2007 Day In The Life
DNS Root Server Analysis

Duane Wessels
Haven Hash
The Measurement Factory/CAIDA

WIDE+CAIDA Workshop #8
July 21, 2007

DITL 2007

- Day In The Life of The Internet. Okay, two days.
- 48 hour period: Jan 9 00:00:00 to Jan 10 23:59:59 UTC
- Primary focus is DNS and root servers, but other data was collected as well.
- We have data from C-, F-, K-, and M-roots, which is the subject of this presentation.
- Data is 740 GB compressed pcap files.
- 10,000,000,000 DNS queries.

Terminology

- Server: a collection of DNS nameservers operating under the same IP address.
  - c.root-servers.net is a server

- Instance: an anycast instance of a server.
  - k-milan is an instance of k.root-servers.net.

  Load-balanced nodes are combined into a single instance.
  - c-lax1a and c-lax1b are load-balanced members of the c-root LAX instance.

- Client: an IP address sending DNS queries.

Merging Pcaps

- First step was to create a single, “merged” pcap stream with all packets in chronological order.
- Created hour-long chunks for all instances, using *tcpdump-join* and *tcpdump-split*. Keep only data within the 48-hour DITL period. Queries only.
- Changed pcap timestamps for instances with known clock skew.
- Rewrote server IP addresses to encode server and instance.
  - e.g., 192.5.5.251 becomes 6.0.0.11 to represent the 11th instance of F-root.
- Merged all hour-long instance files into timestamp-sorted files with *mergecap*.

Analysis Software

- C++ program reads pcap files and keeps various counters.
- Runs at about 40,000 packets/second, or about 80% the rate of "pcap time."
  - i.e., takes 60 hours to analyze 48 hours of data.
- Needs about 3GB RAM.
- Data goes into Postgres
- SQL SELECT statements and perl scripts produce data for ploting with *ploticus*.

II 1) Average rates of requests.

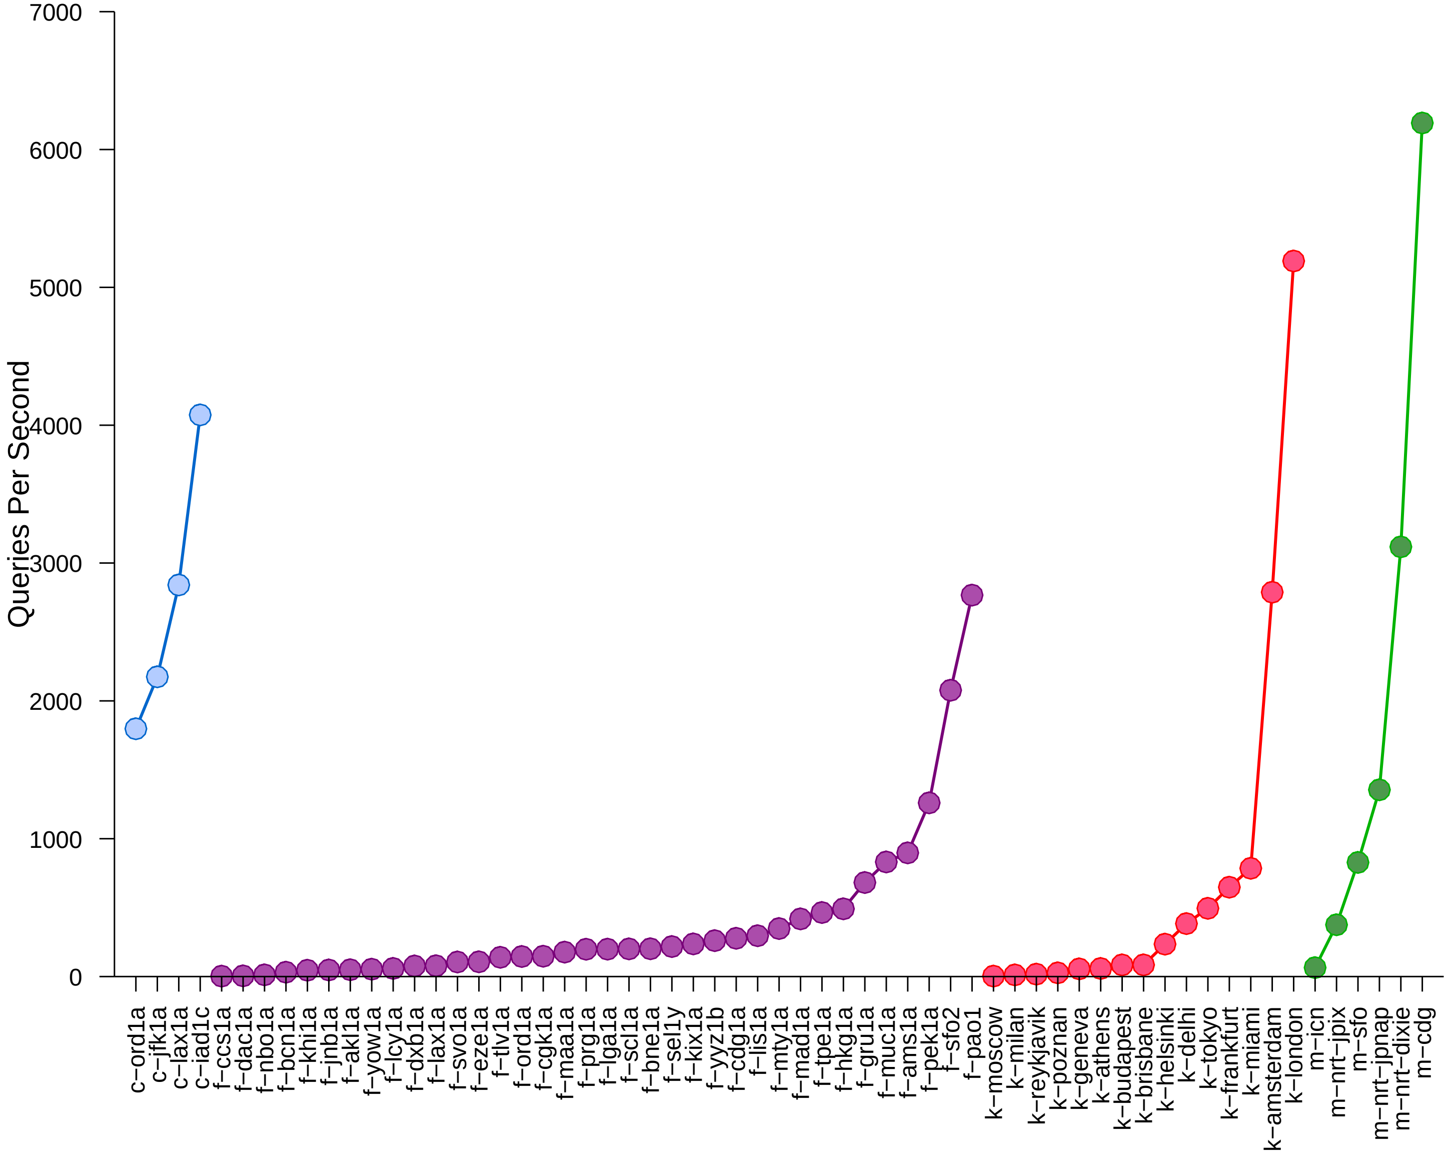

II 1) Average rates of requests.

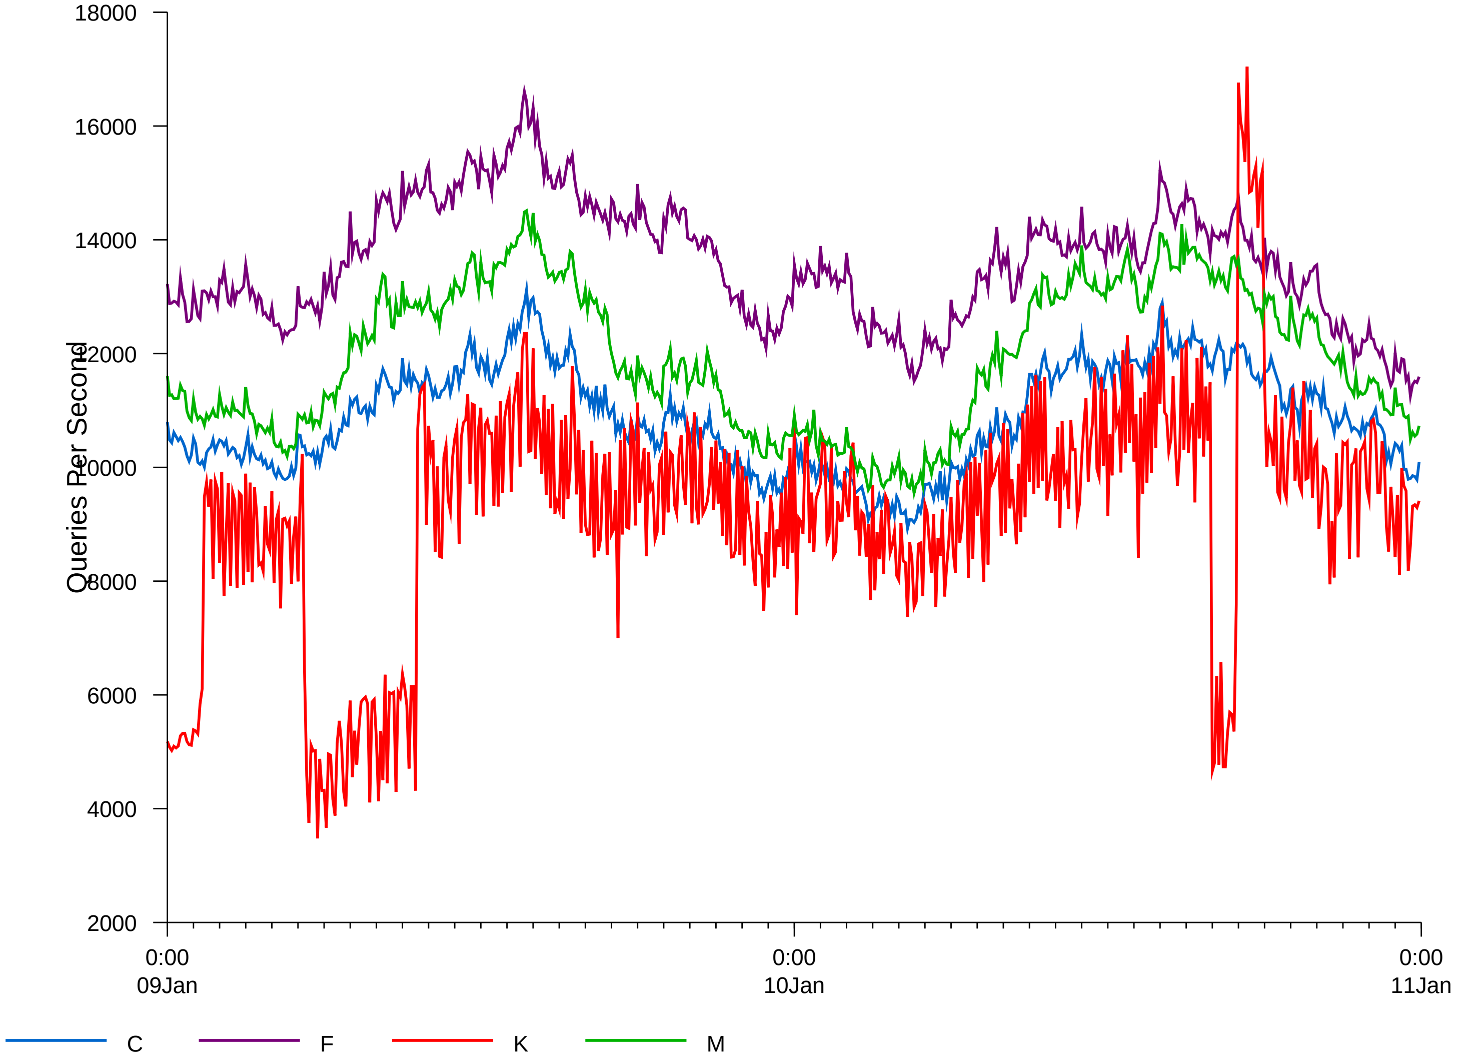

II 2) The average number of clients per second seen at each instance.

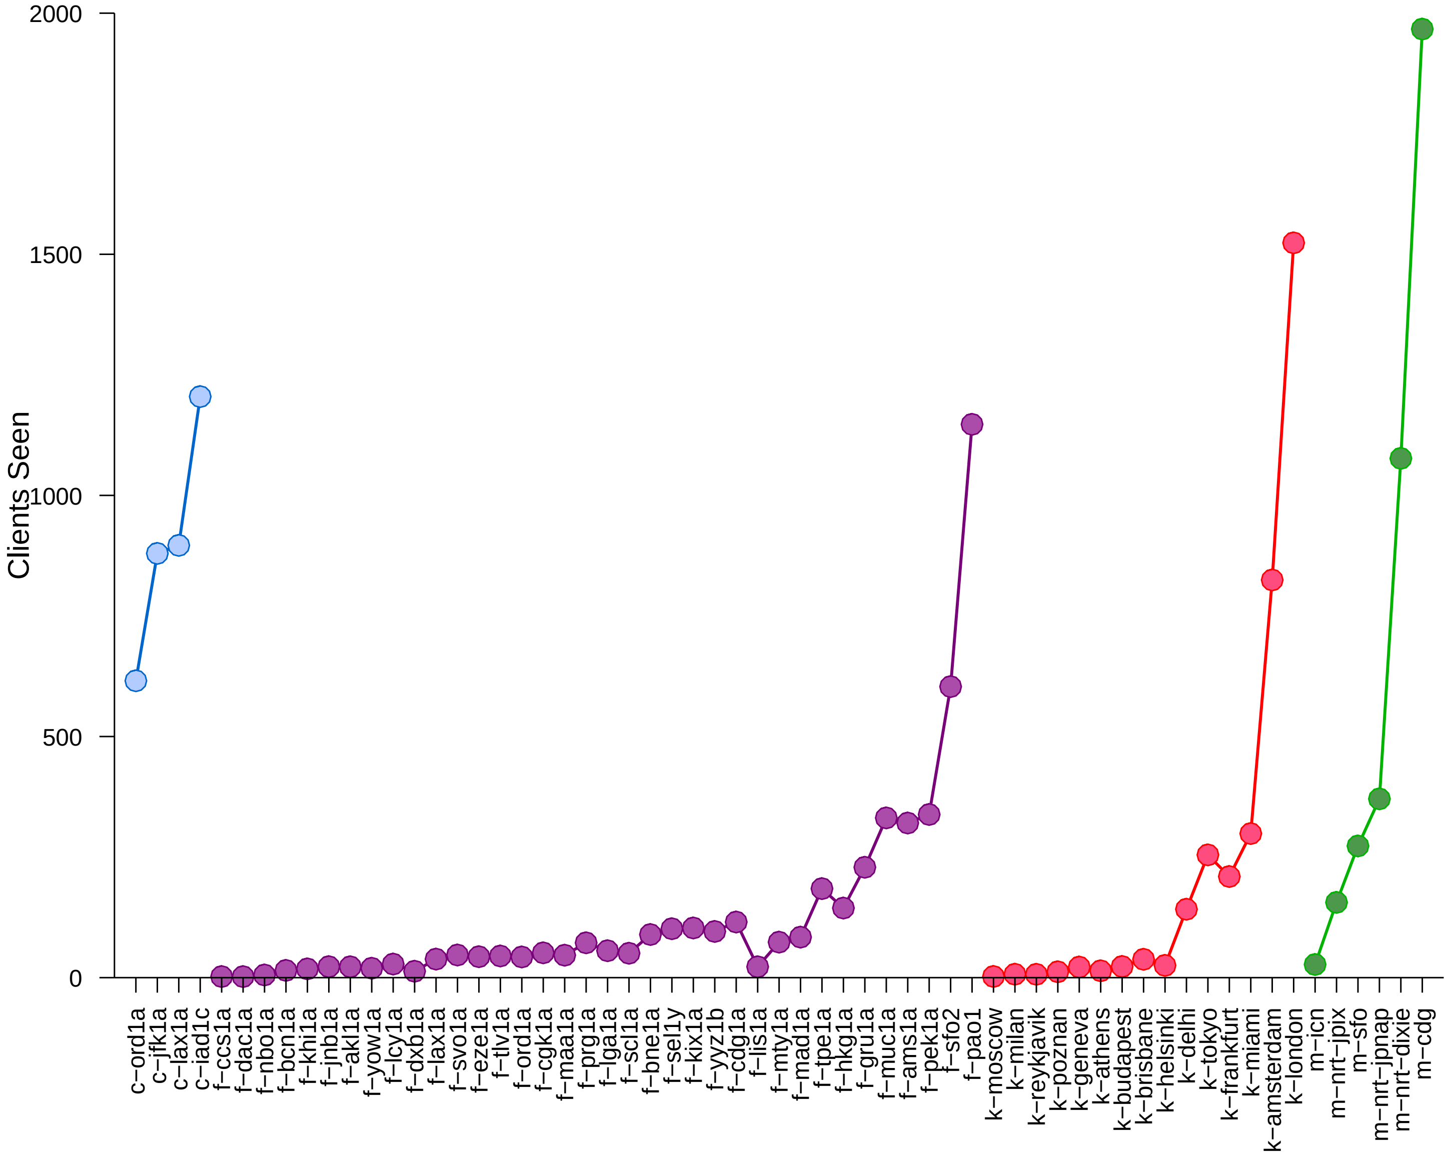

II 2) The number of clients per second seen at each C−root instance.

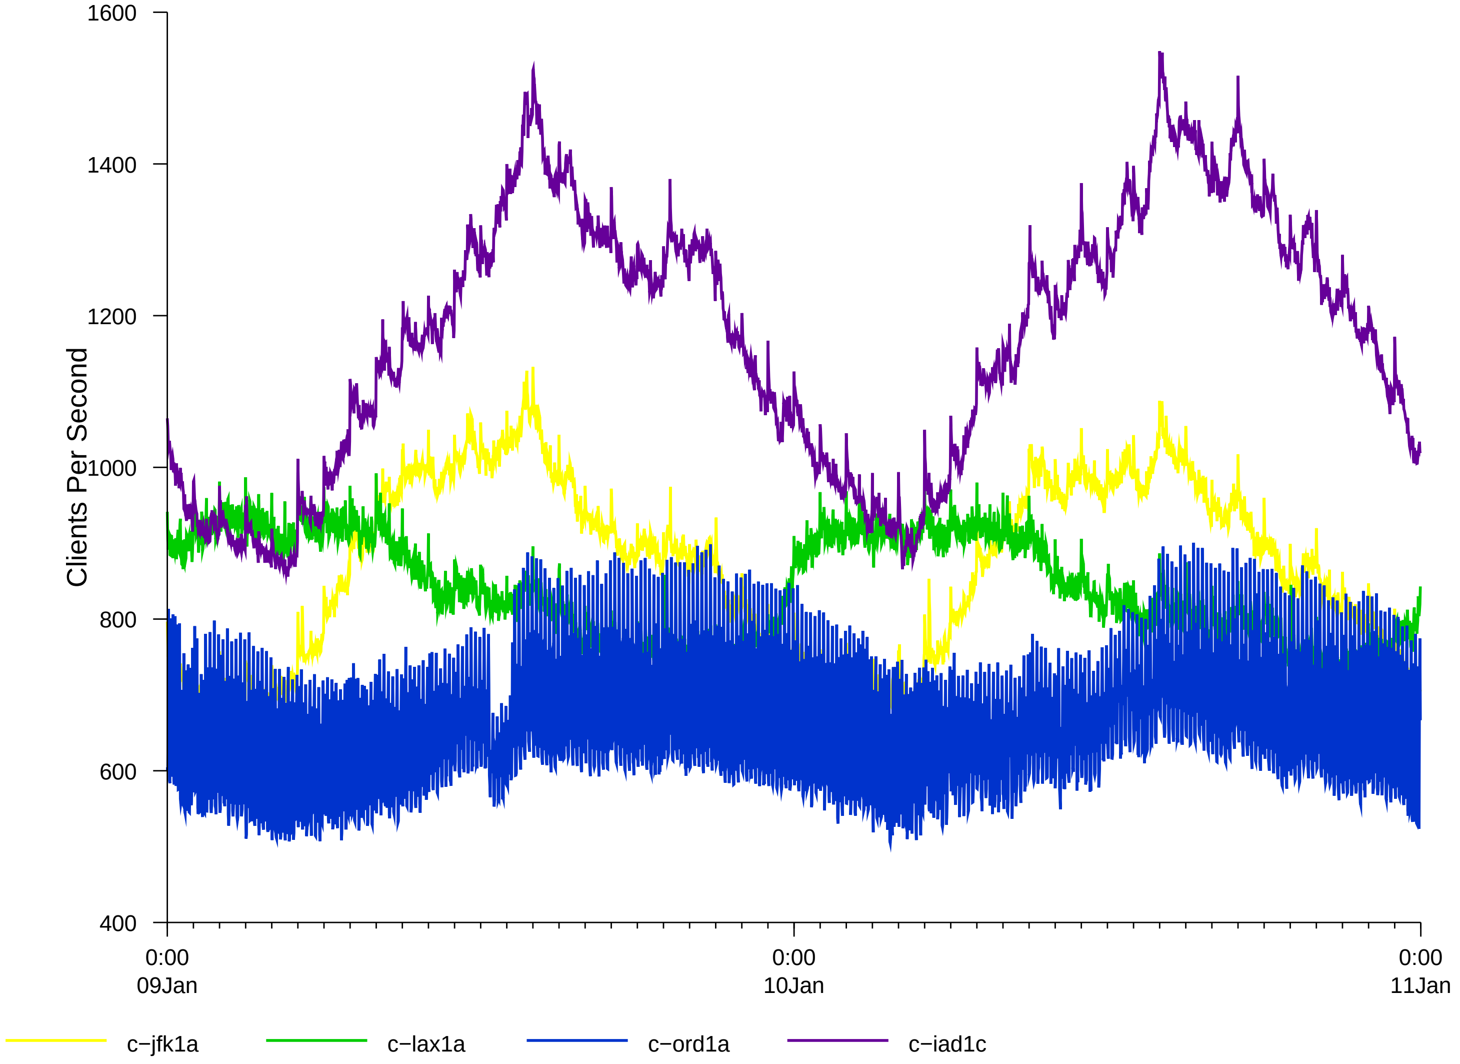

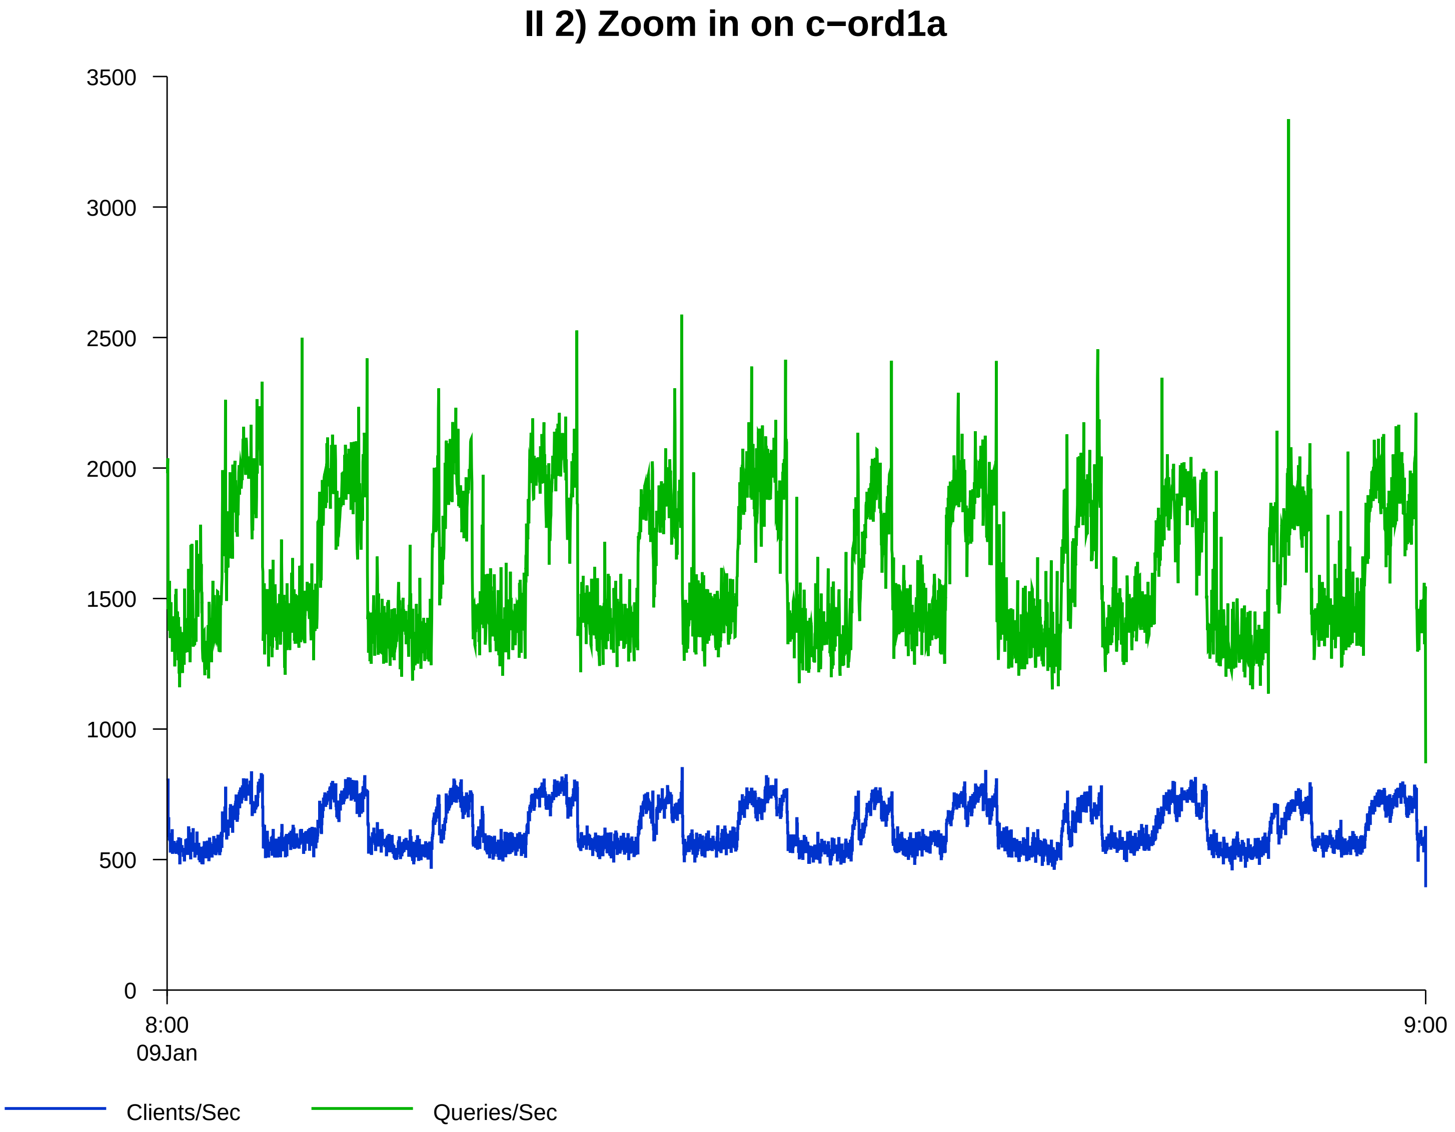
II 2) Zoom in on c−ord1a
3500
3000
2500
2000
1500
1000
500
0
8:00
09Jan
9:00
Clients/Sec
Queries/Sec

The cause??

```
Date: Thu, 11 Jan 2007 01:03:47 +0000
From: Paul Vixie <paul@vix.com>
To: wessels@0arc.isc.org
Subject: oops

#ord1a.c:i386# jobs
[1]  + Running                       ./tcpdump -s 0 -n -w oarc.tcpd. -z gzip
-P 5 host c.root-servers.net
#ord1a.c:i386# kill %1

626638995 packets captured
667208048 packets received by filter
15549 packets dropped by kernel
[1]    Done                          ./tcpdump -s 0 -n -w oarc.tcpd. -z gzip
-P 5 host c.root-servers.net

i had two tcpdumps running on one of the c-root boxes...
```

II 2) The number of clients per second seen at each F−root instance.

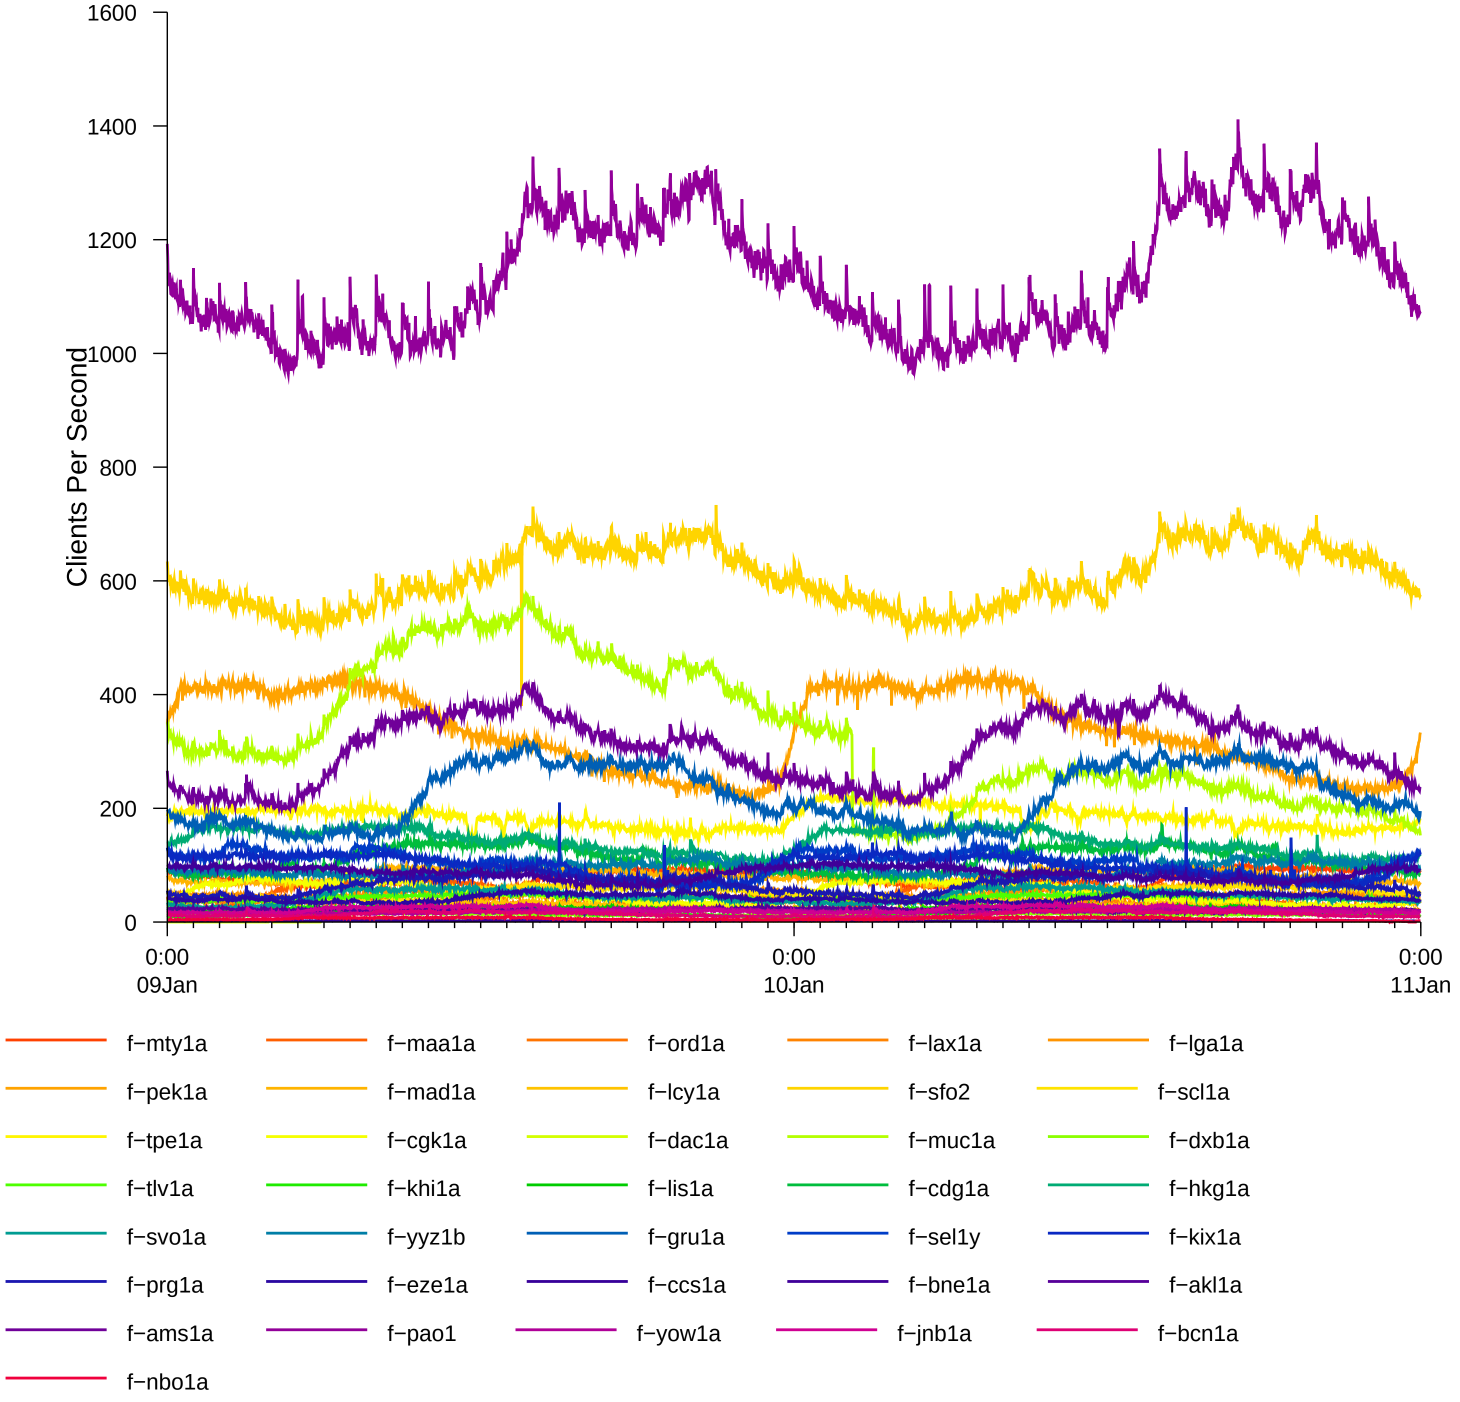

II 2) The number of clients per second seen at each K−root instance.

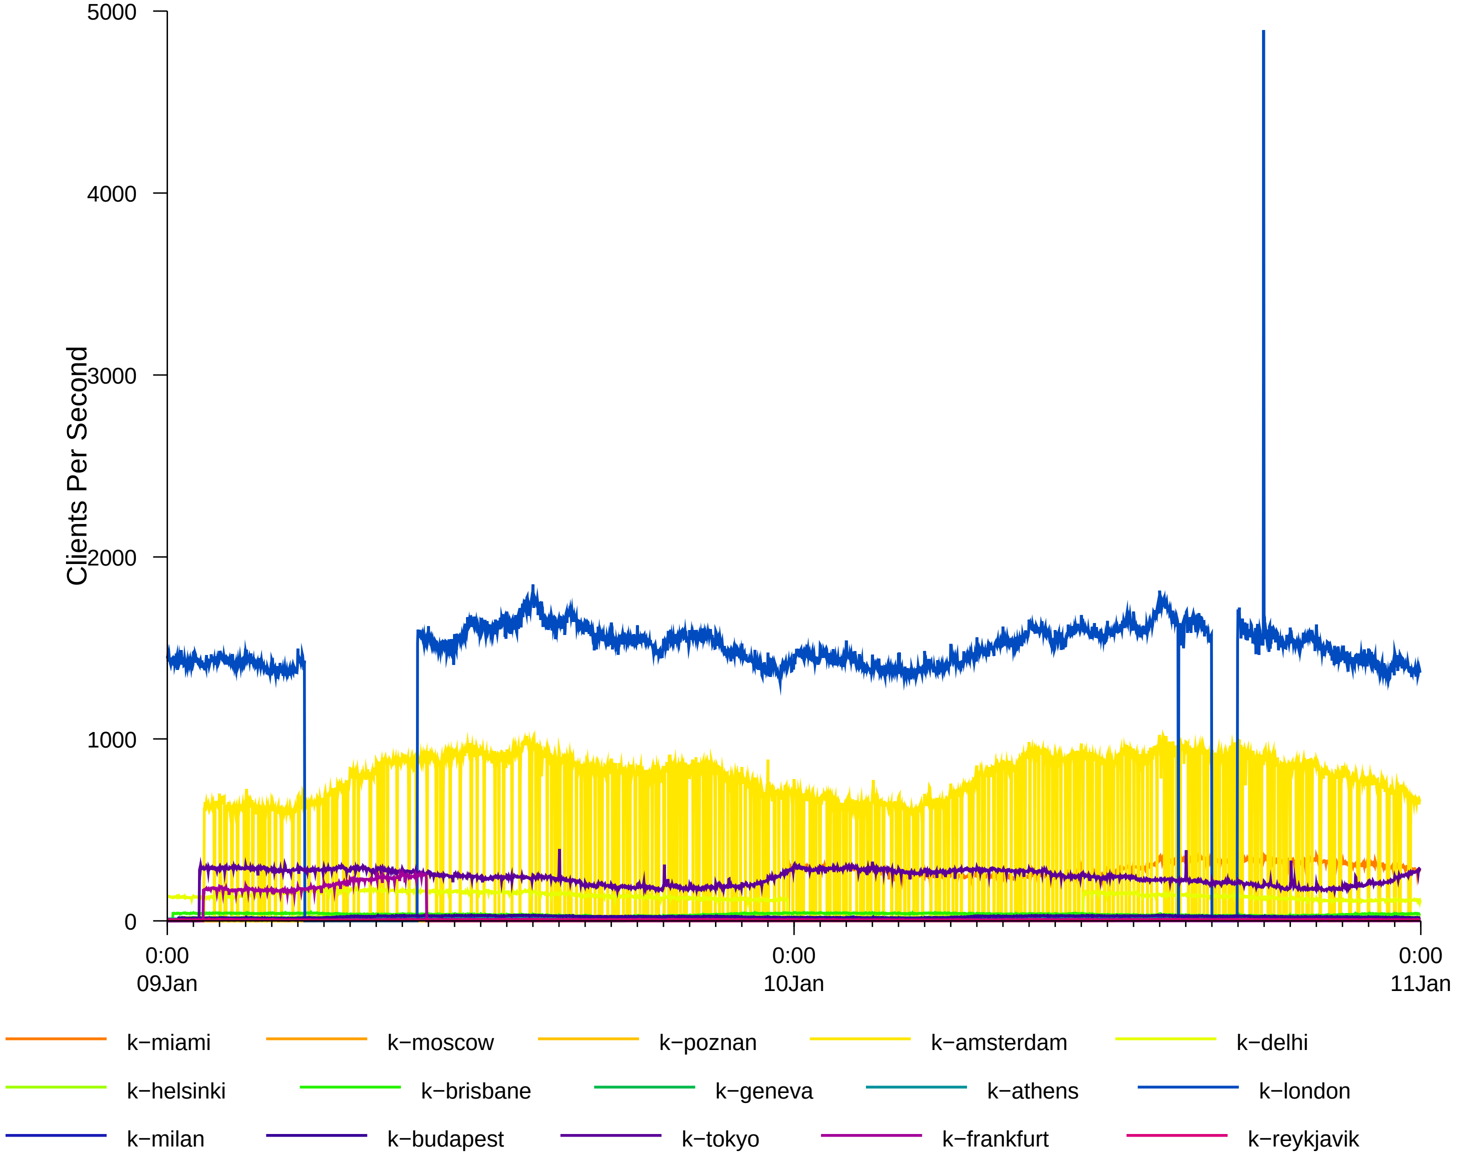

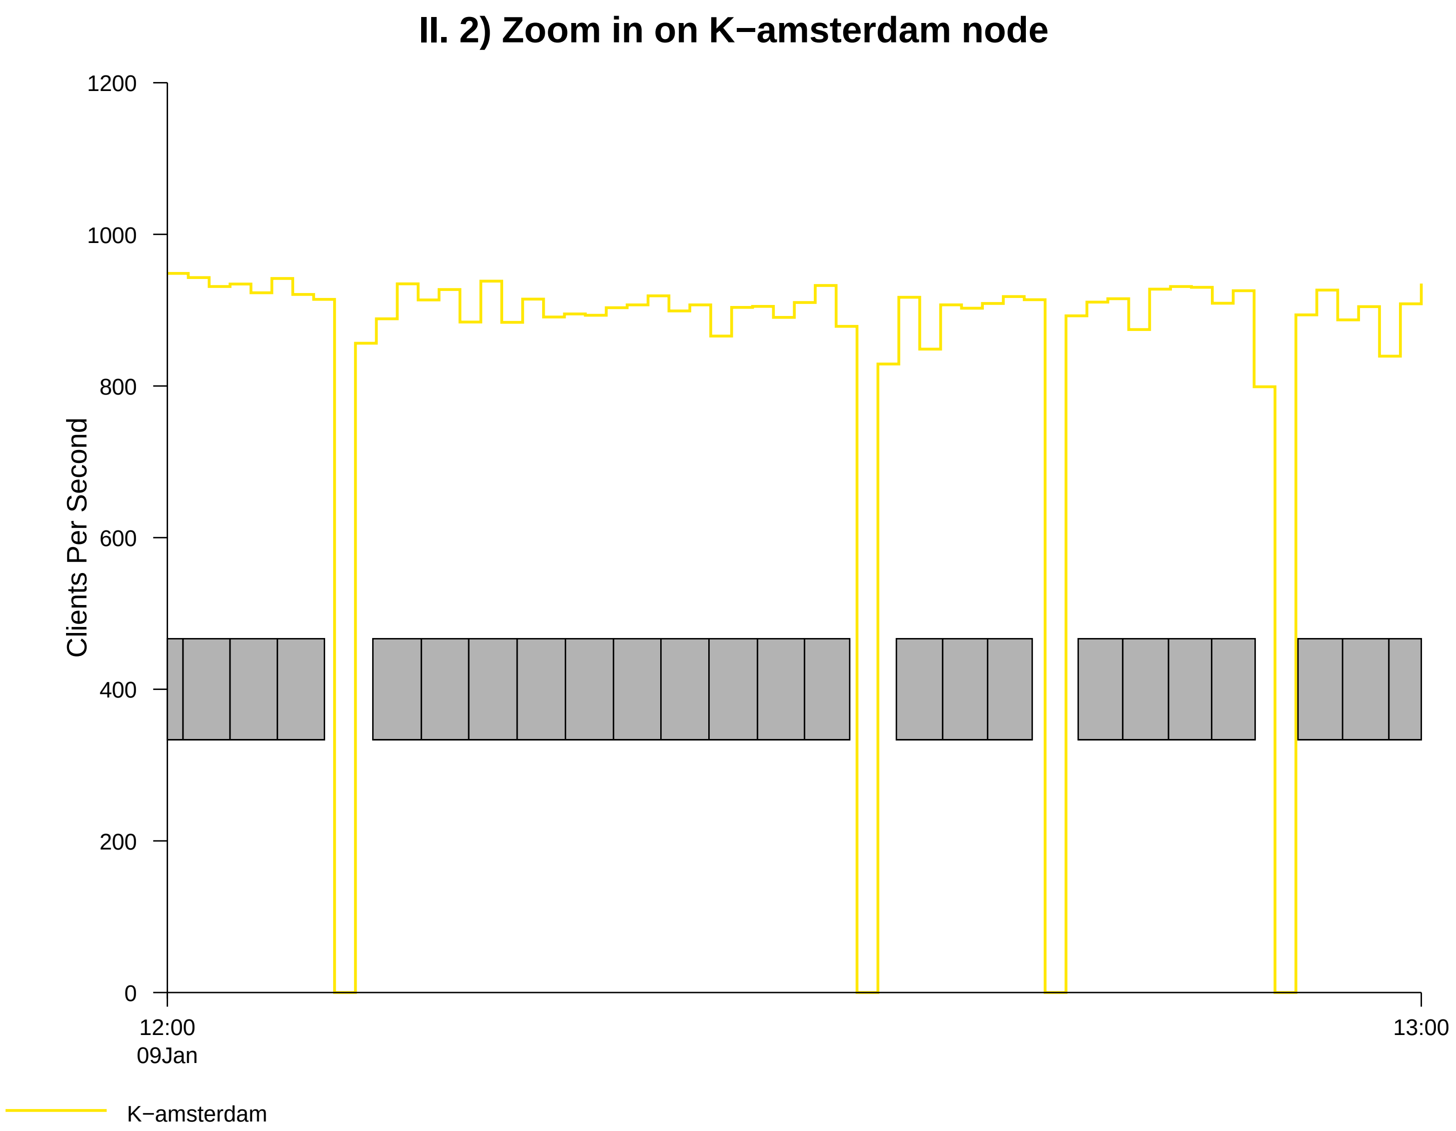
II. 2) Zoom in on K−amsterdam node
Clients Per Second
1200
1000
800
600
400
200
0
12:00
09Jan
13:00
K−amsterdam

II 2) The number of clients per second seen at each M−root instance.

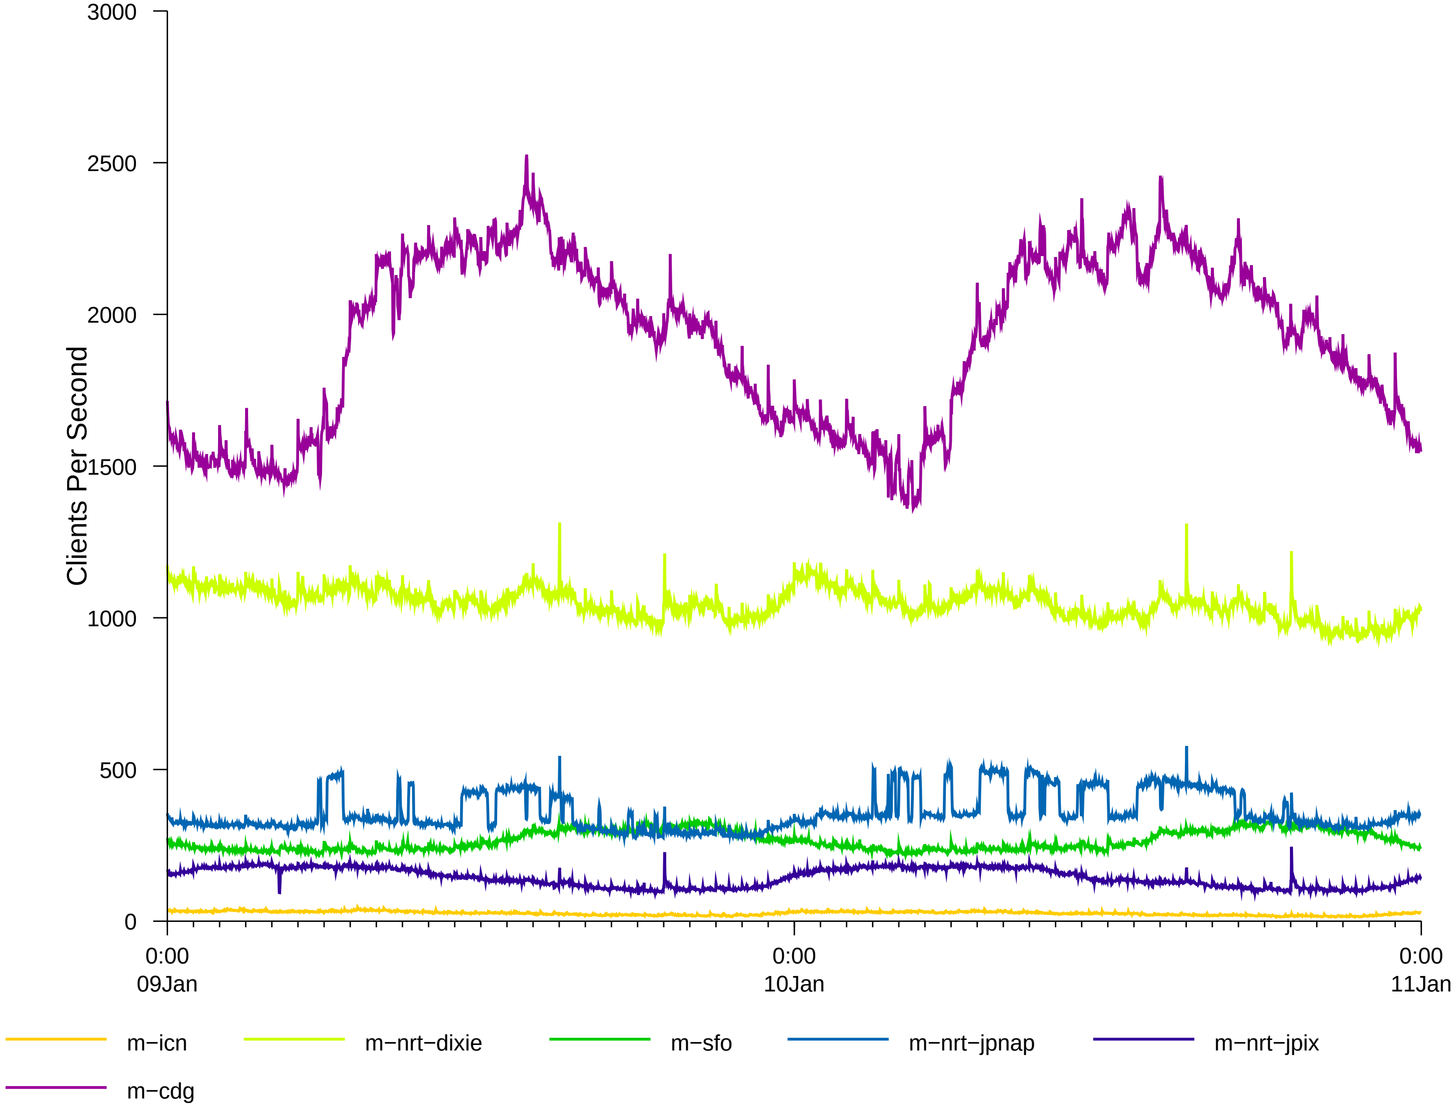

II 3) Topological coverage by ASes.

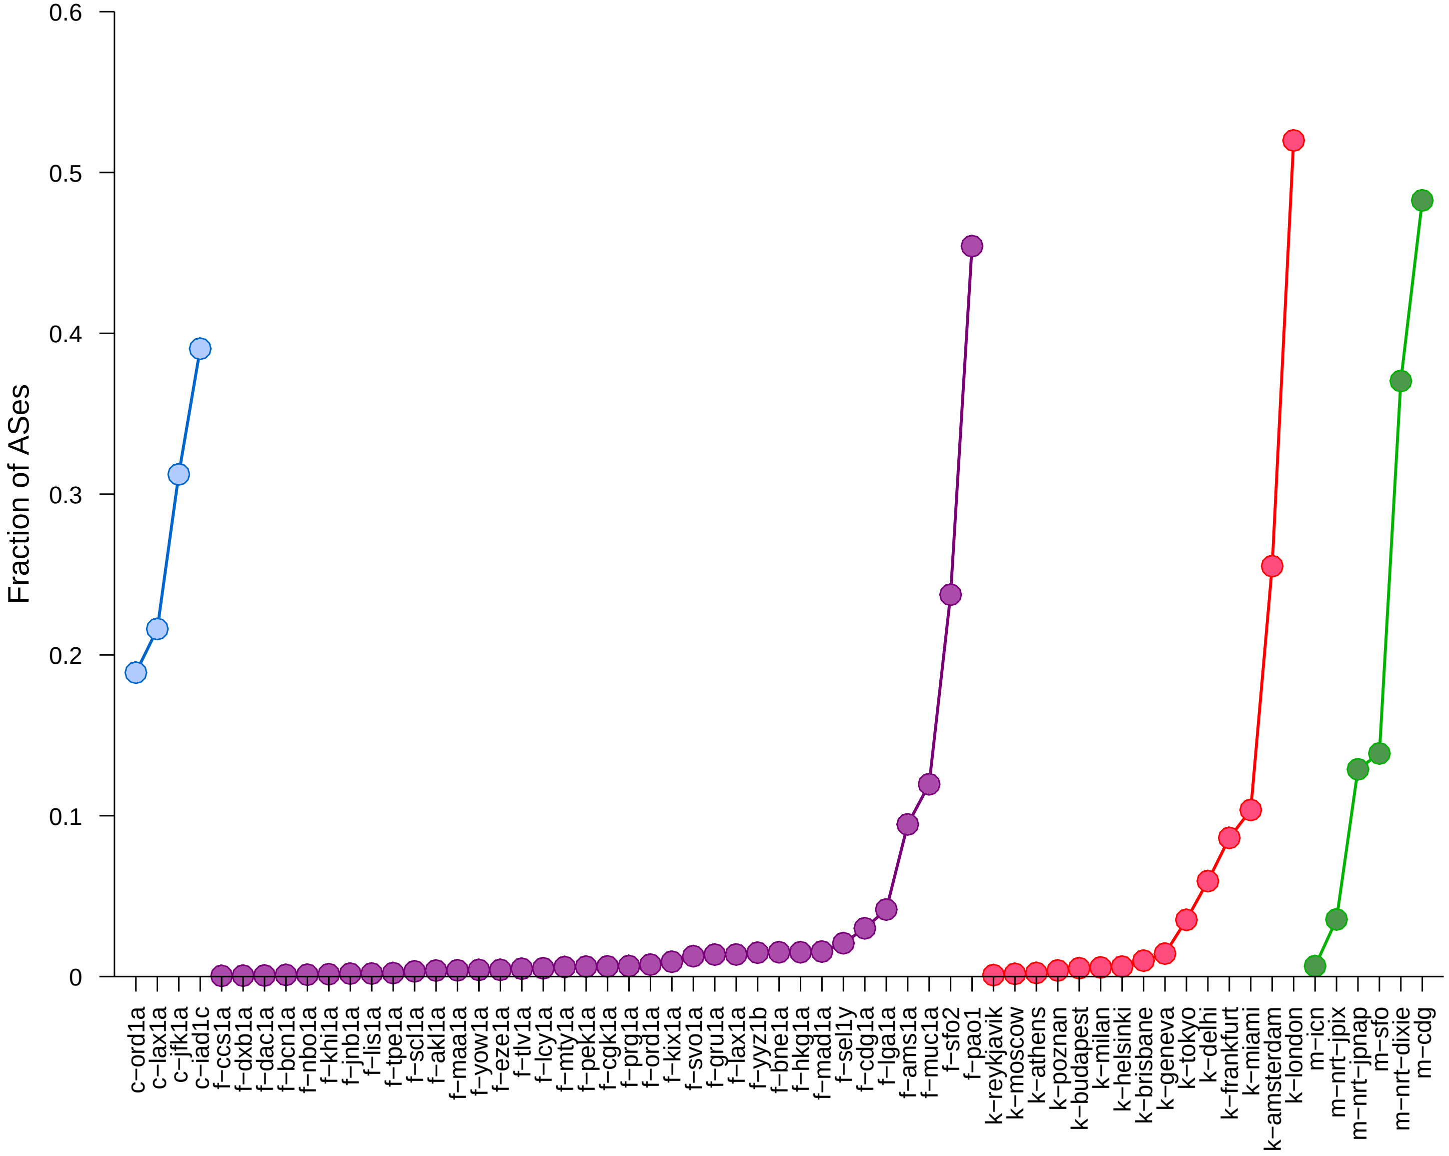

II 4) Topological coverage by prefixes.

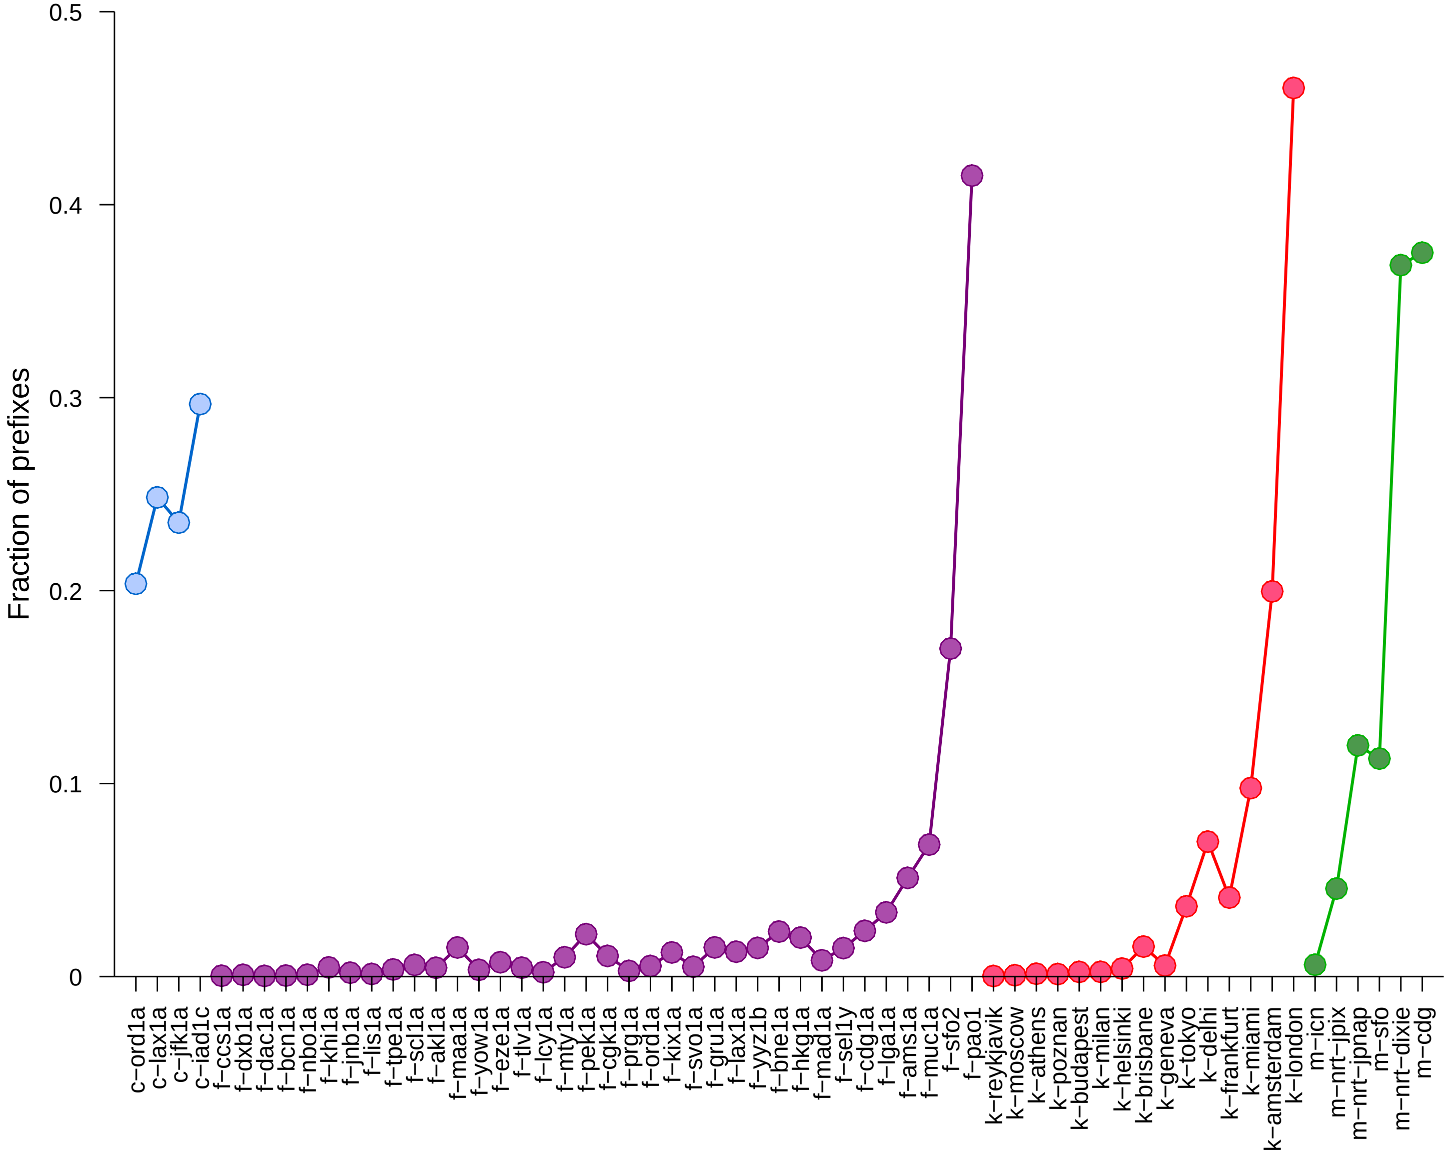

III 1) Clients distribution by RIR for each instance

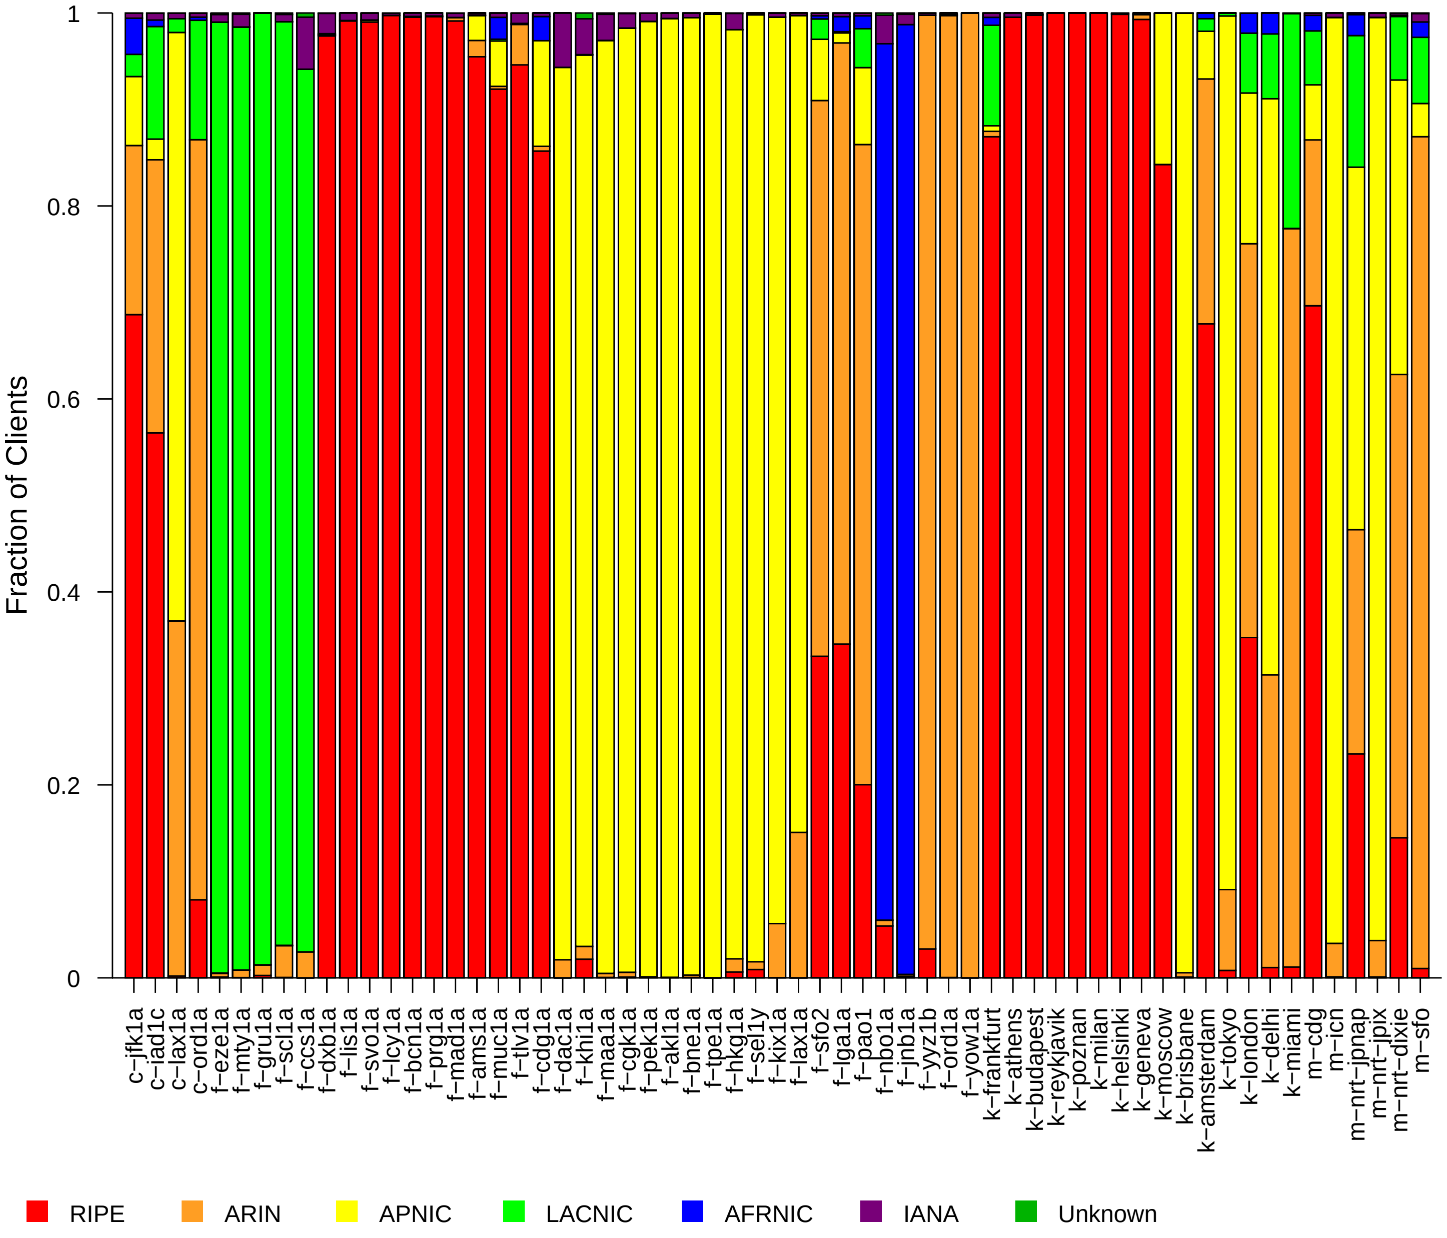

IV 1) Distribution of users binned by query rate intervals for C−root.

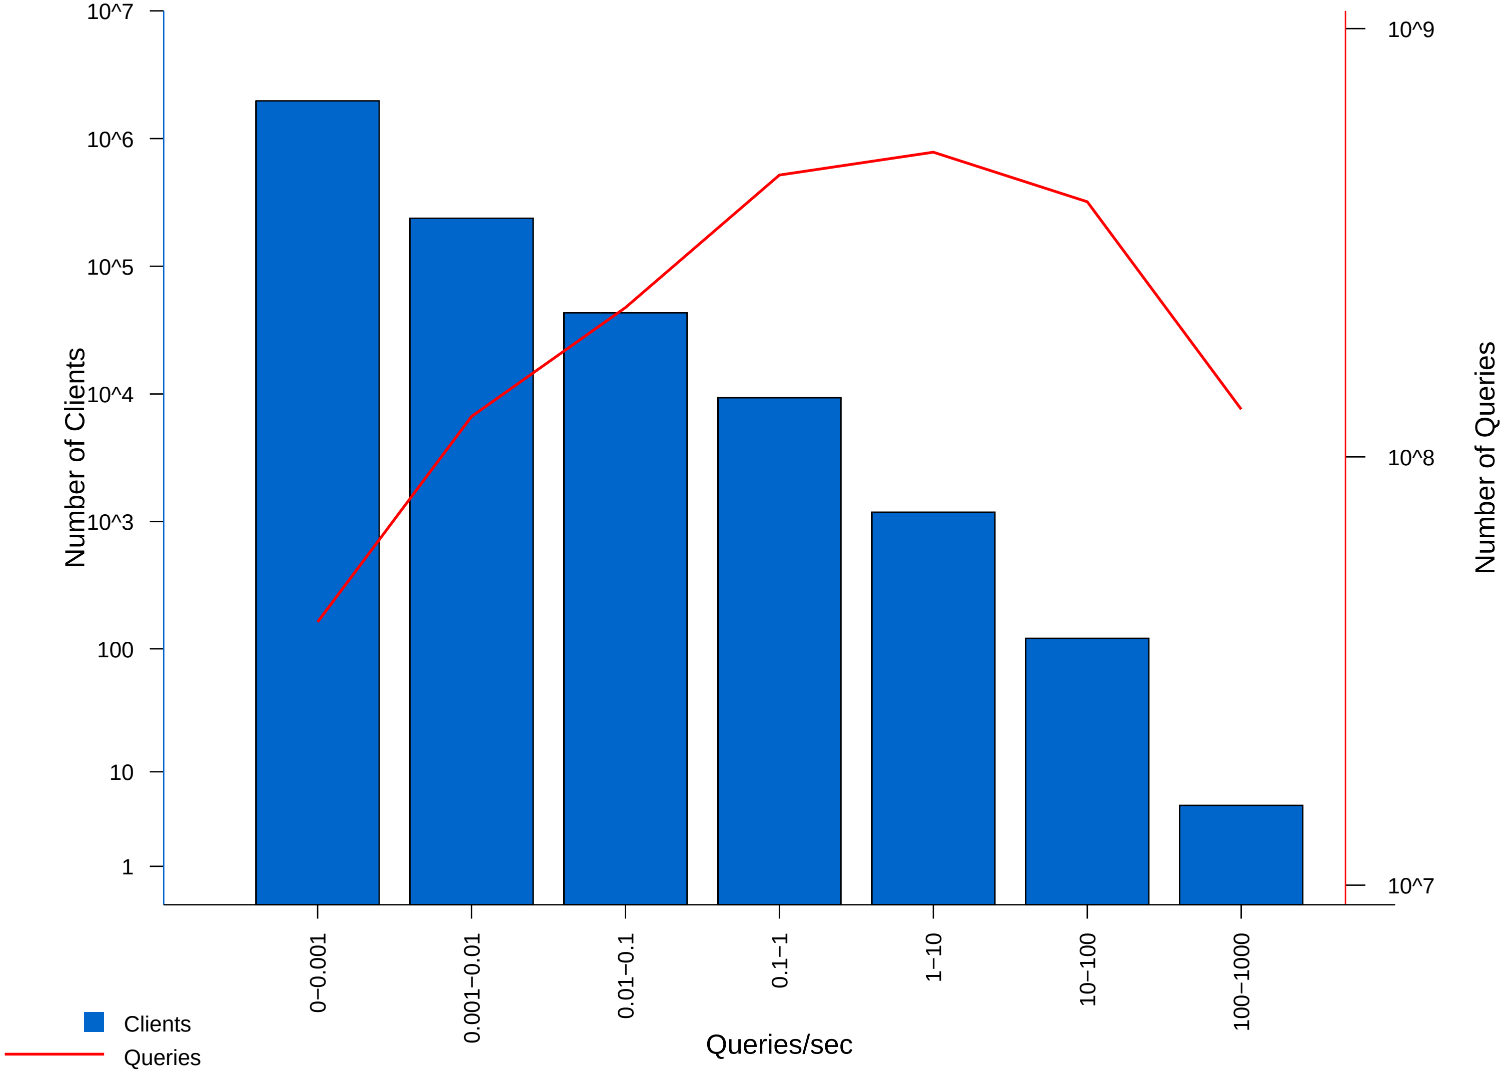

IV 1) Distribution of users binned by query rate intervals for F−root.

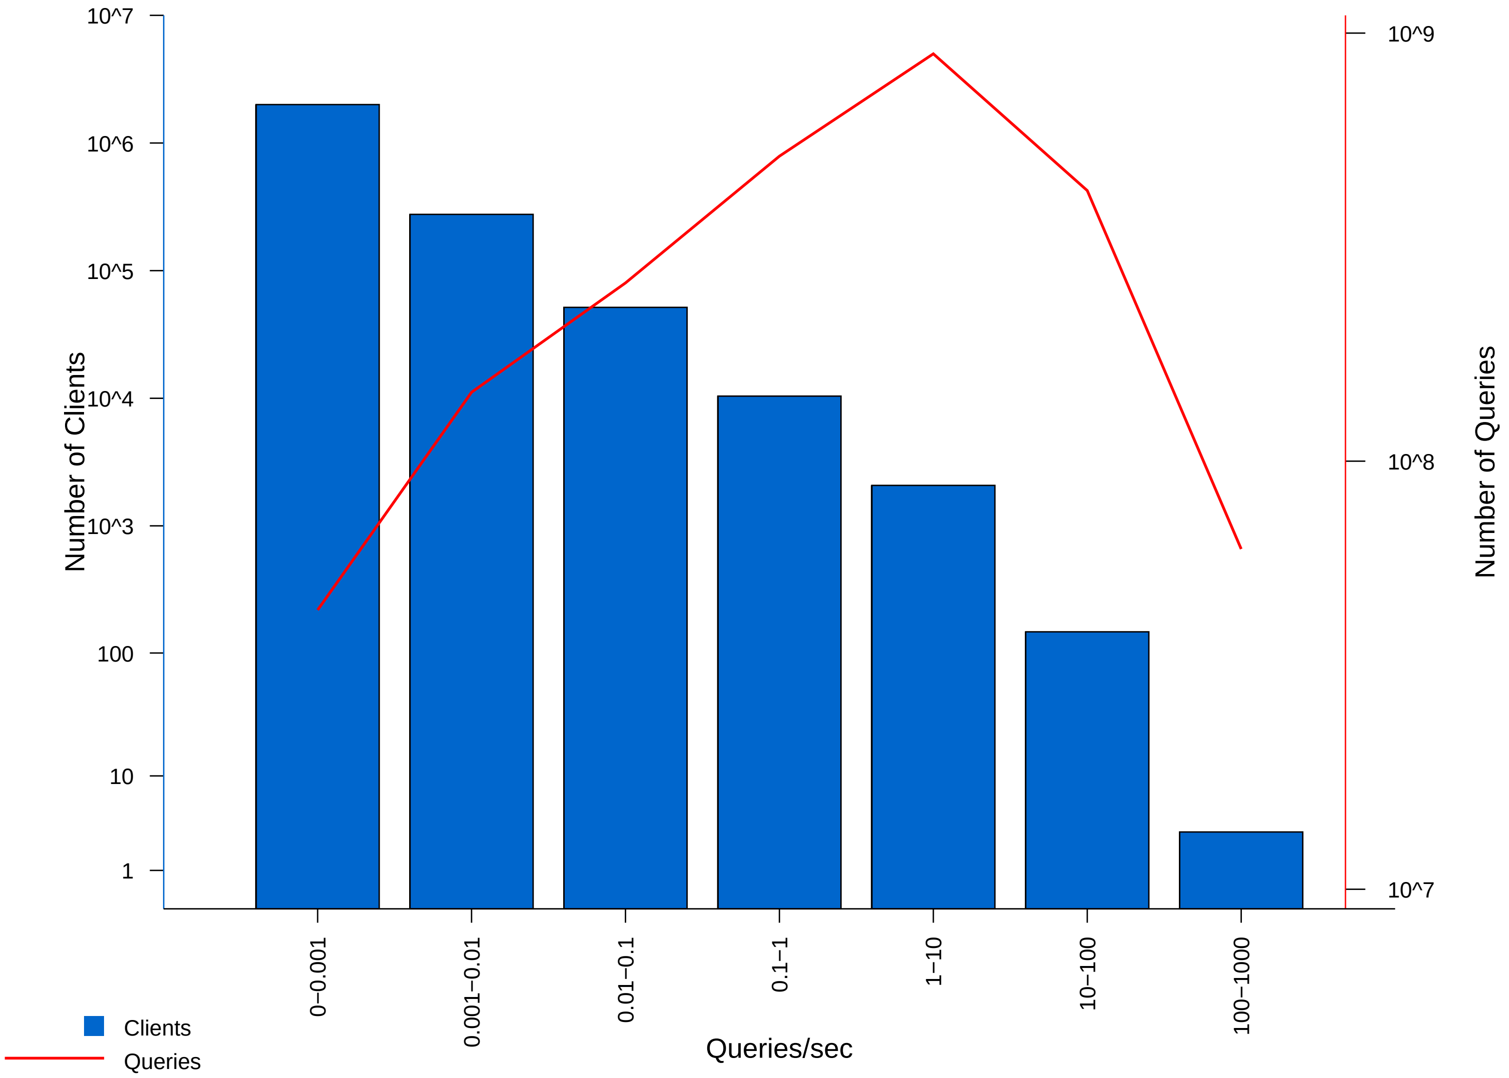

IV 1) Distribution of users binned by query rate intervals for K−root.

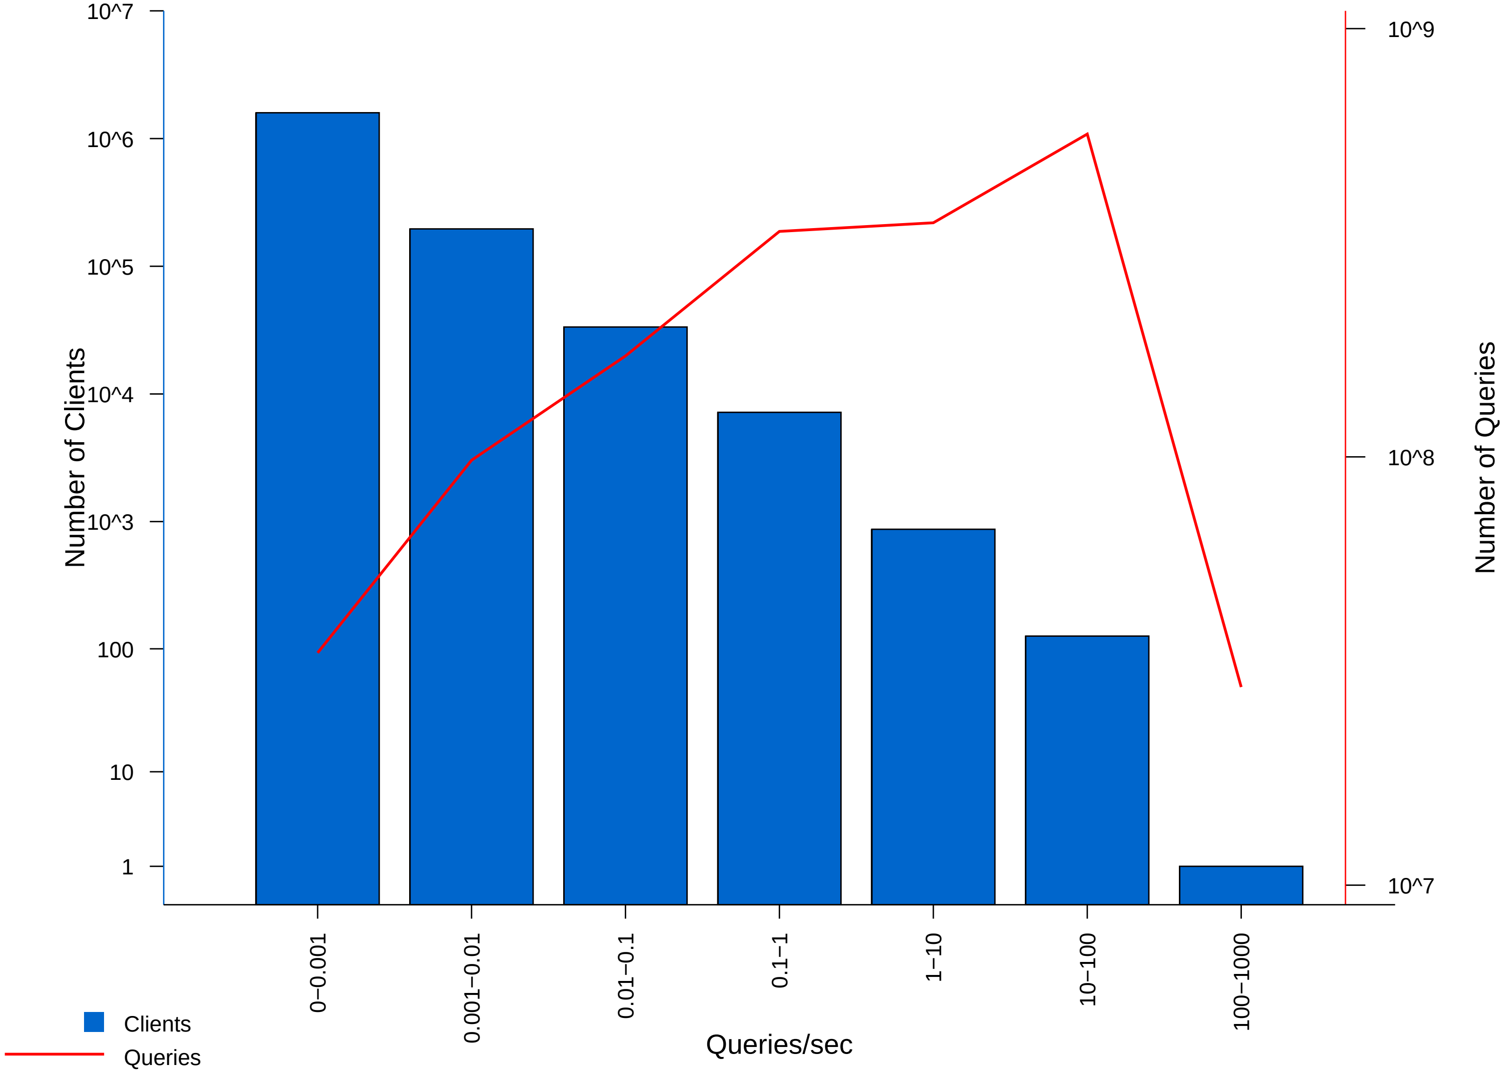

IV 1) Distribution of users binned by query rate intervals for M−root.

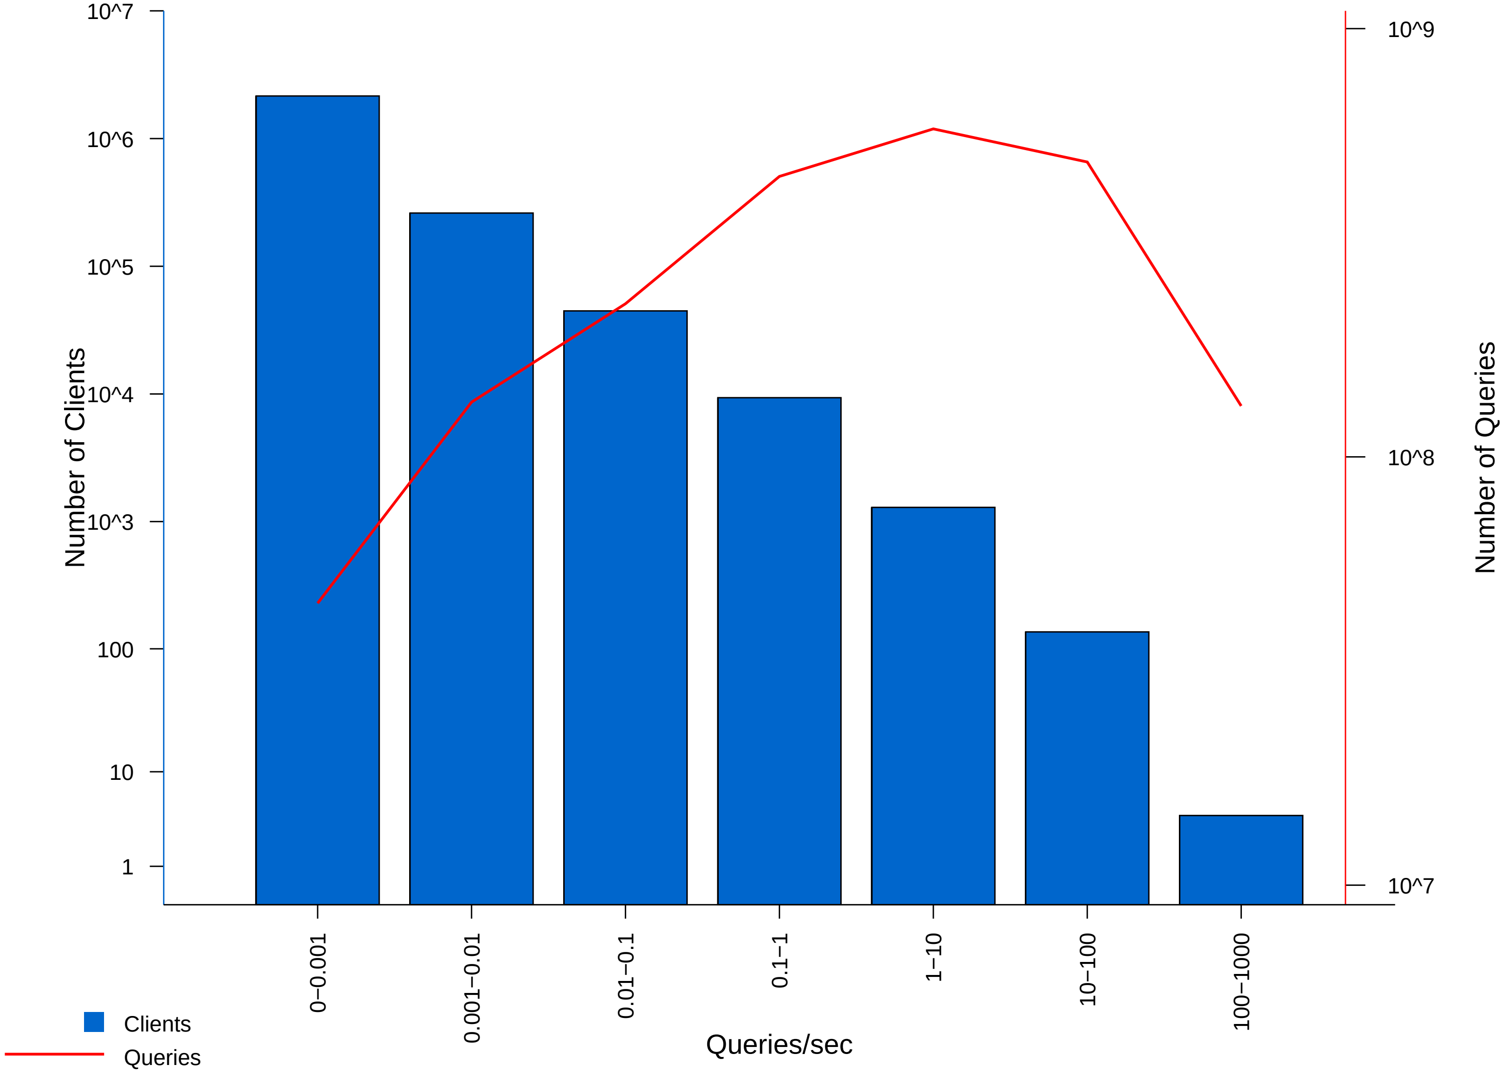

IV 3) Breakdown by query types

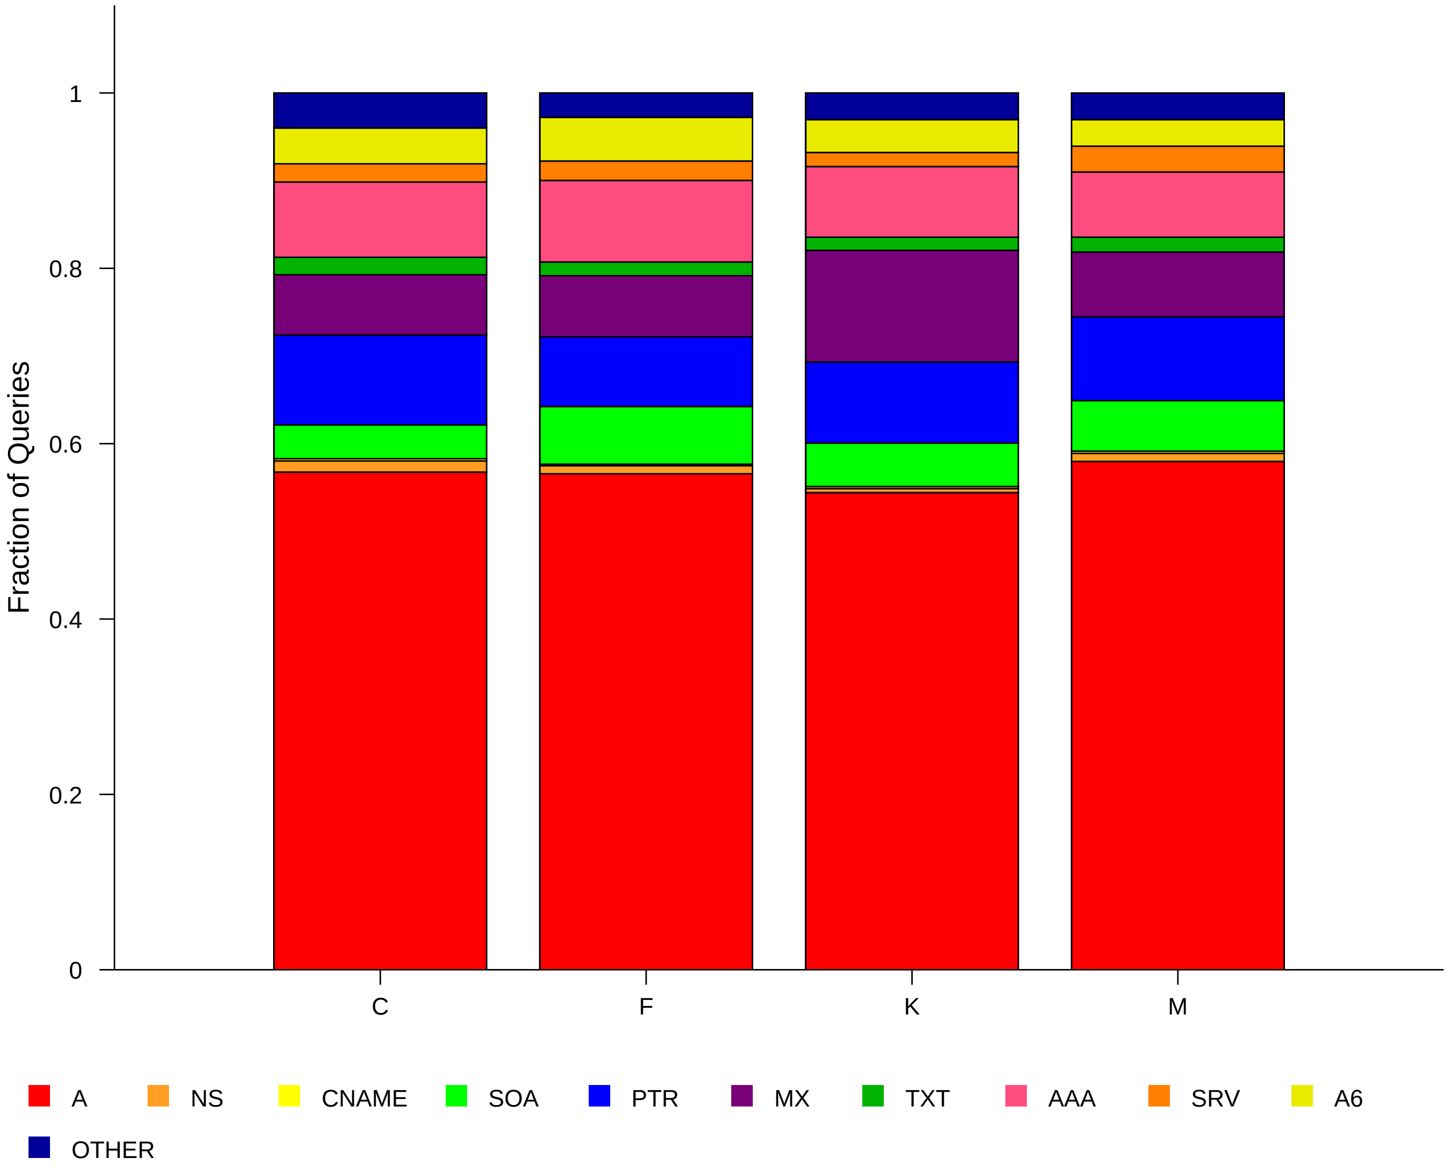

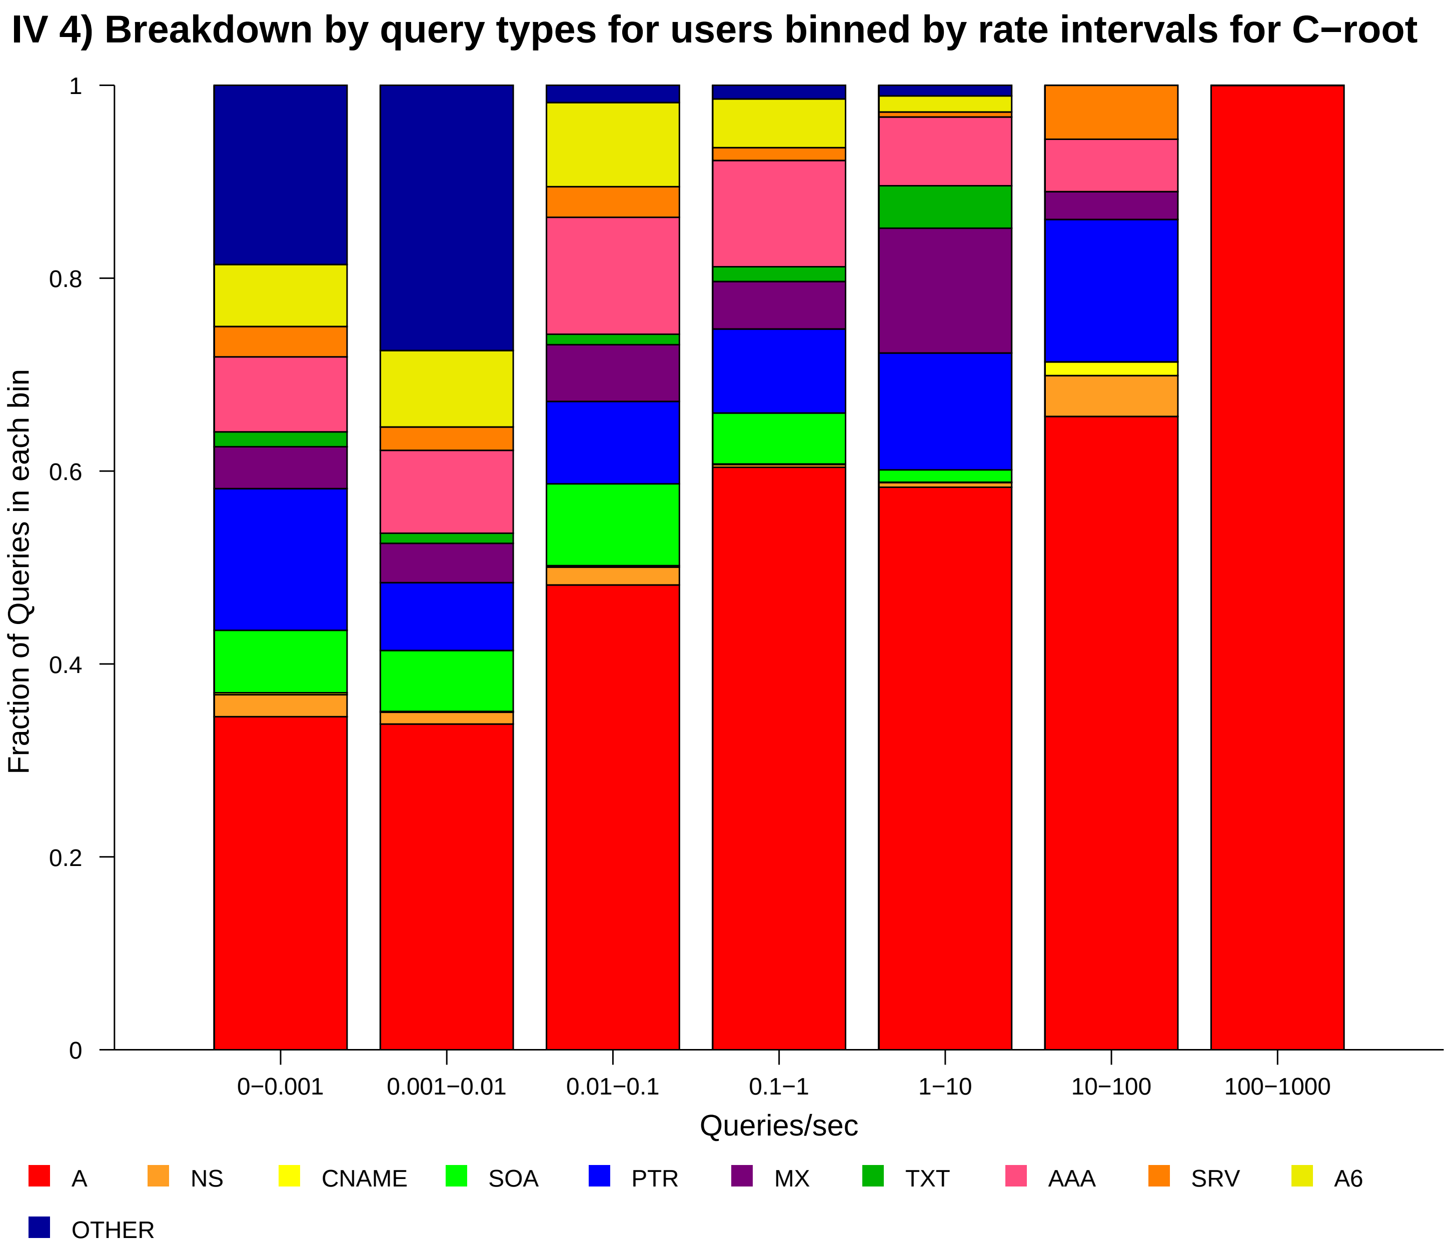

IV 4) Breakdown by query types for users binned by rate intervals for C−root

IV 4) Breakdown by query types for users binned by rate intervals for F−root

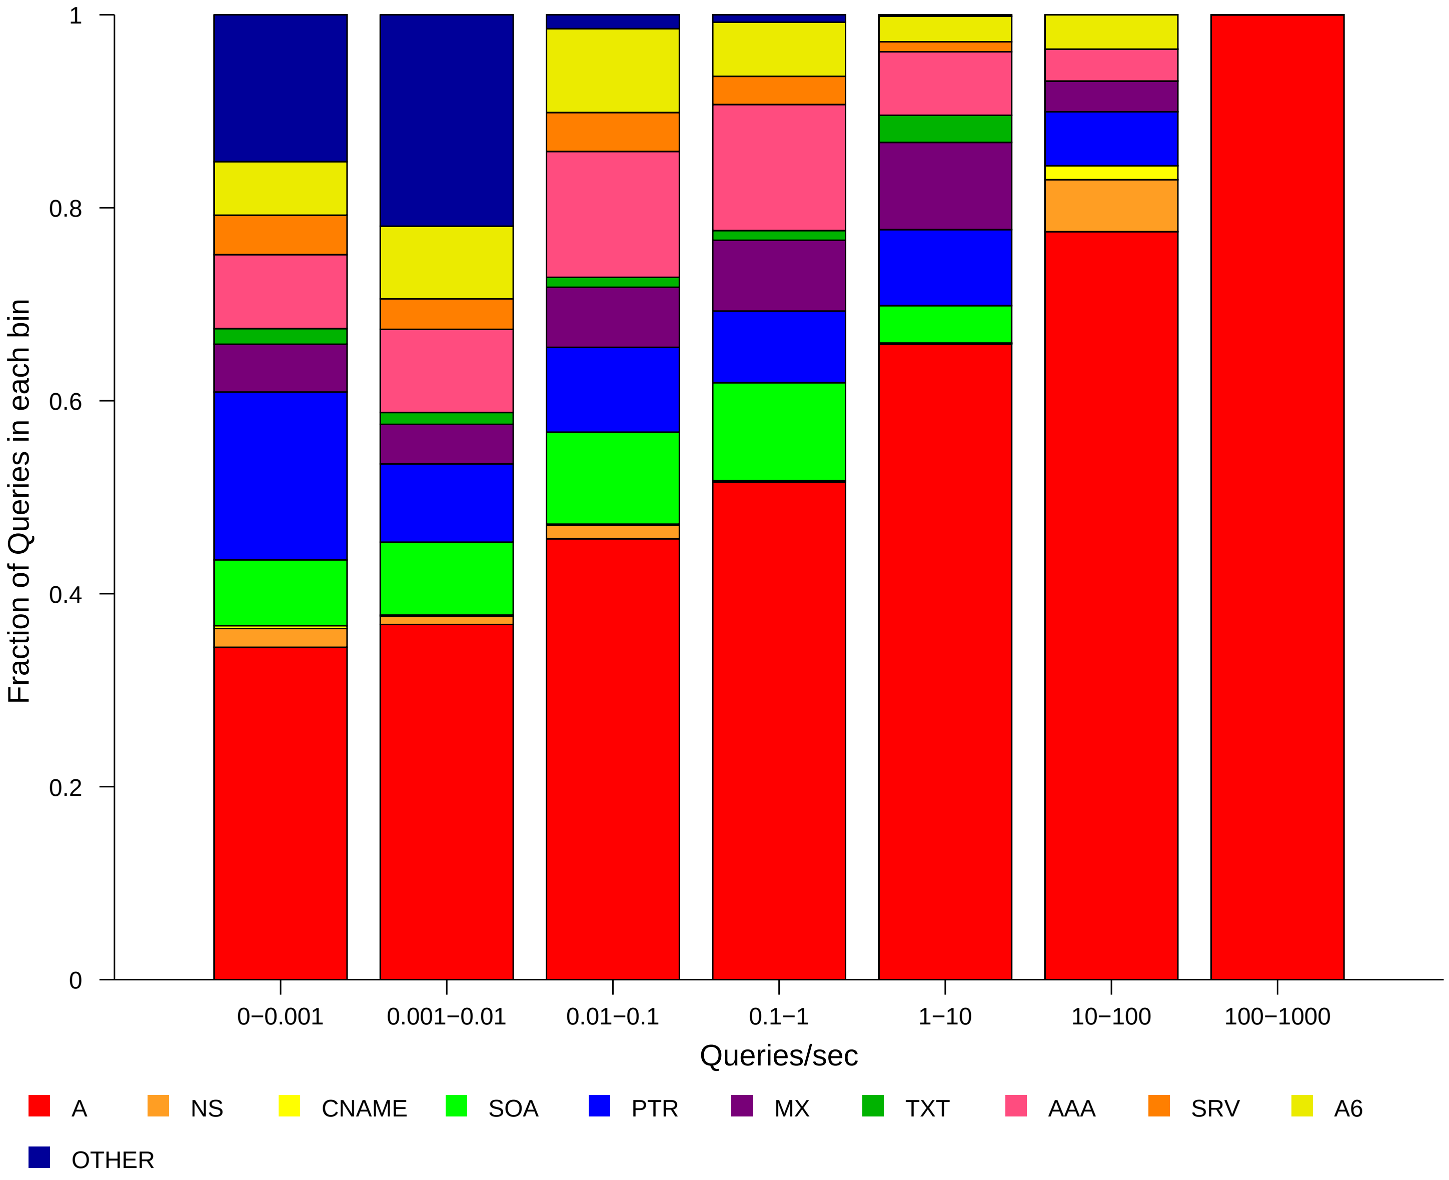

IV 4) Breakdown by query types for users binned by rate intervals for K−root

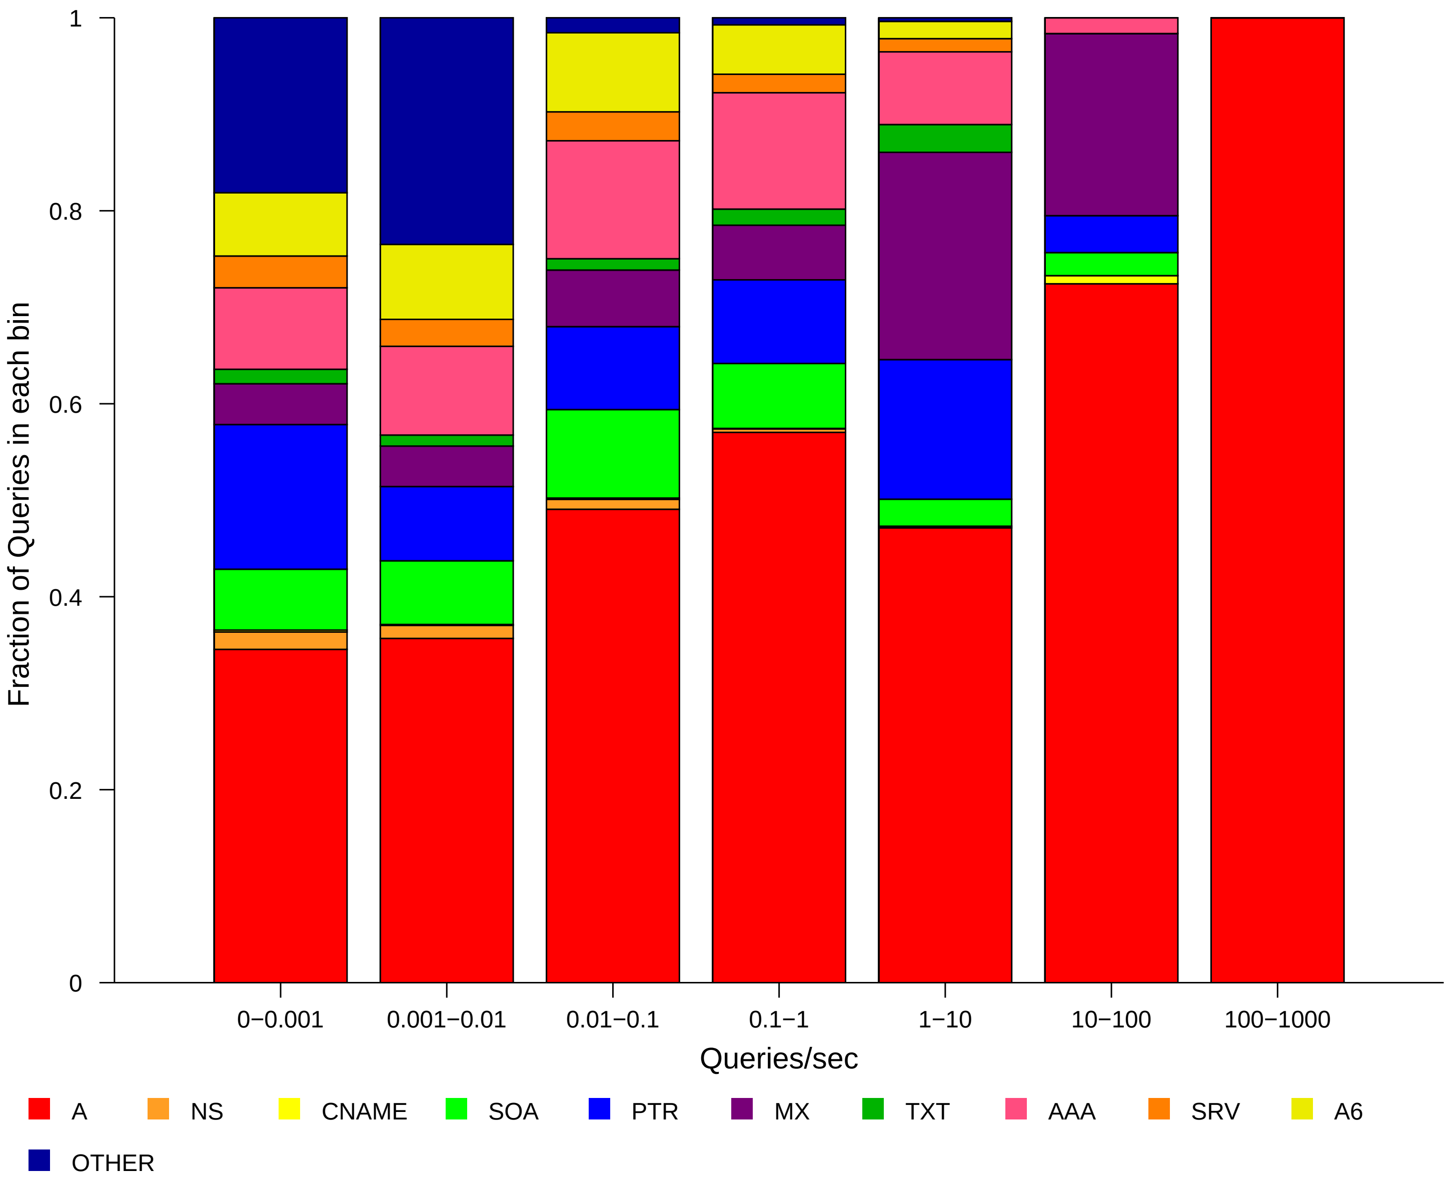

IV 4) Breakdown by query types for users binned by rate intervals for M−root

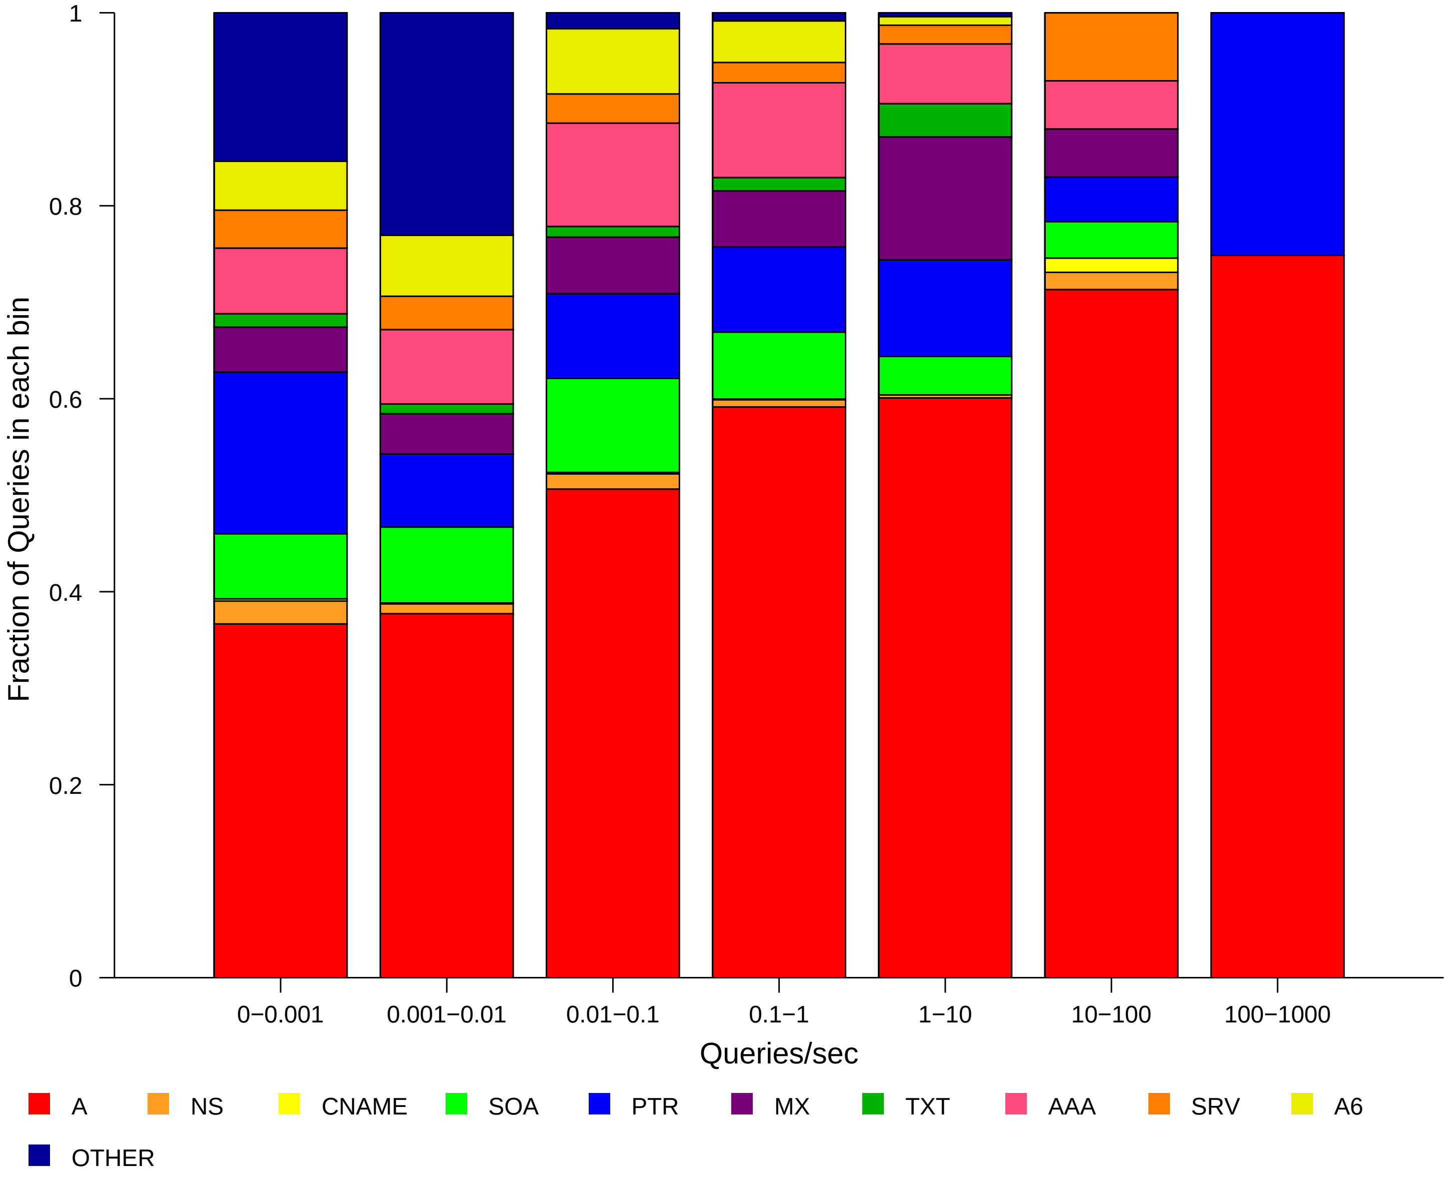

The End