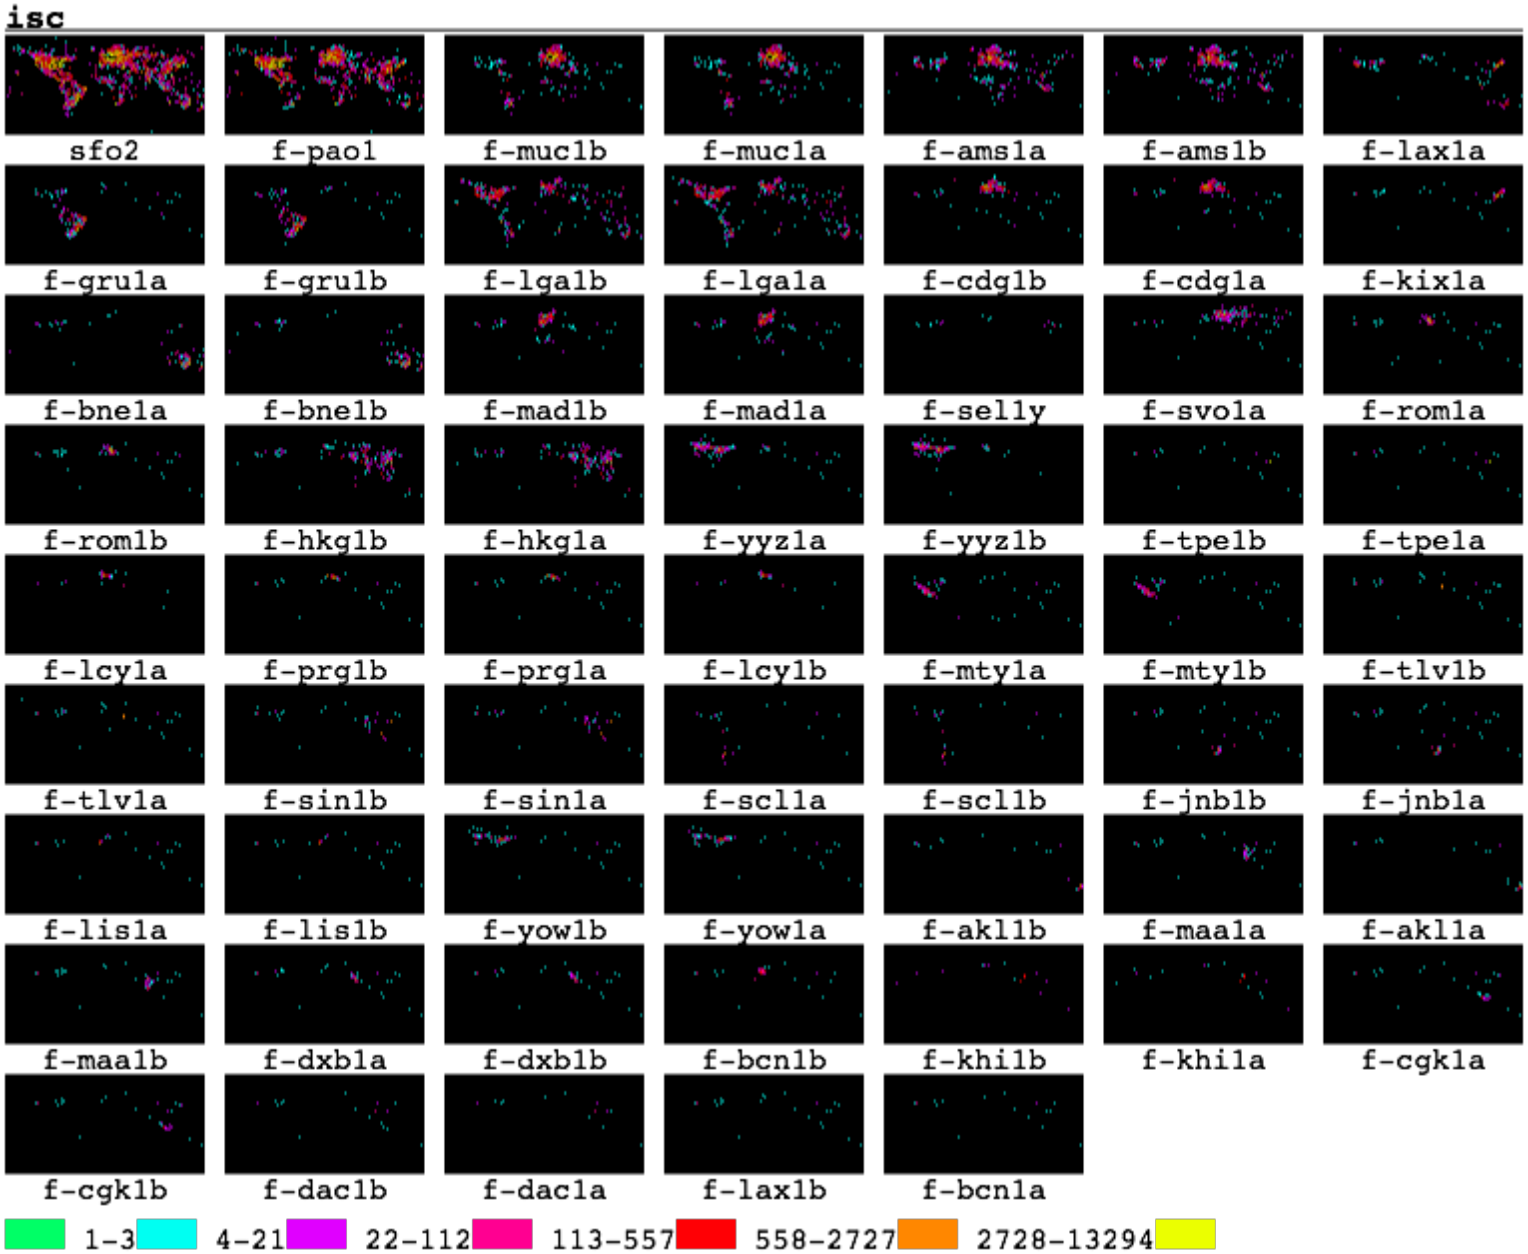

- **DNS anycast analysis – AS coverage per instance**
  ratio = ASs seen by instance / ASs seen by all

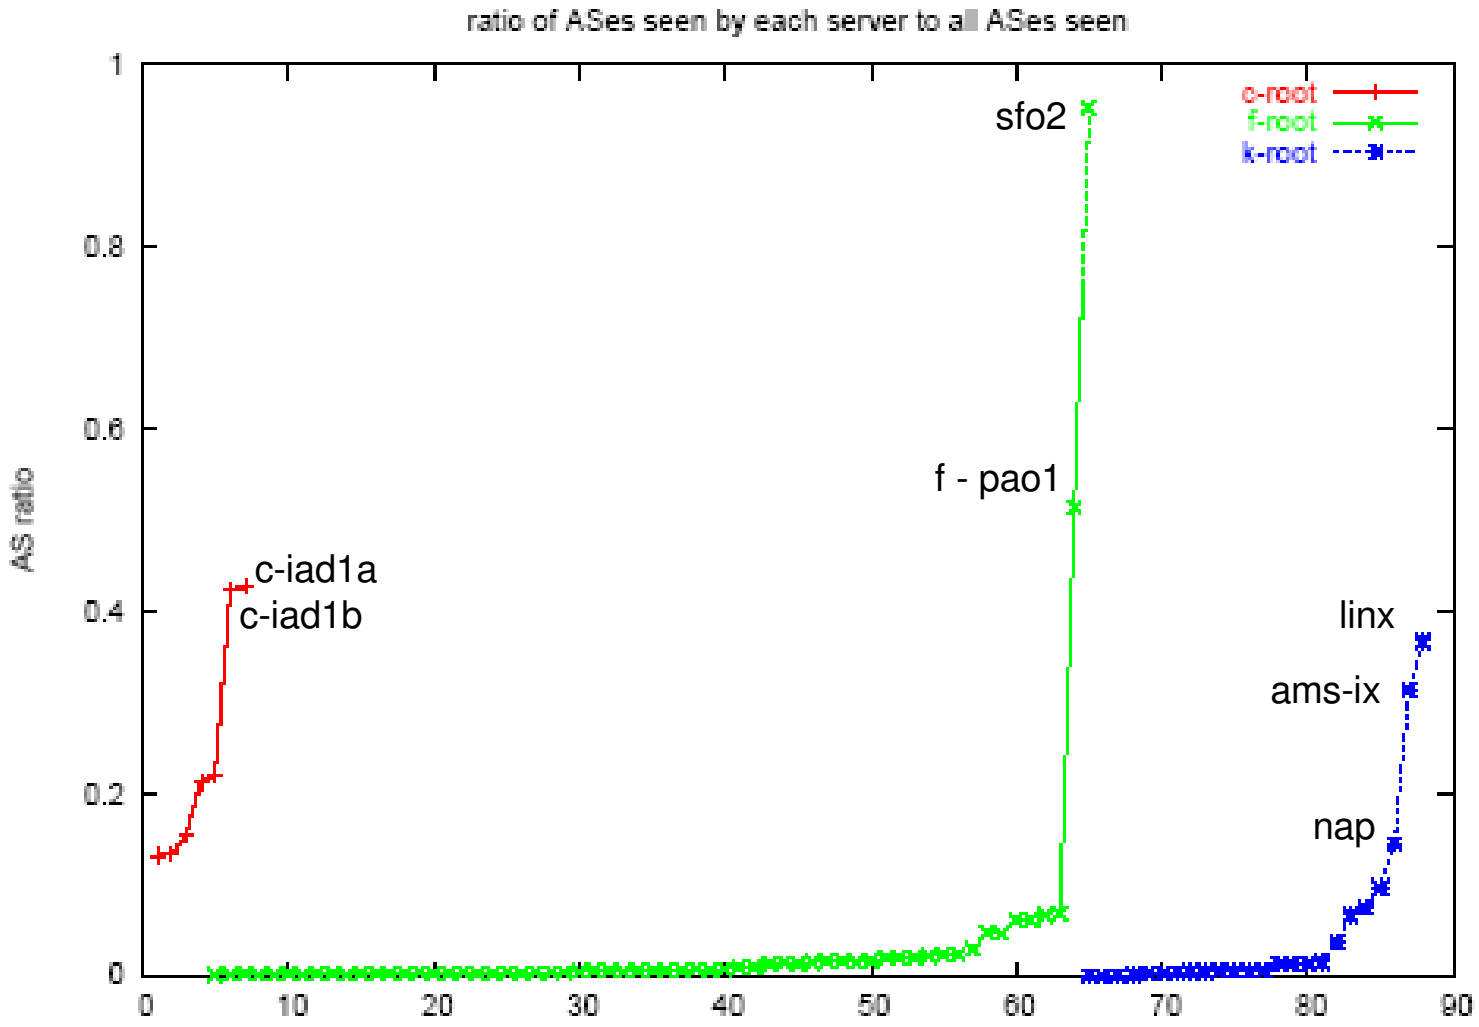

- **DNS anycast analysis – prefix coverage per instance**
  ratio = prefixes seen by instance / prefixes seen by all

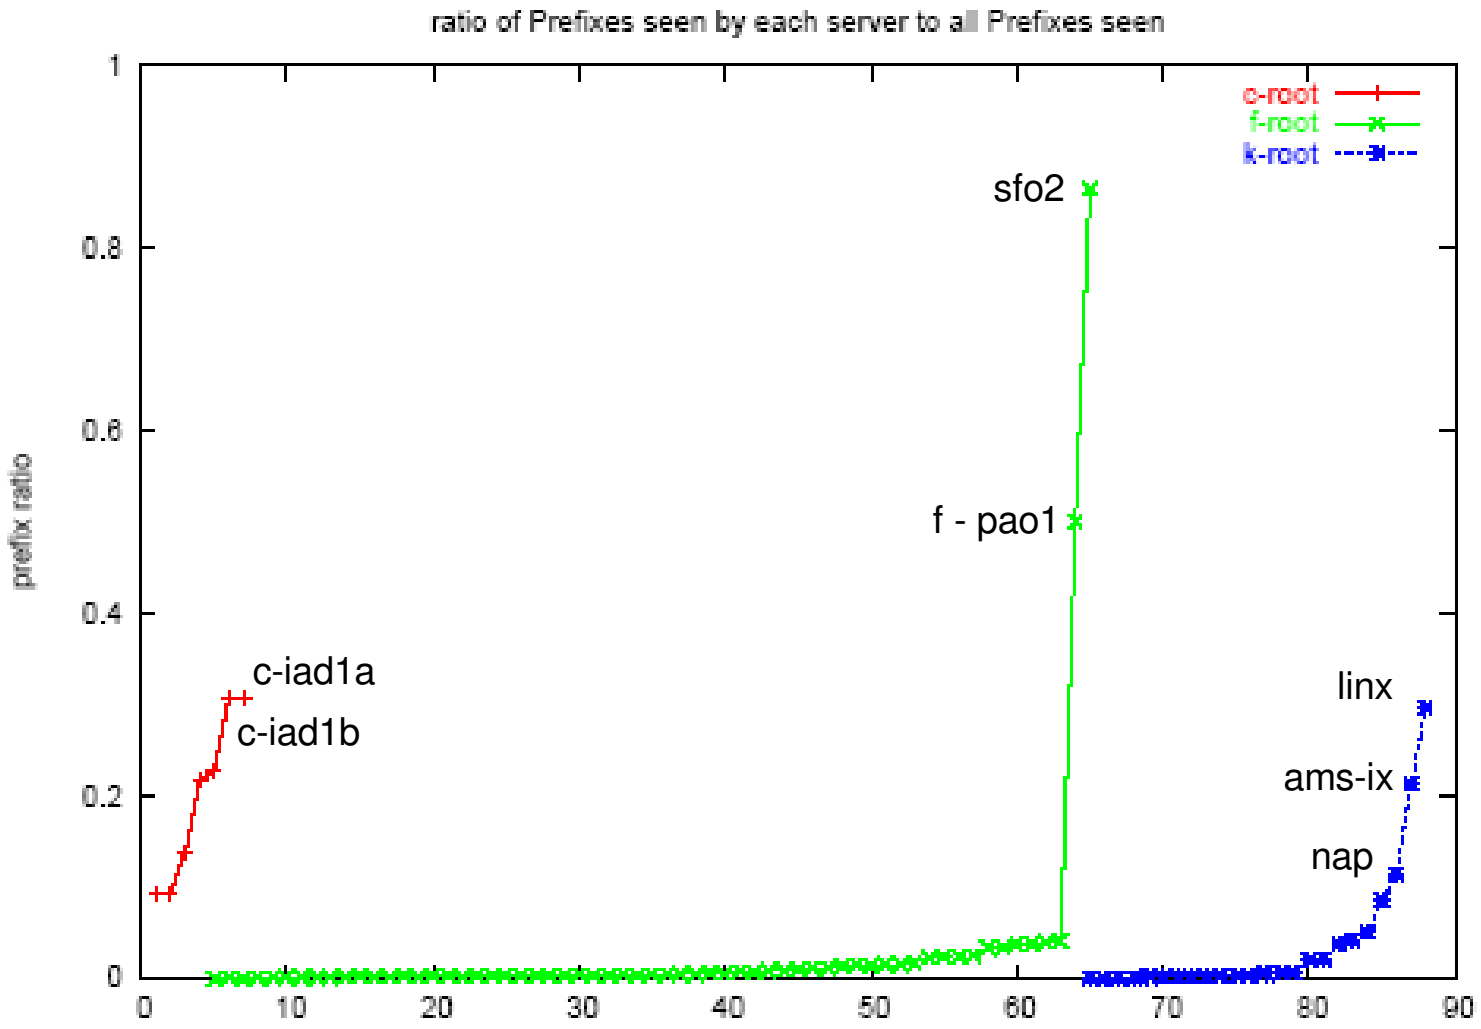

- **DNS anycast analysis – client (IP address) coverage per instance**
  ratio = clients seen by instance / clients seen by all

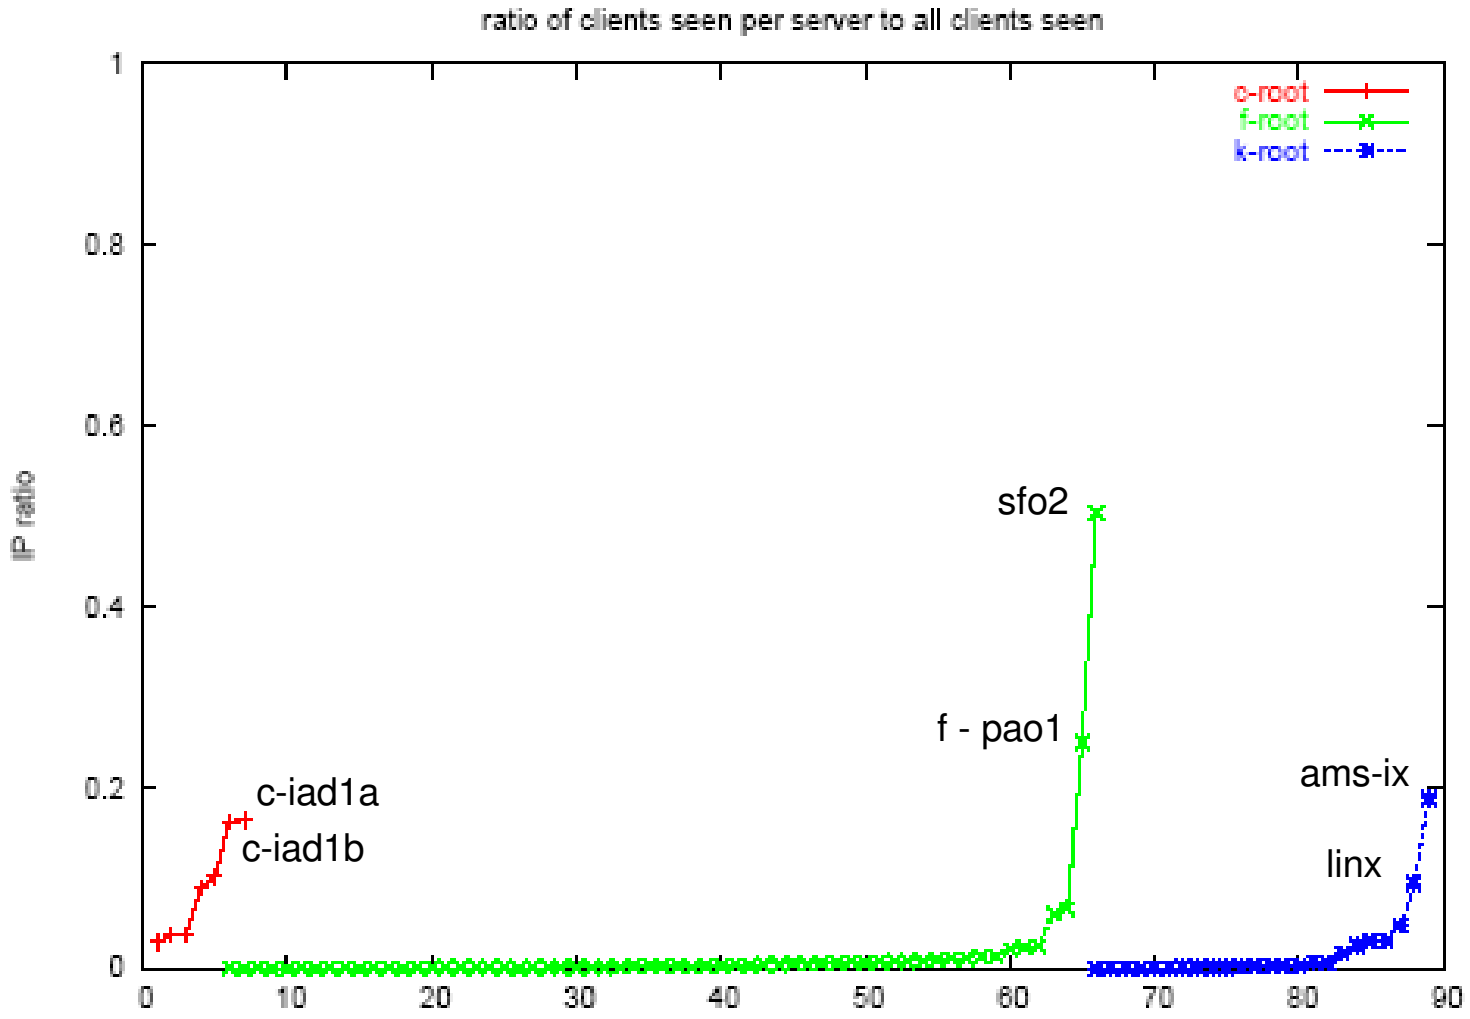

- **Size distribution of requests and responses**

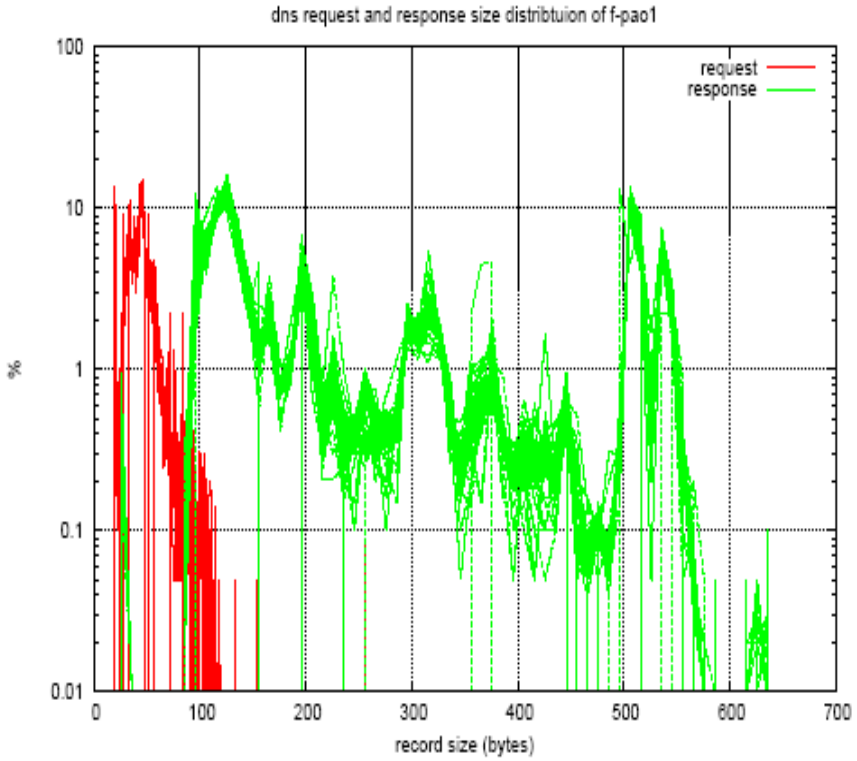

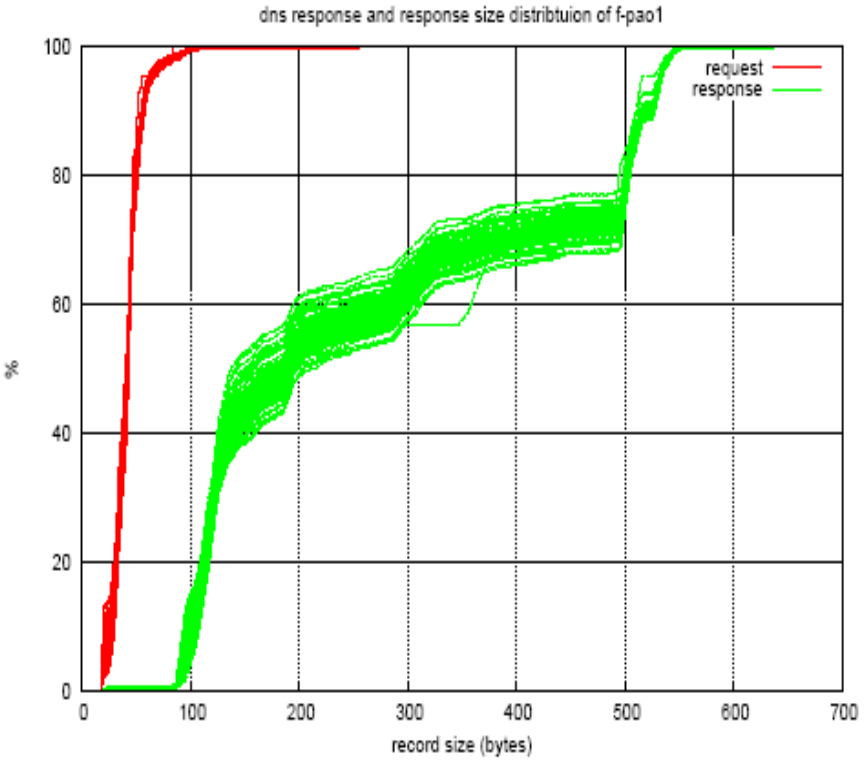

pdf (y: log scale)

cdf (y: linear scale)

Request and response size distributions every 5m at Palo Alto

-- 80% of requests are shorter than 50B,

-- 80% of responses are shorter than 500B

- **Size distribution of UDP queries for gTLD and ccTLD, over all the OARC datasets**

The ratios of the query sizes larger than 100 bytes are tiny, not shown here…

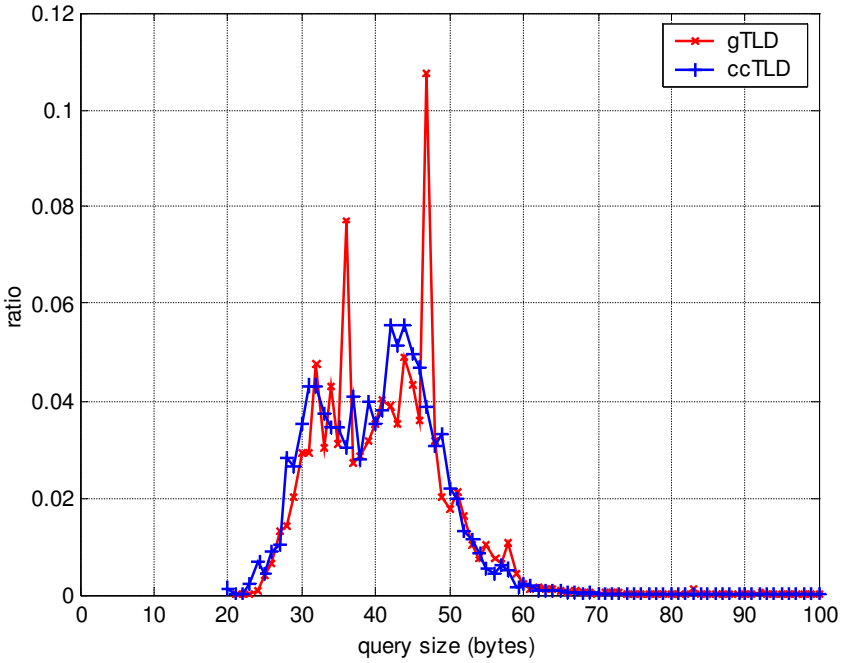

pdf

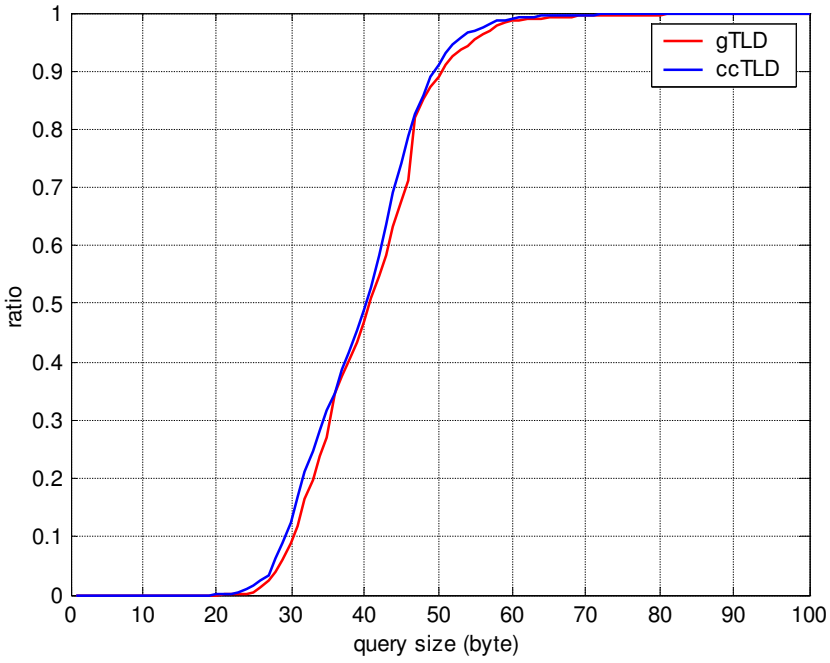

cdf

	# of queries	min size (B)	max size (B)	avg. size (B)
gTLD	2.08G	21	1279 (faked?)	41
ccTLD	0.98G	20	1279 (faked?)	40

Overall, 90% of the queries are shorter than 50 bytes.

- **Fraction of genuine TCP DNS traffic**
  - We saw few genuine TCP requests
  - Most of TCP "requests" (to port 53) are bogus (syn, fin, ack, … no payload)
  - Further analysis needed
- **Growth in and impact of DNSSEC**
  - More analysis is needed. For example, what % of queries/responses include DNSSEC-related RRs?
  - Only a few instances collected response data

Recommendations for future collections

- **We want as much data as possible**
  - **Both TCP and UDP for port 53**
  - **Both queries and responses**
  - **From all root and gTLD instances**
  - **Over 48 hour long, mid-week preferred**
  - **Synchronized to UTC time**
- **Why so greedy?**
  - **TCP accounts for a small fraction of DNS traffic, but may well increase**
  - **Responses require significant storage, but are necessary to answer questions about DNSSEC usage**
  - **Differences between 'global' and 'local' instances**
  - **Average daily traffic, diurnal patterns, anomalies within a day**

Thanks

- **To my colleagues at CAIDA (who did most of the actual work)**
  - **Brad Huffaker**
  - **Ziqian Liu**
  - **Marina Fomenkova**

- **And to the OARC team, ISC and root operators (who collected the traces)**

- **Questions?**