DNS

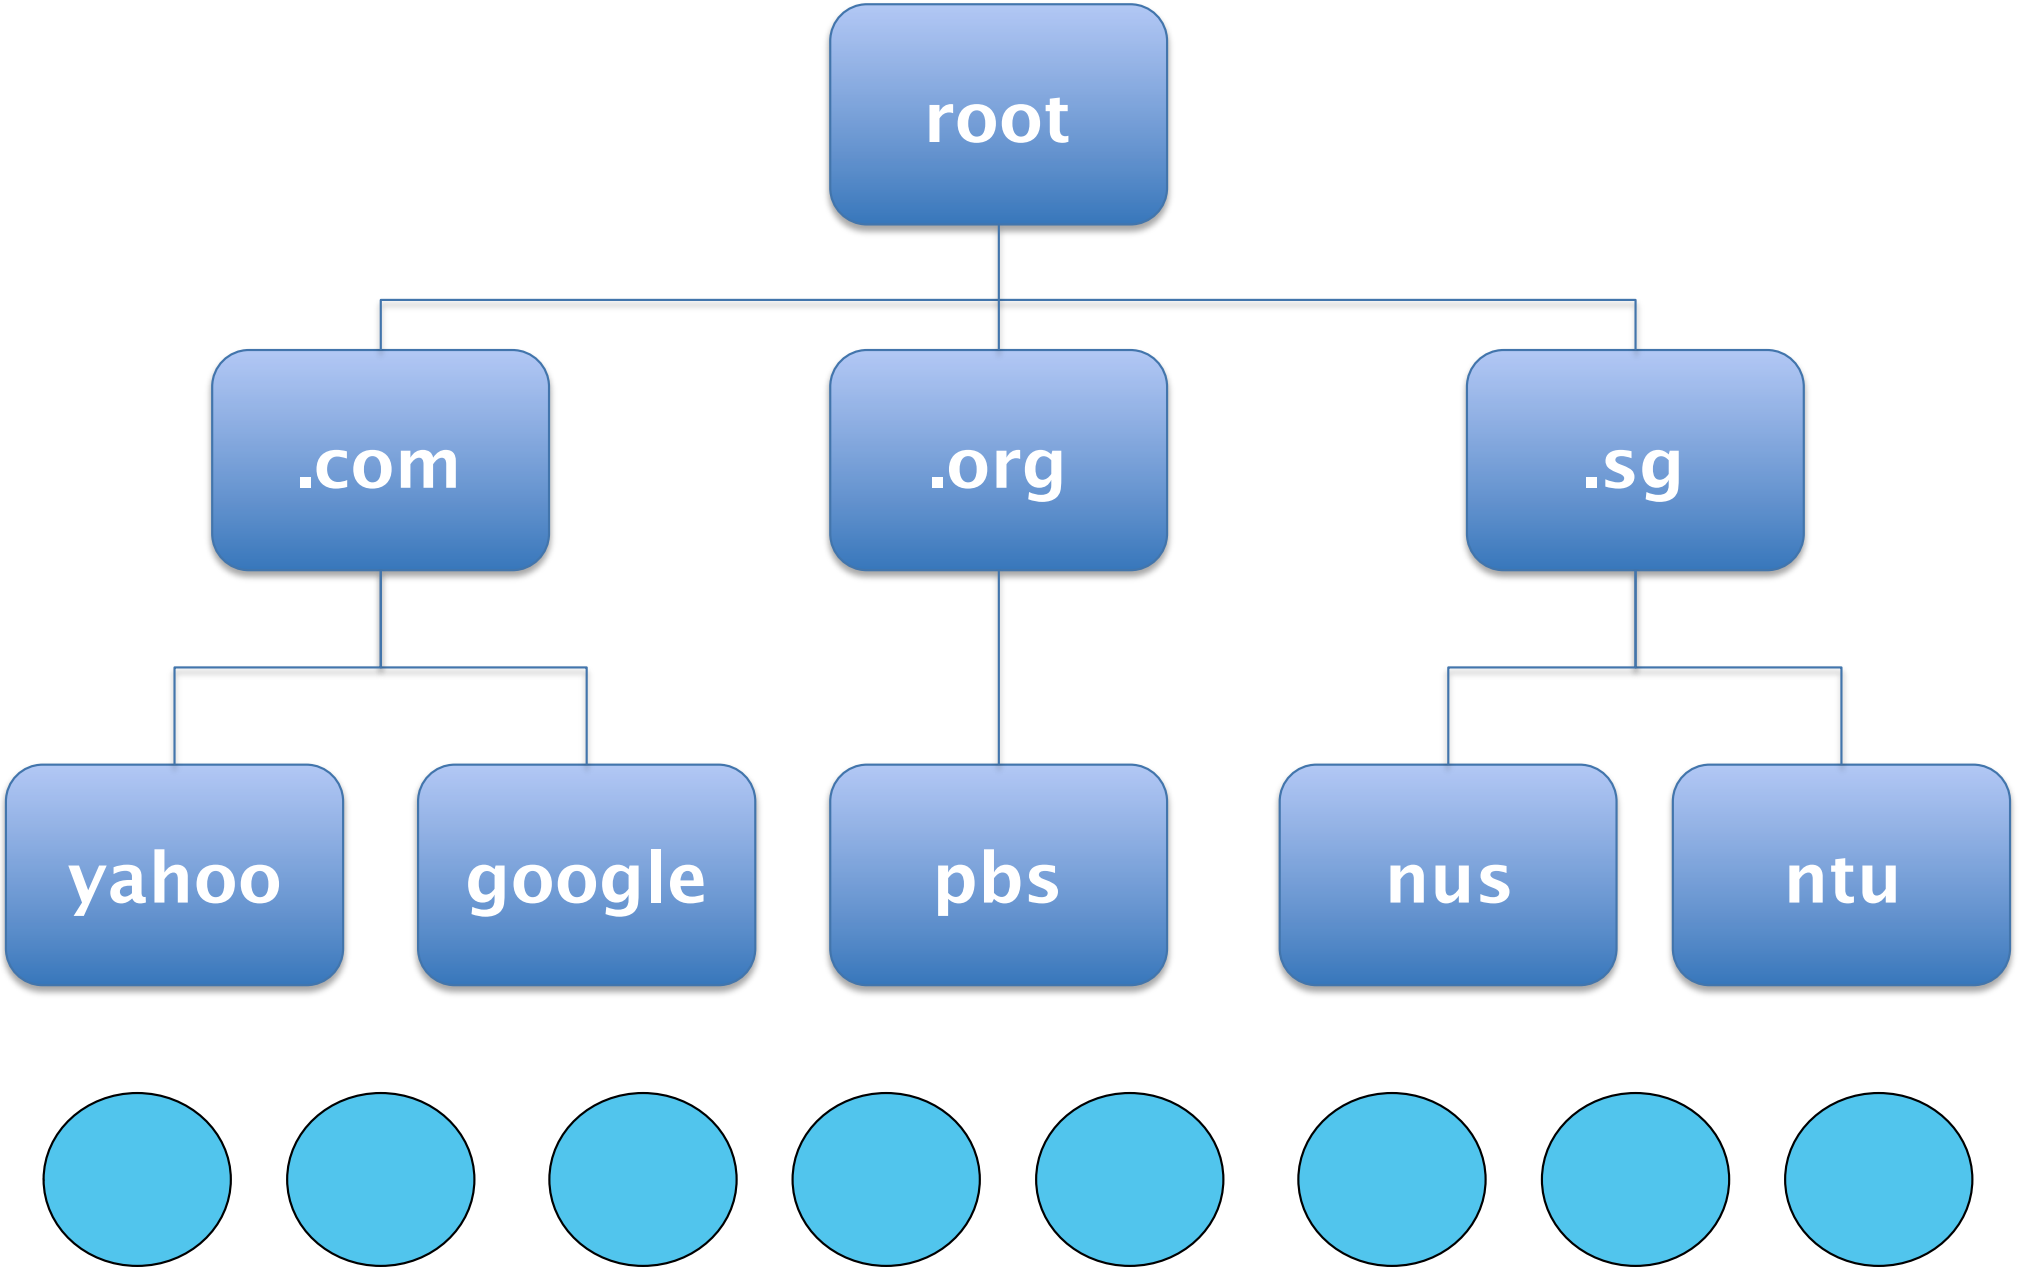
root
.com
.org
.sg
yahoo
google
pbs
nus
ntu

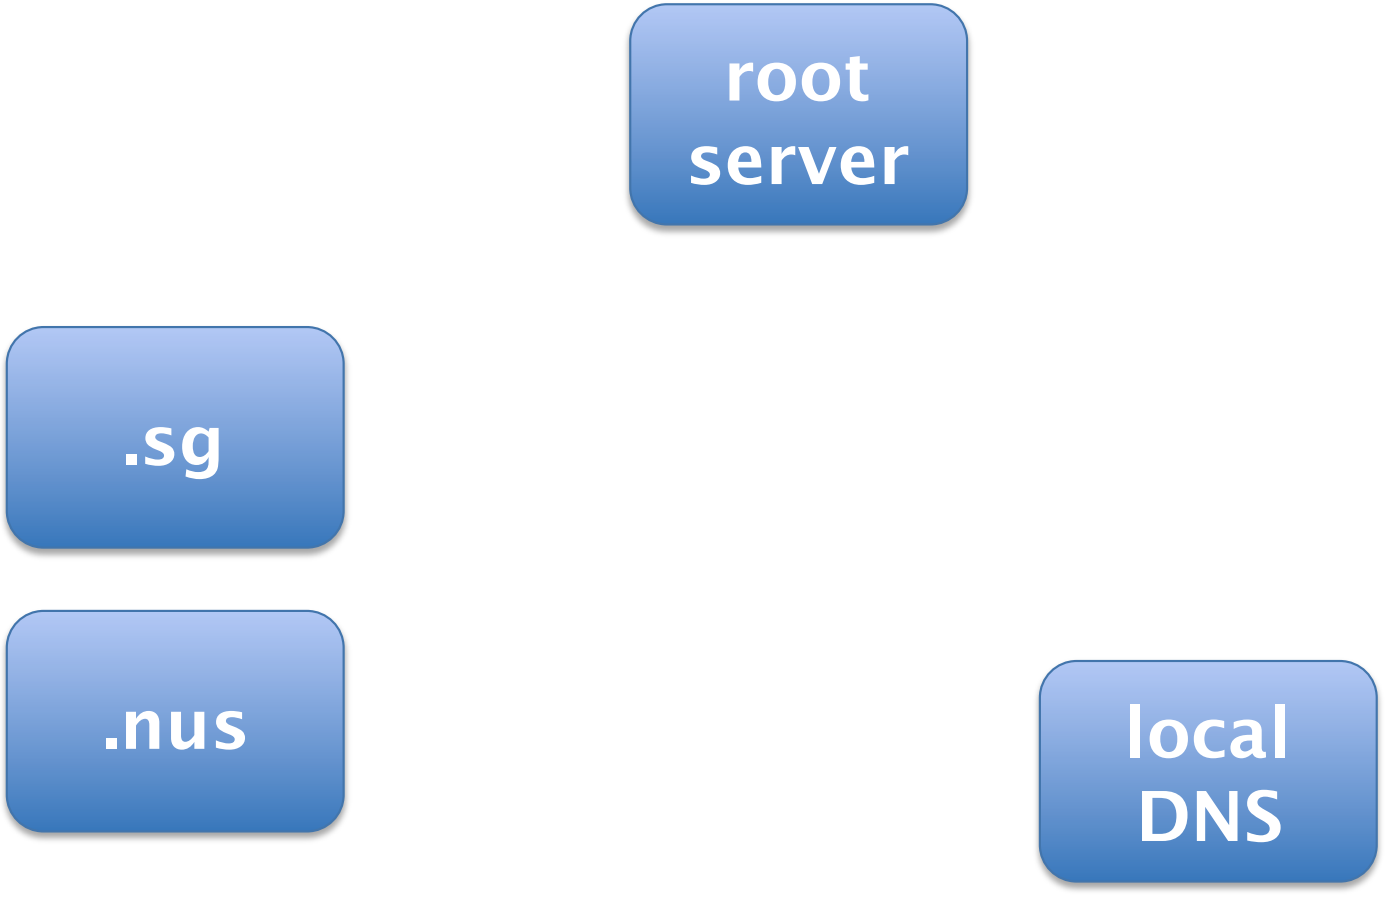


```
dig +trace www.nus.sg
```

uses **UDP**

Resource Record

(name, TTL, class, type, value)

```
ns1.nus.edu.sg.   6562   IN   A    137.132.123.4
```

TTL-based Caching

Resource Type

A NS PTR MX CNAME …

```
ns1.nus.edu.sg.  6562  IN  A   137.132.123.4
```

DNS-based Load Balancing

DNS-based Server Selection

how fast?
how efficient?
is caching useful?
how big should TTL be?

"DNS Performance and Effectiveness of Caching"

J Jung et. al.

IEEE TON 2002

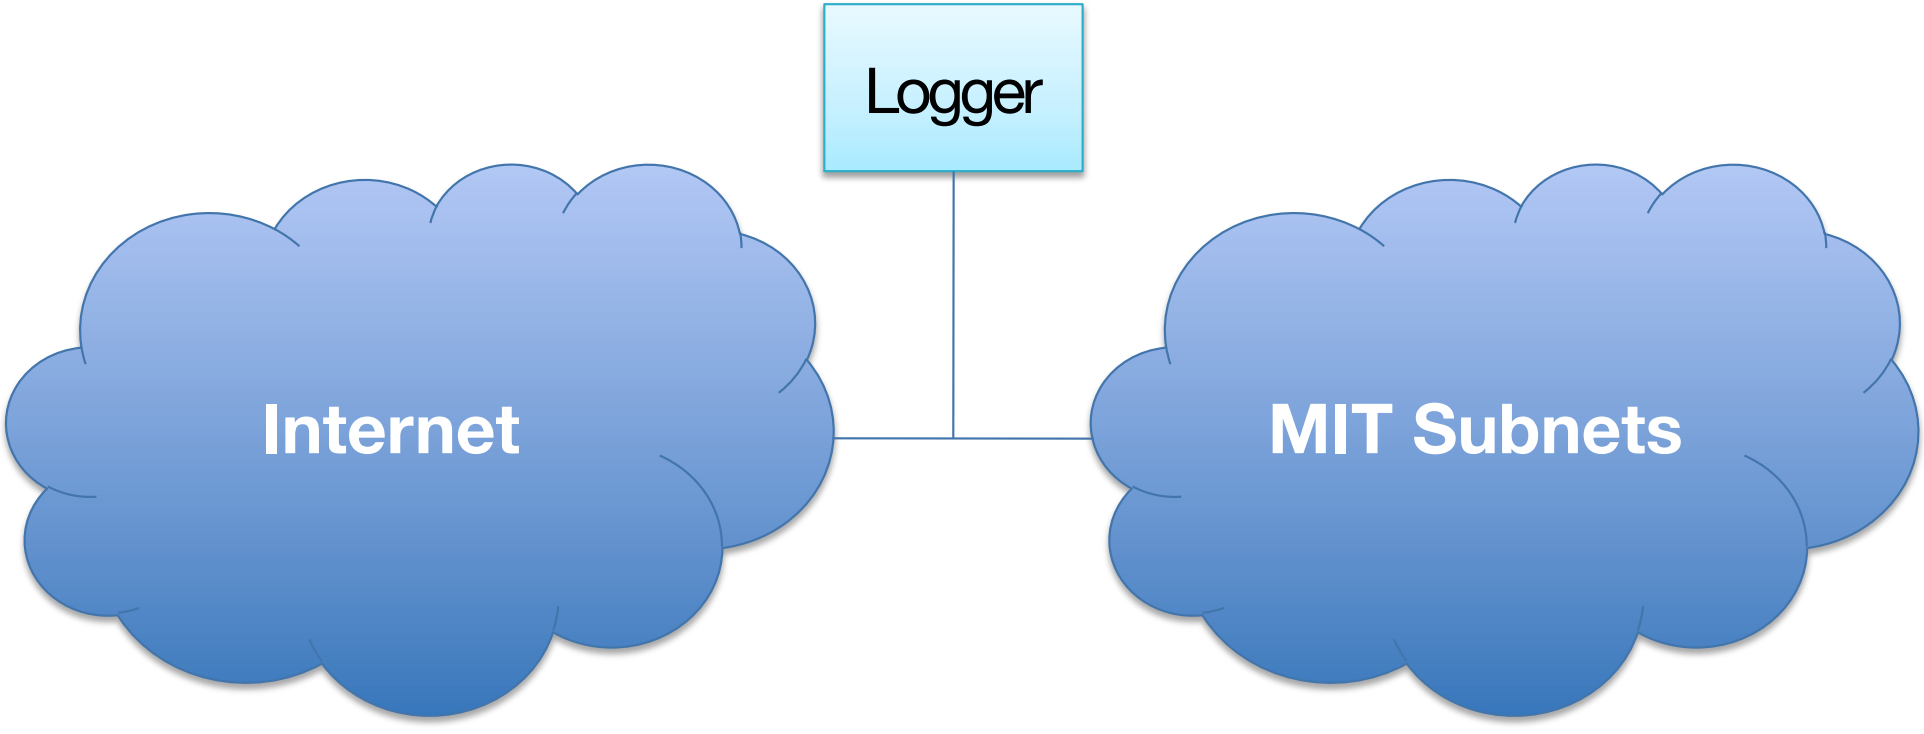
Logger
Internet
MIT Subnets

3 traces

1 week each

DNS query/respose

TCP SYN/FIN/RST

basic analysis

DNS Query Types for mit-jan00

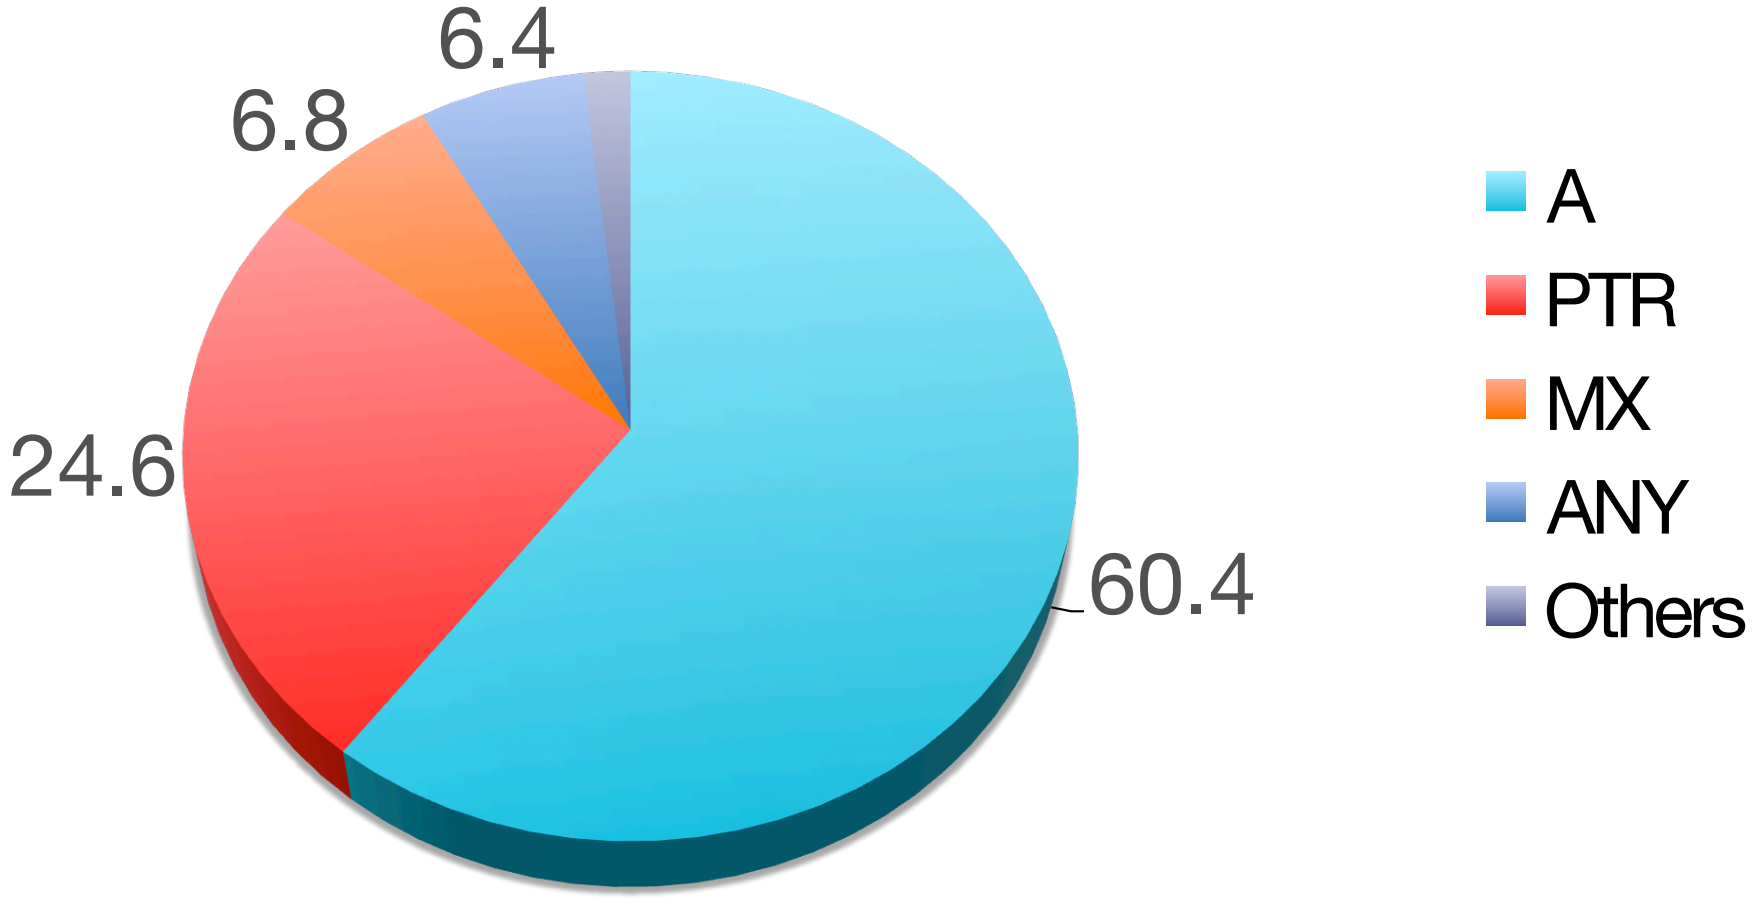

496,802

lookups followed by TCP connections

3,619,173

TCP connections

80-87%

DNS cache hit ratio
for A lookup

not very high since it includes Web browser's caching too

how fast is DNS lookup?

lookup latency ≈
time between first query
and last response

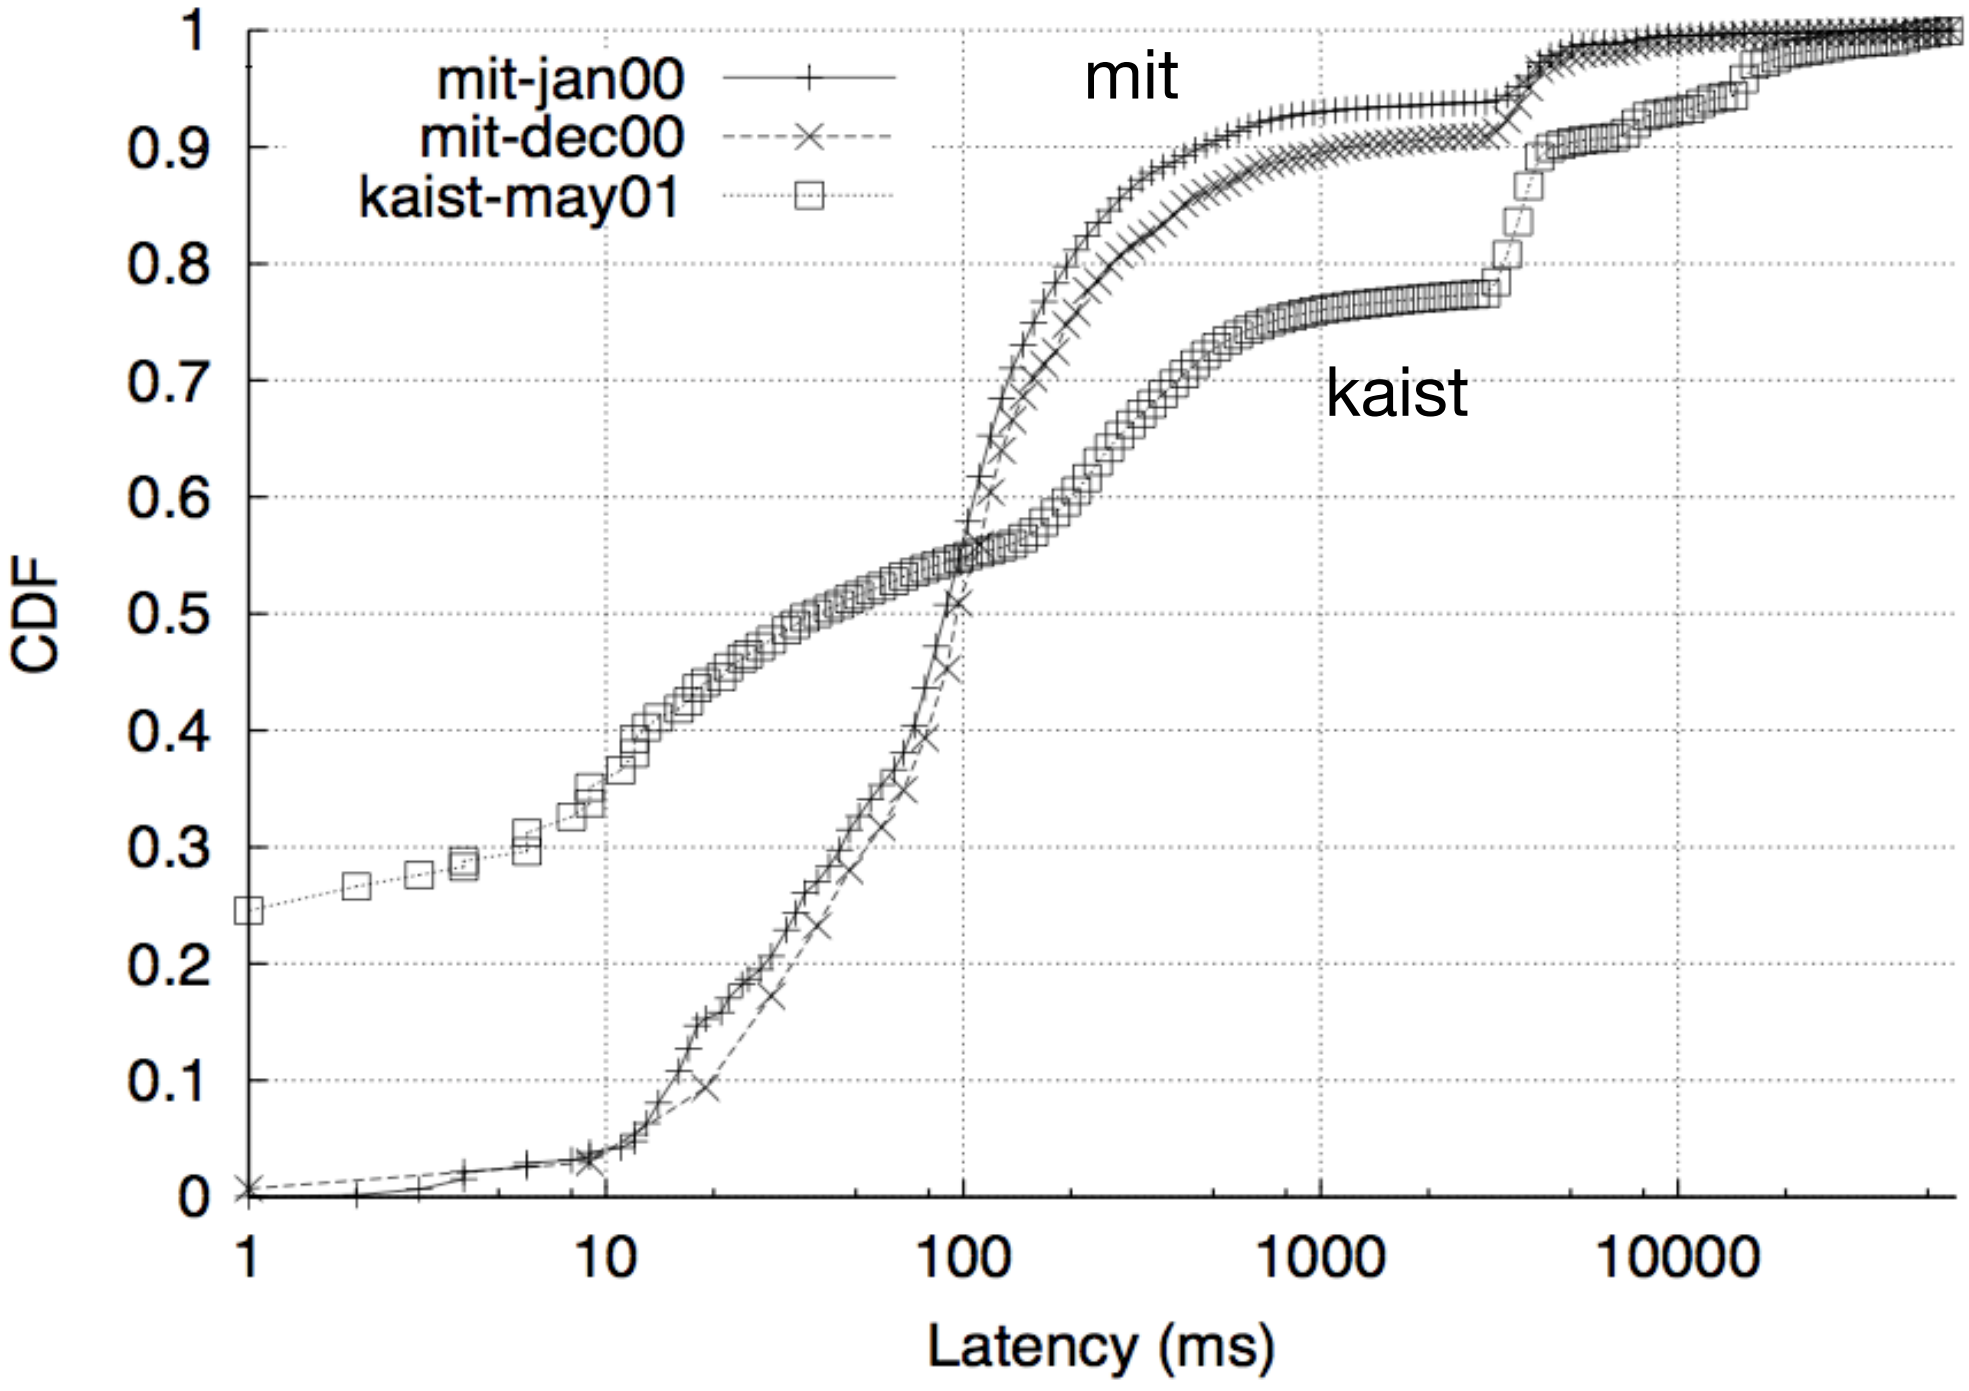
mit-jan00
mit-dec00
kaist-may01
mit
kaist
CDF
Latency (ms)
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
1
10
100
1000
10000

Number of referals per lookup

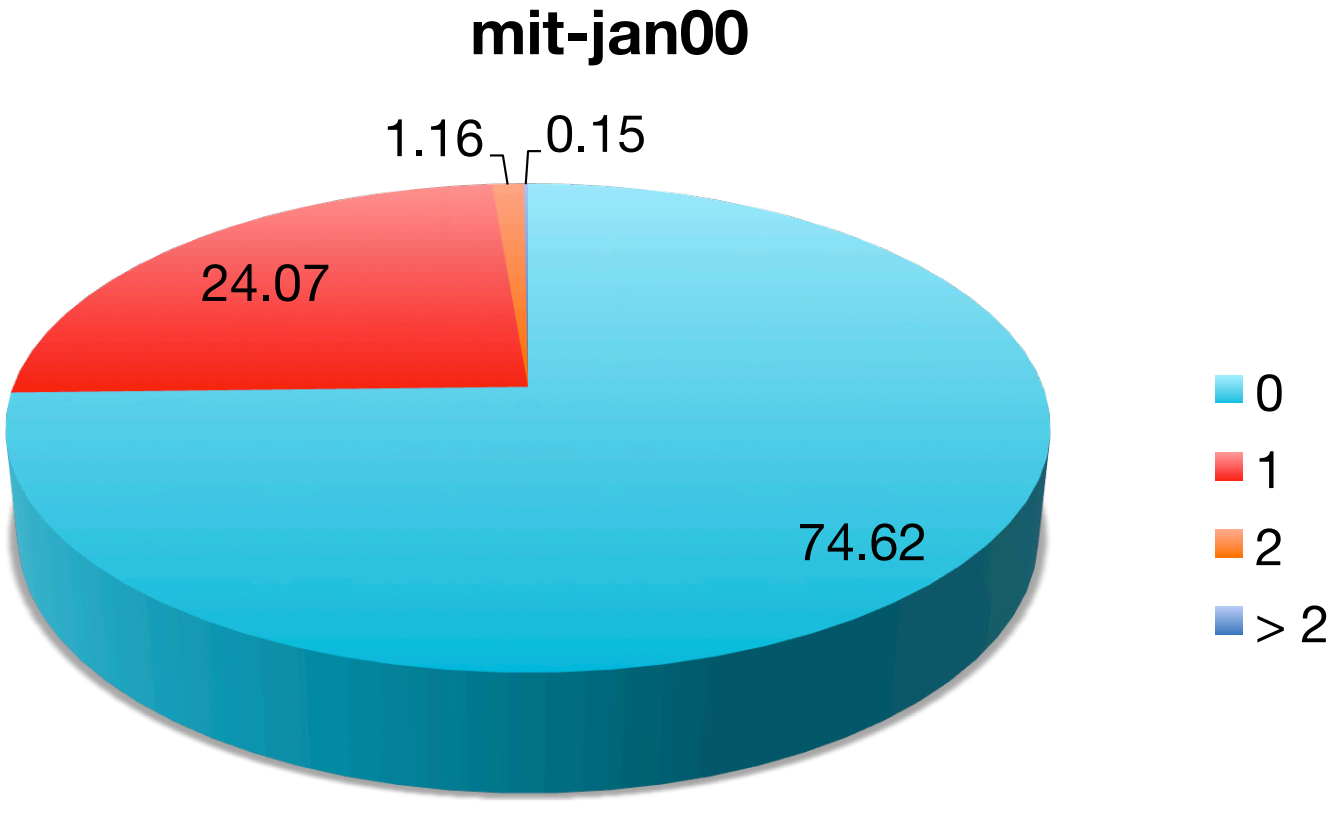

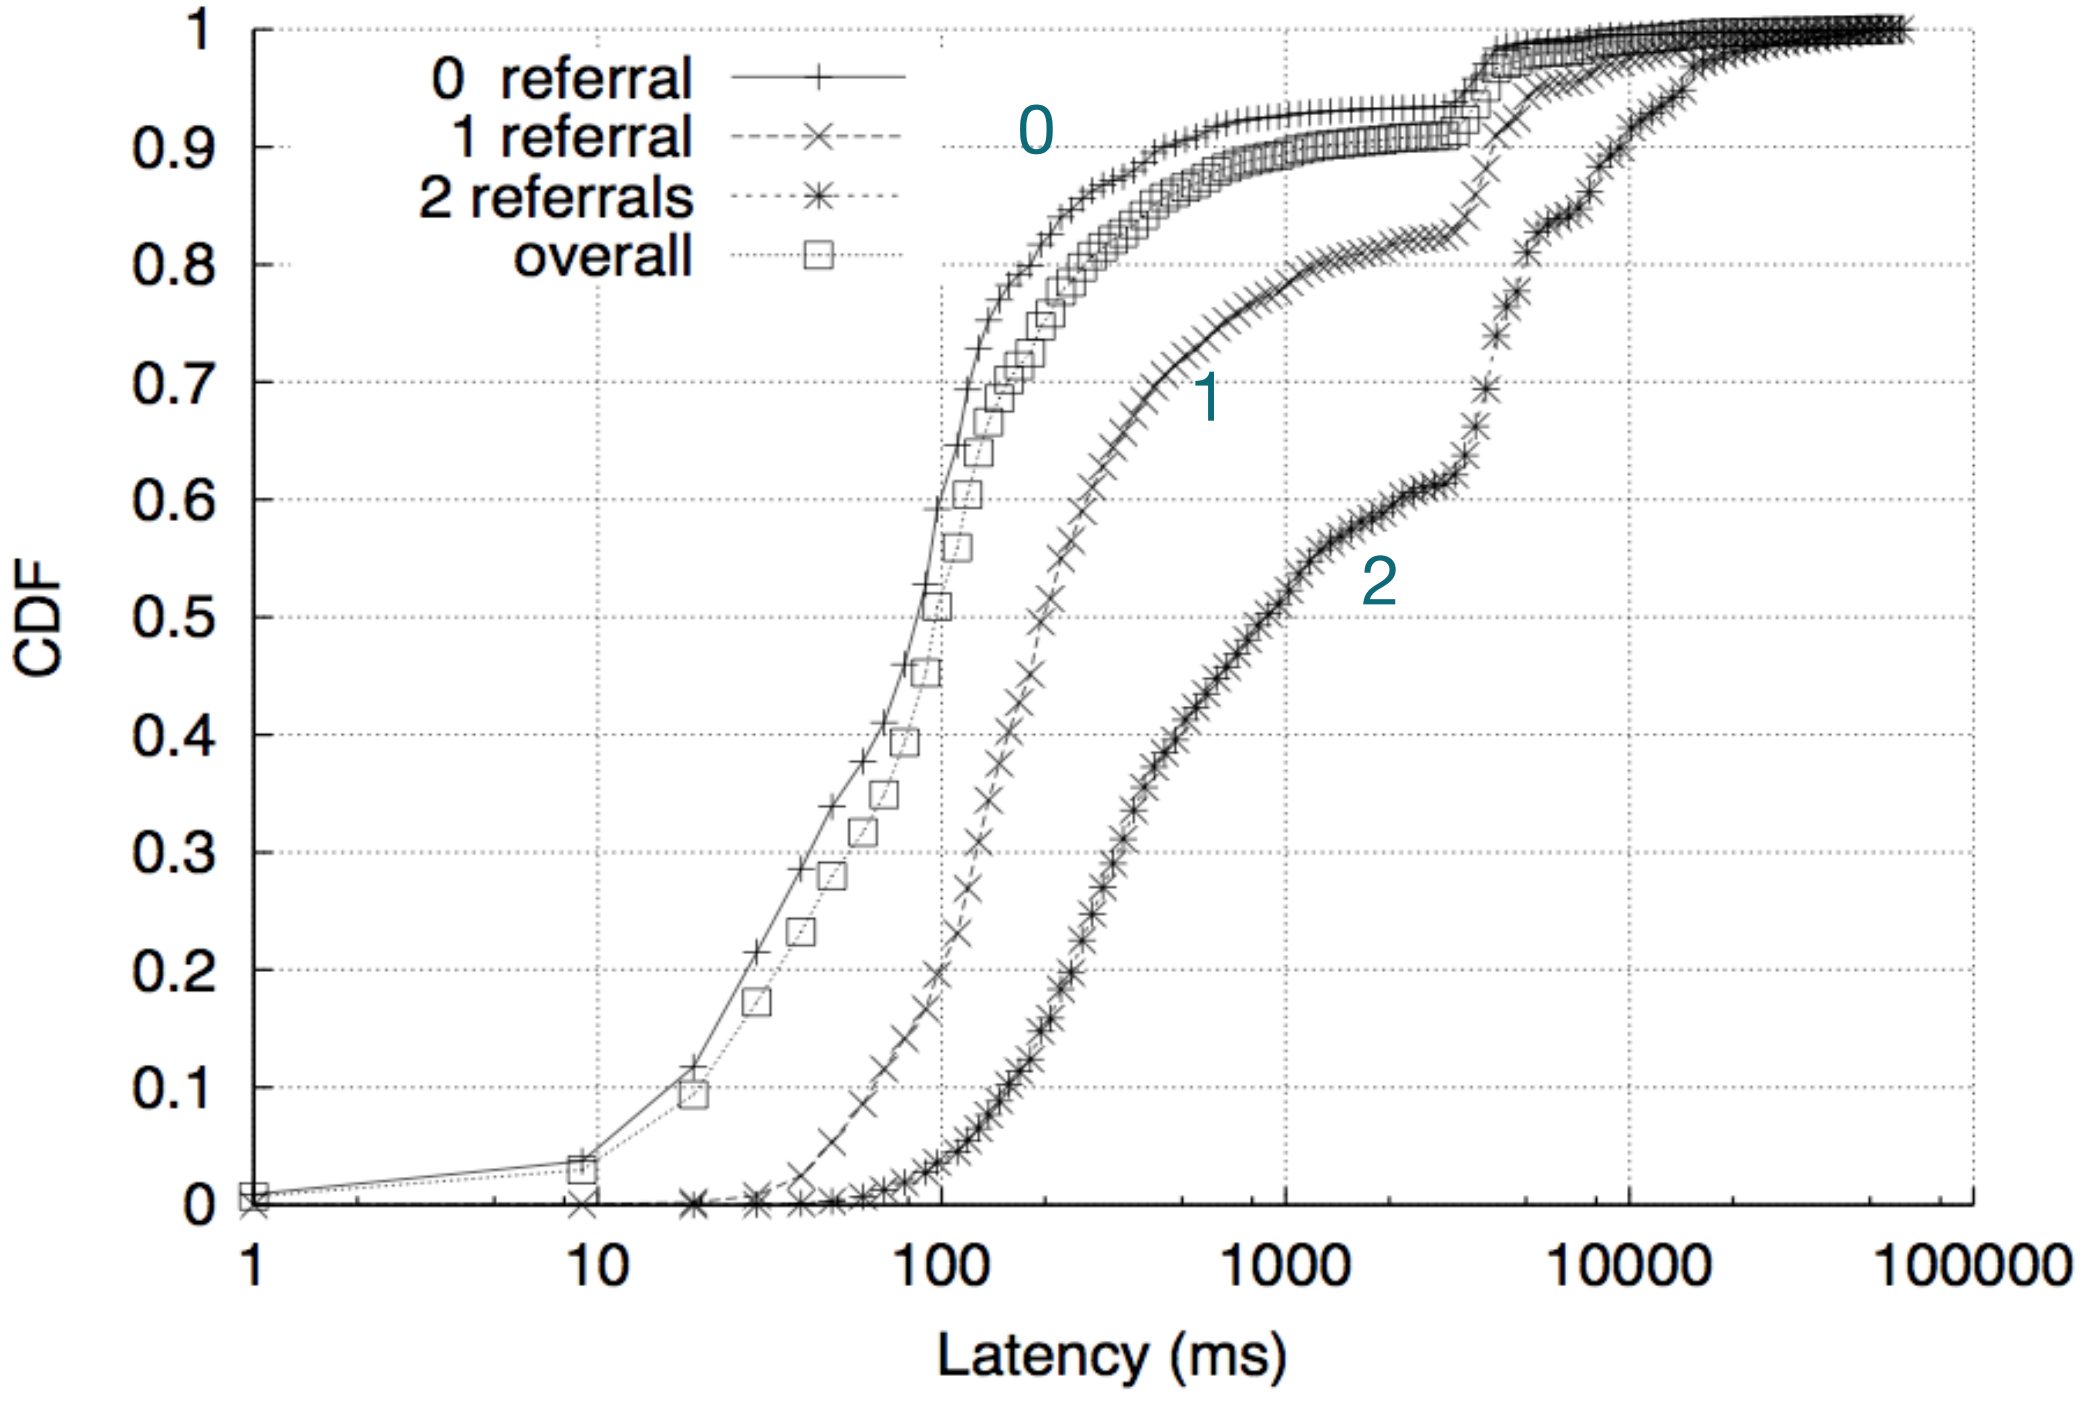
0 referral
1 referral
2 referrals
overall
0
1
2
CDF
Latency (ms)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
10
100
1000
10000
100000

NS Cache

miss: first query to gTLD or root servers

hit: otherwise

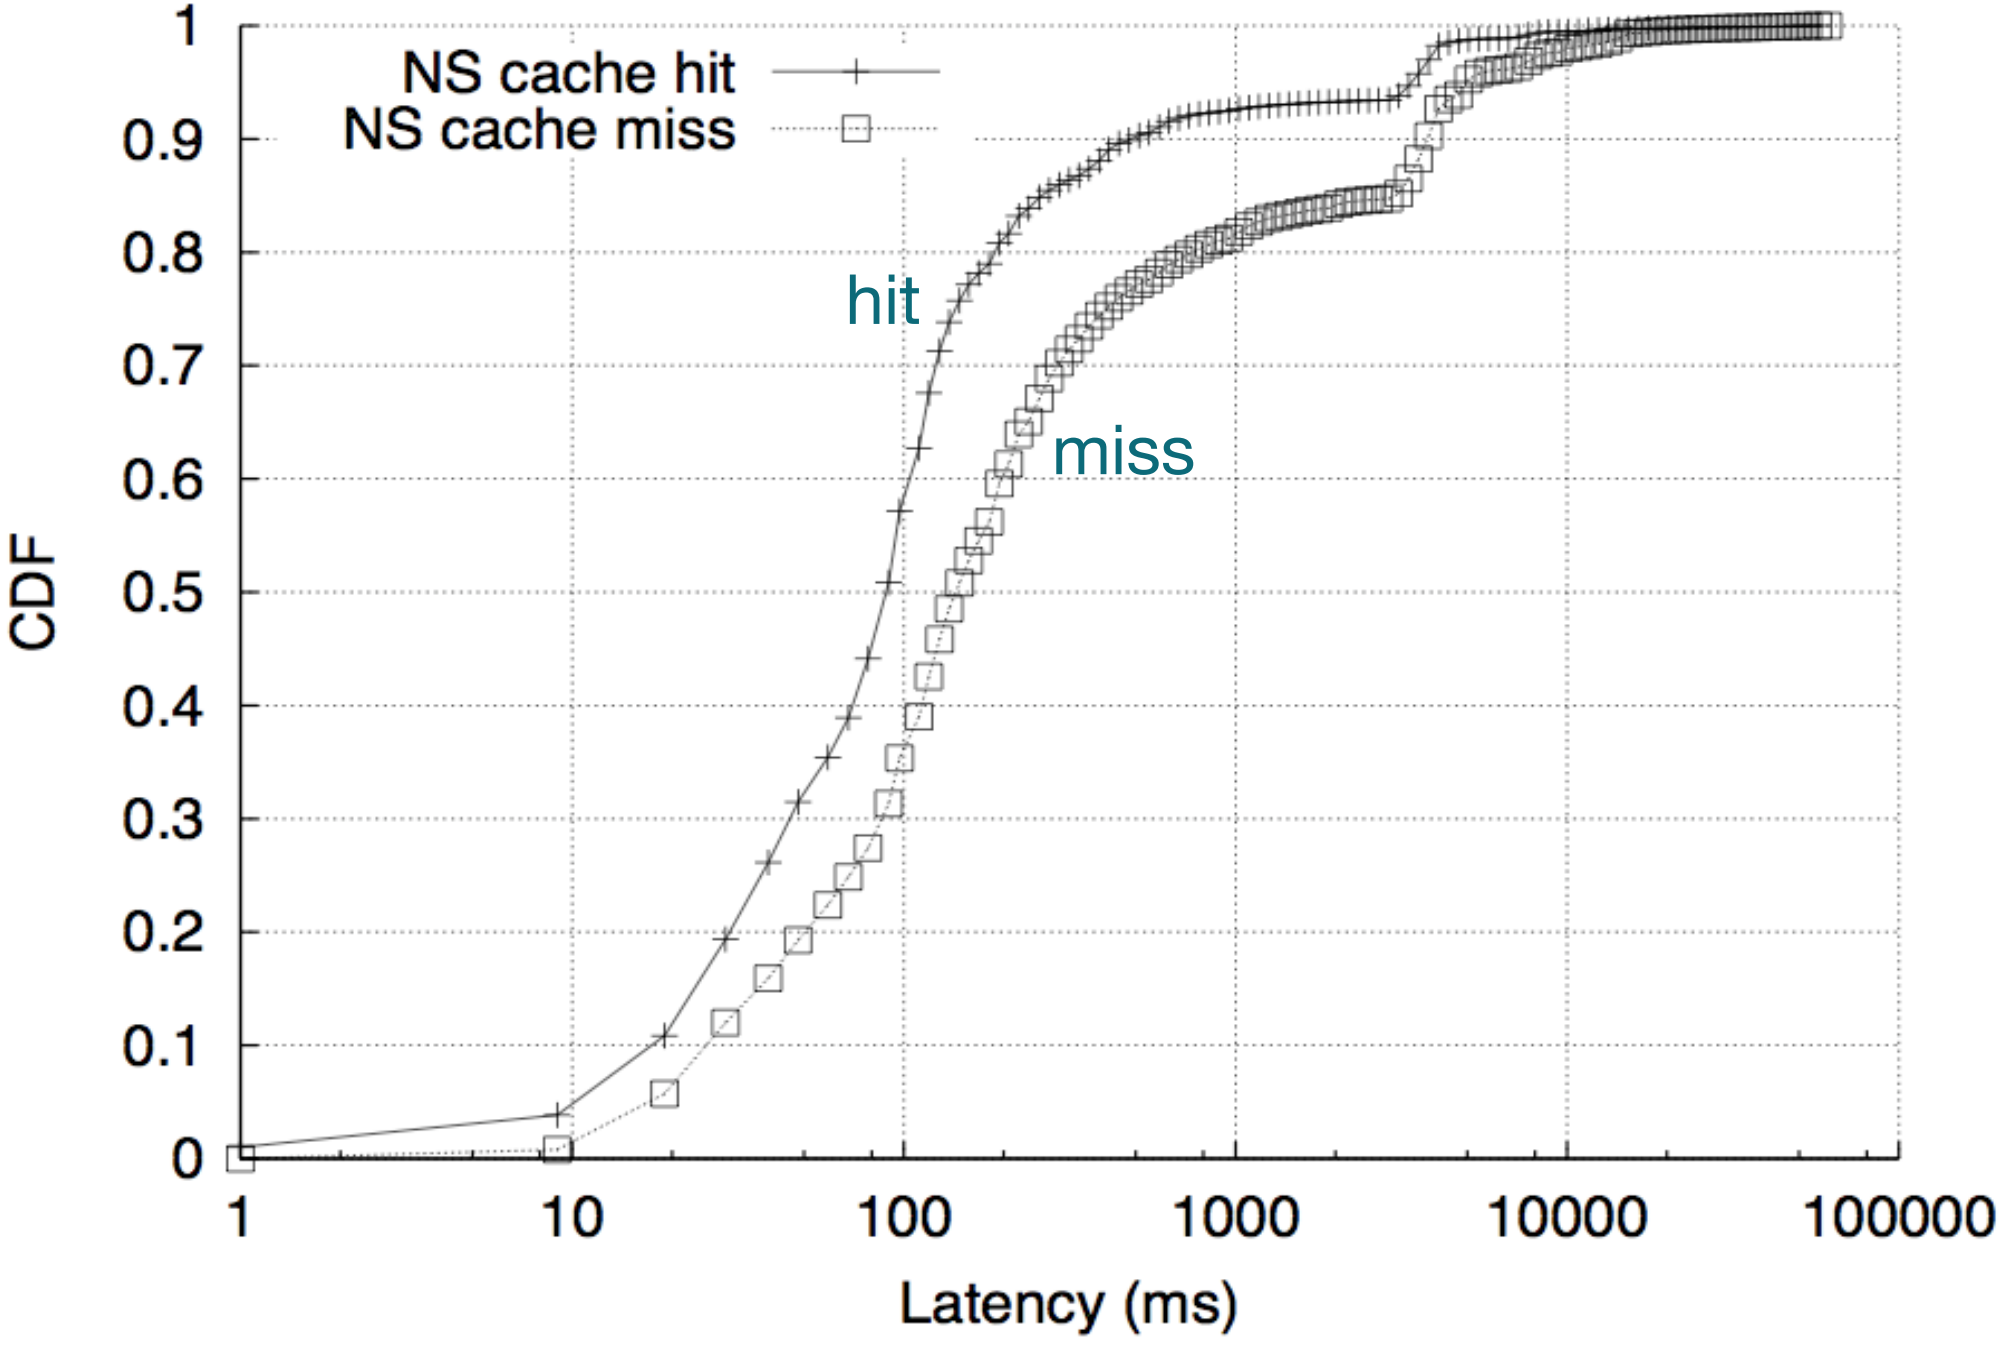
NS cache hit
NS cache miss
hit
miss
CDF
Latency (ms)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
10
100
1000
10000
100000

2,486,104

iterative lookups

6,039,582

query packets

significant number of retransmissions

~24%

of lookups are unanswered

persistent retransmission
referral loops

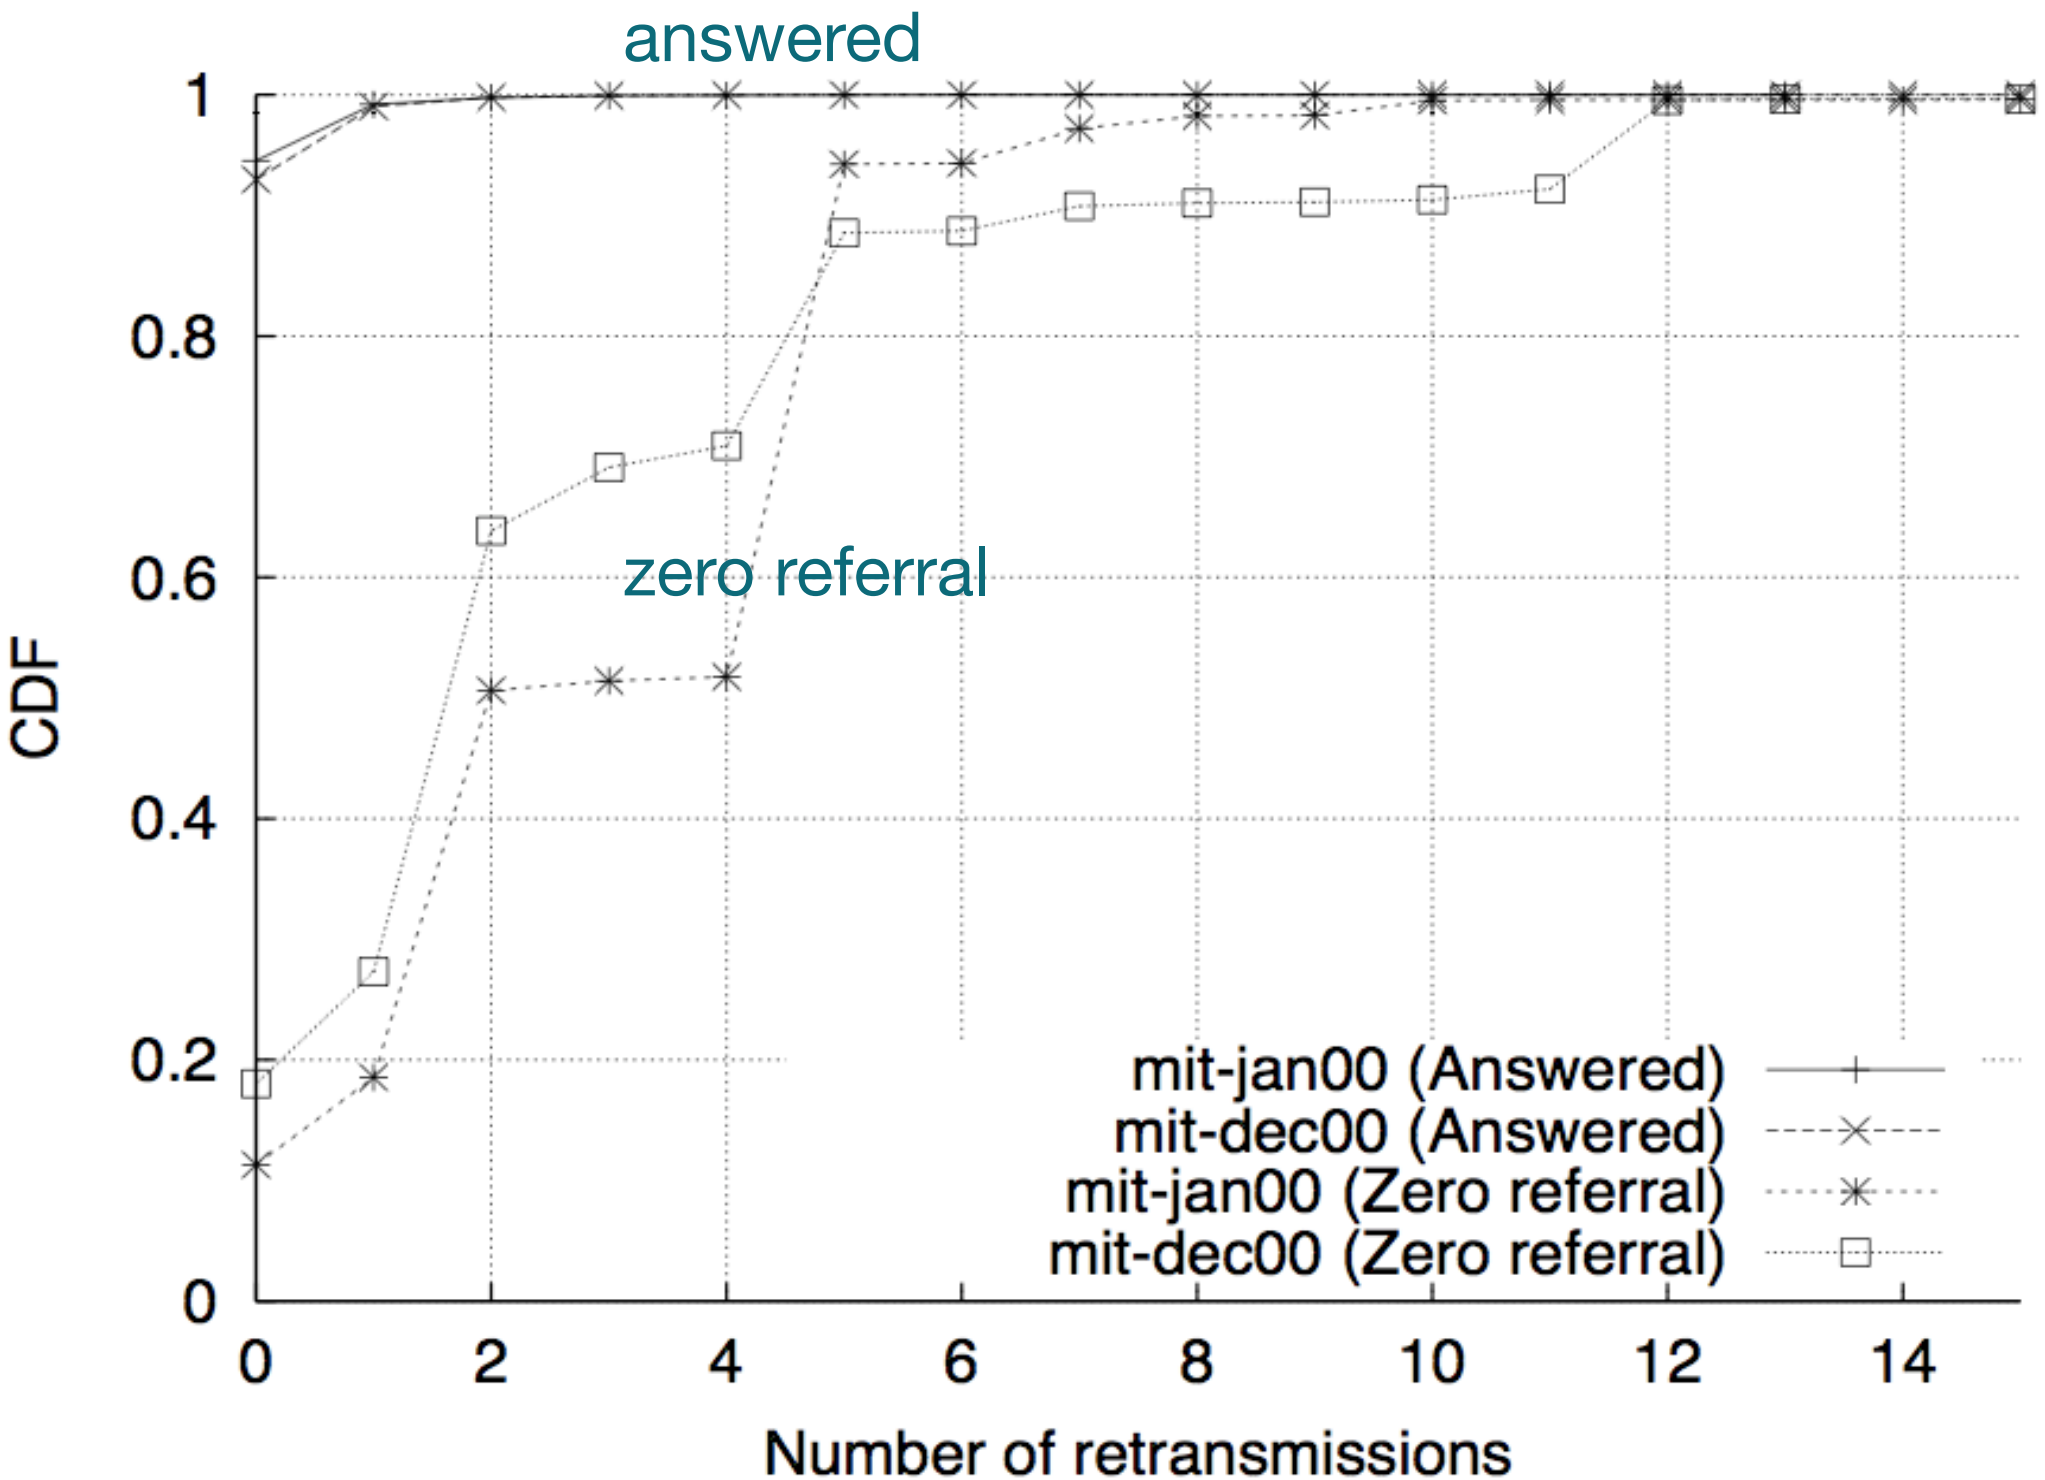
answered
zero referral
CDF
Number of retransmissions
0
0.2
0.4
0.6
0.8
1
0
2
4
6
8
10
12
14
mit-jan00 (Answered)
mit-dec00 (Answered)
mit-jan00 (Zero referral)
mit-dec00 (Zero referral)

no need to retransmit beyond 2-3 retransmissions!

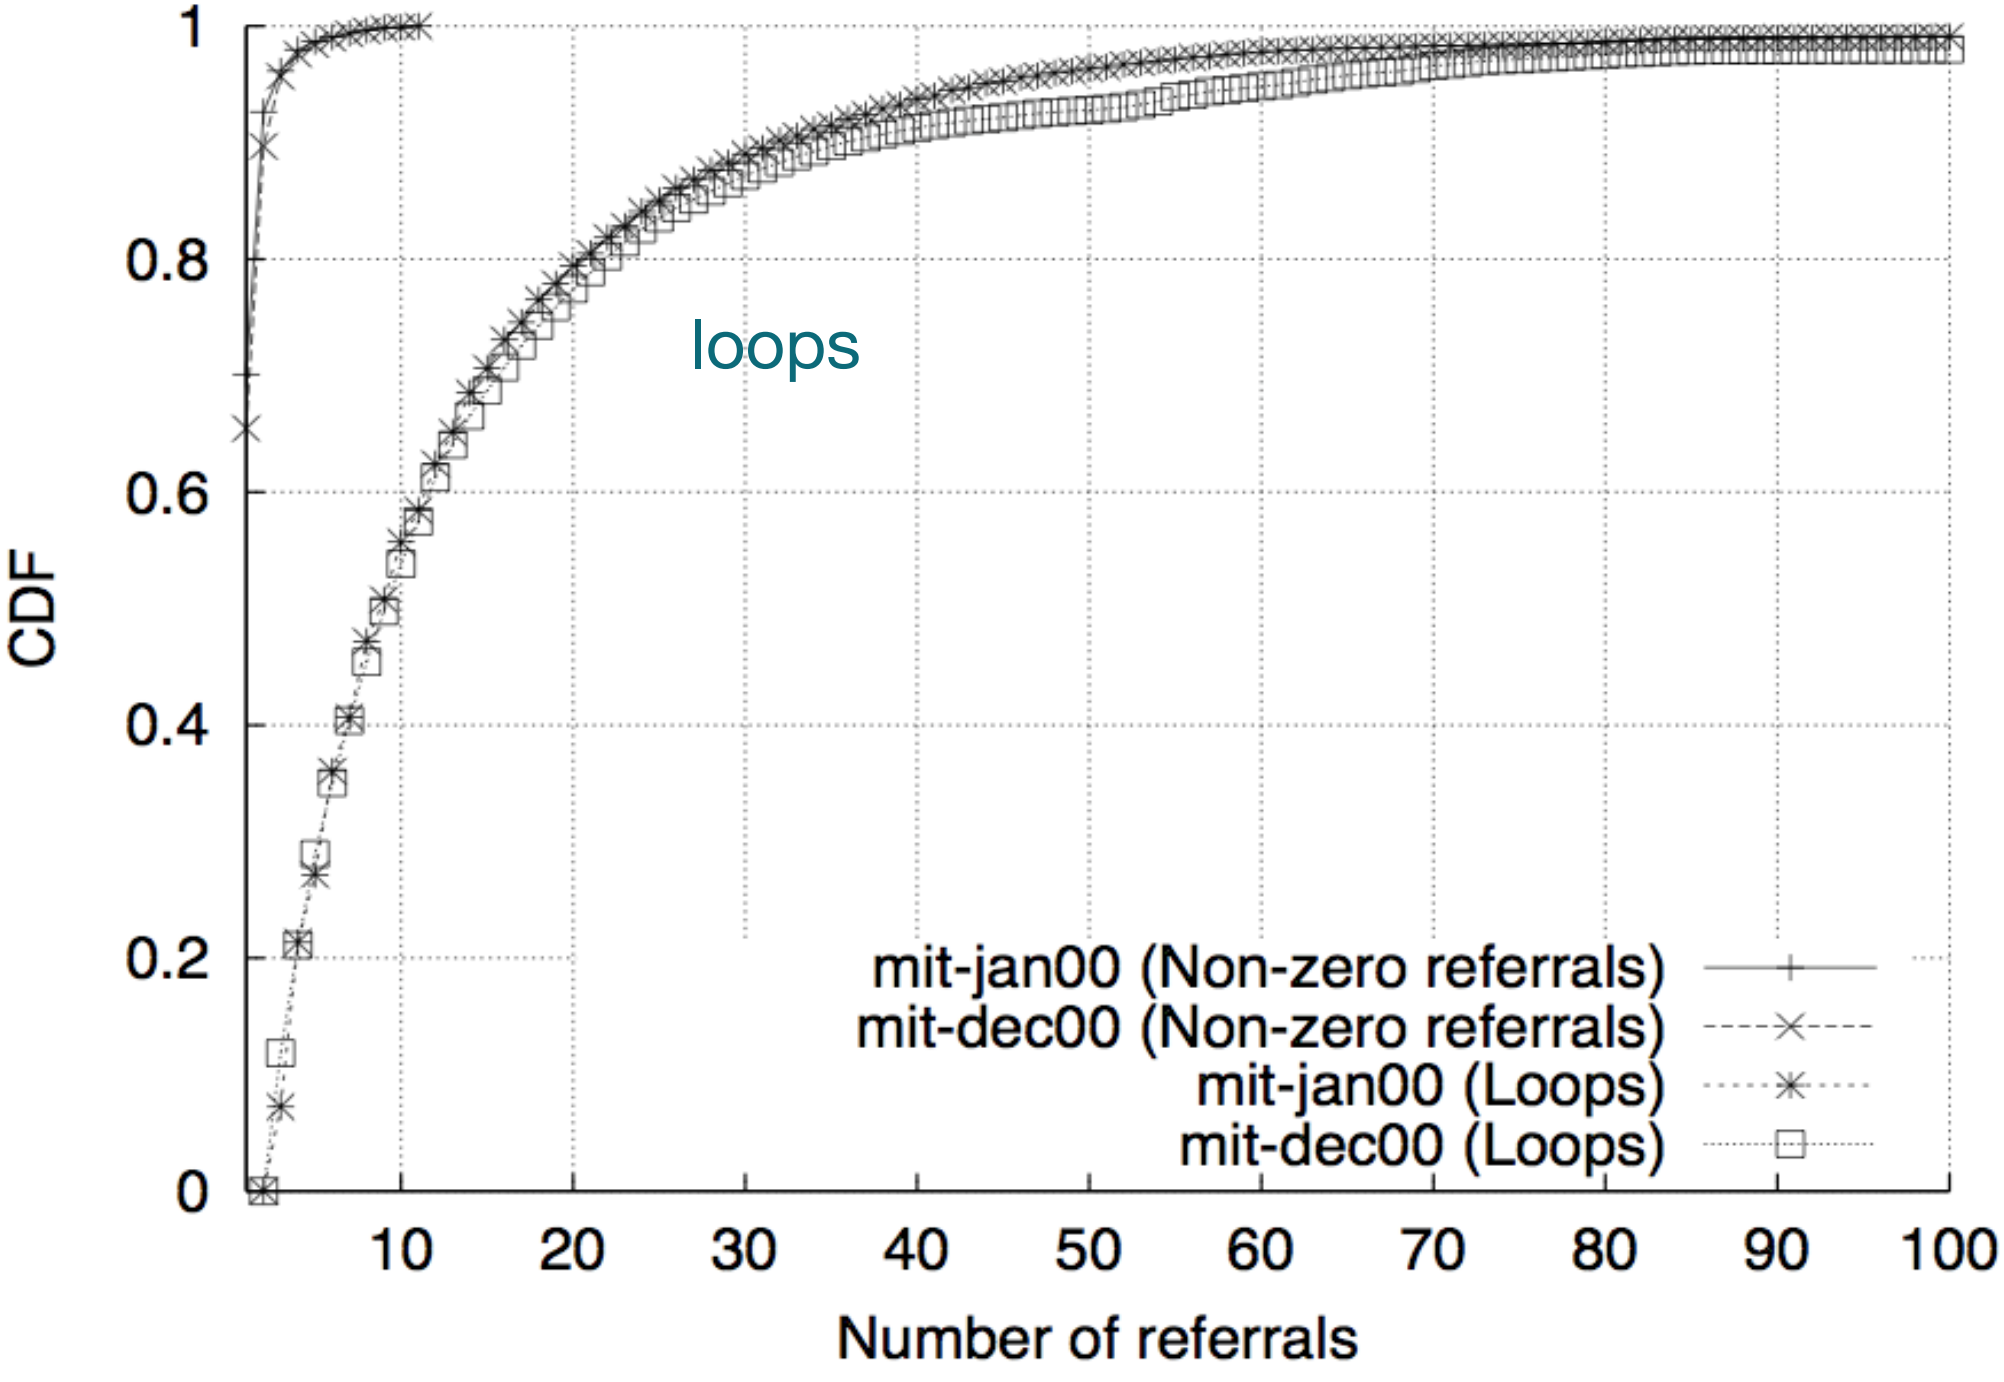
loops
CDF
Number of referrals
0
0.2
0.4
0.6
0.8
1
10
20
30
40
50
60
70
80
90
100
mit-jan00 (Non-zero referrals)
mit-dec00 (Non-zero referrals)
mit-jan00 (Loops)
mit-dec00 (Loops)

each loop generated on average

10

query packets

~60%

of all queries are generated by unanswered lookup

popularity of domain name follows

Zipf

distribution

Frequency of the x-th most popular item is

$$P_i \propto \frac{1}{x^{\alpha}}$$

popularity of web pages
population of countries
occurances of English words
:
:
:
:

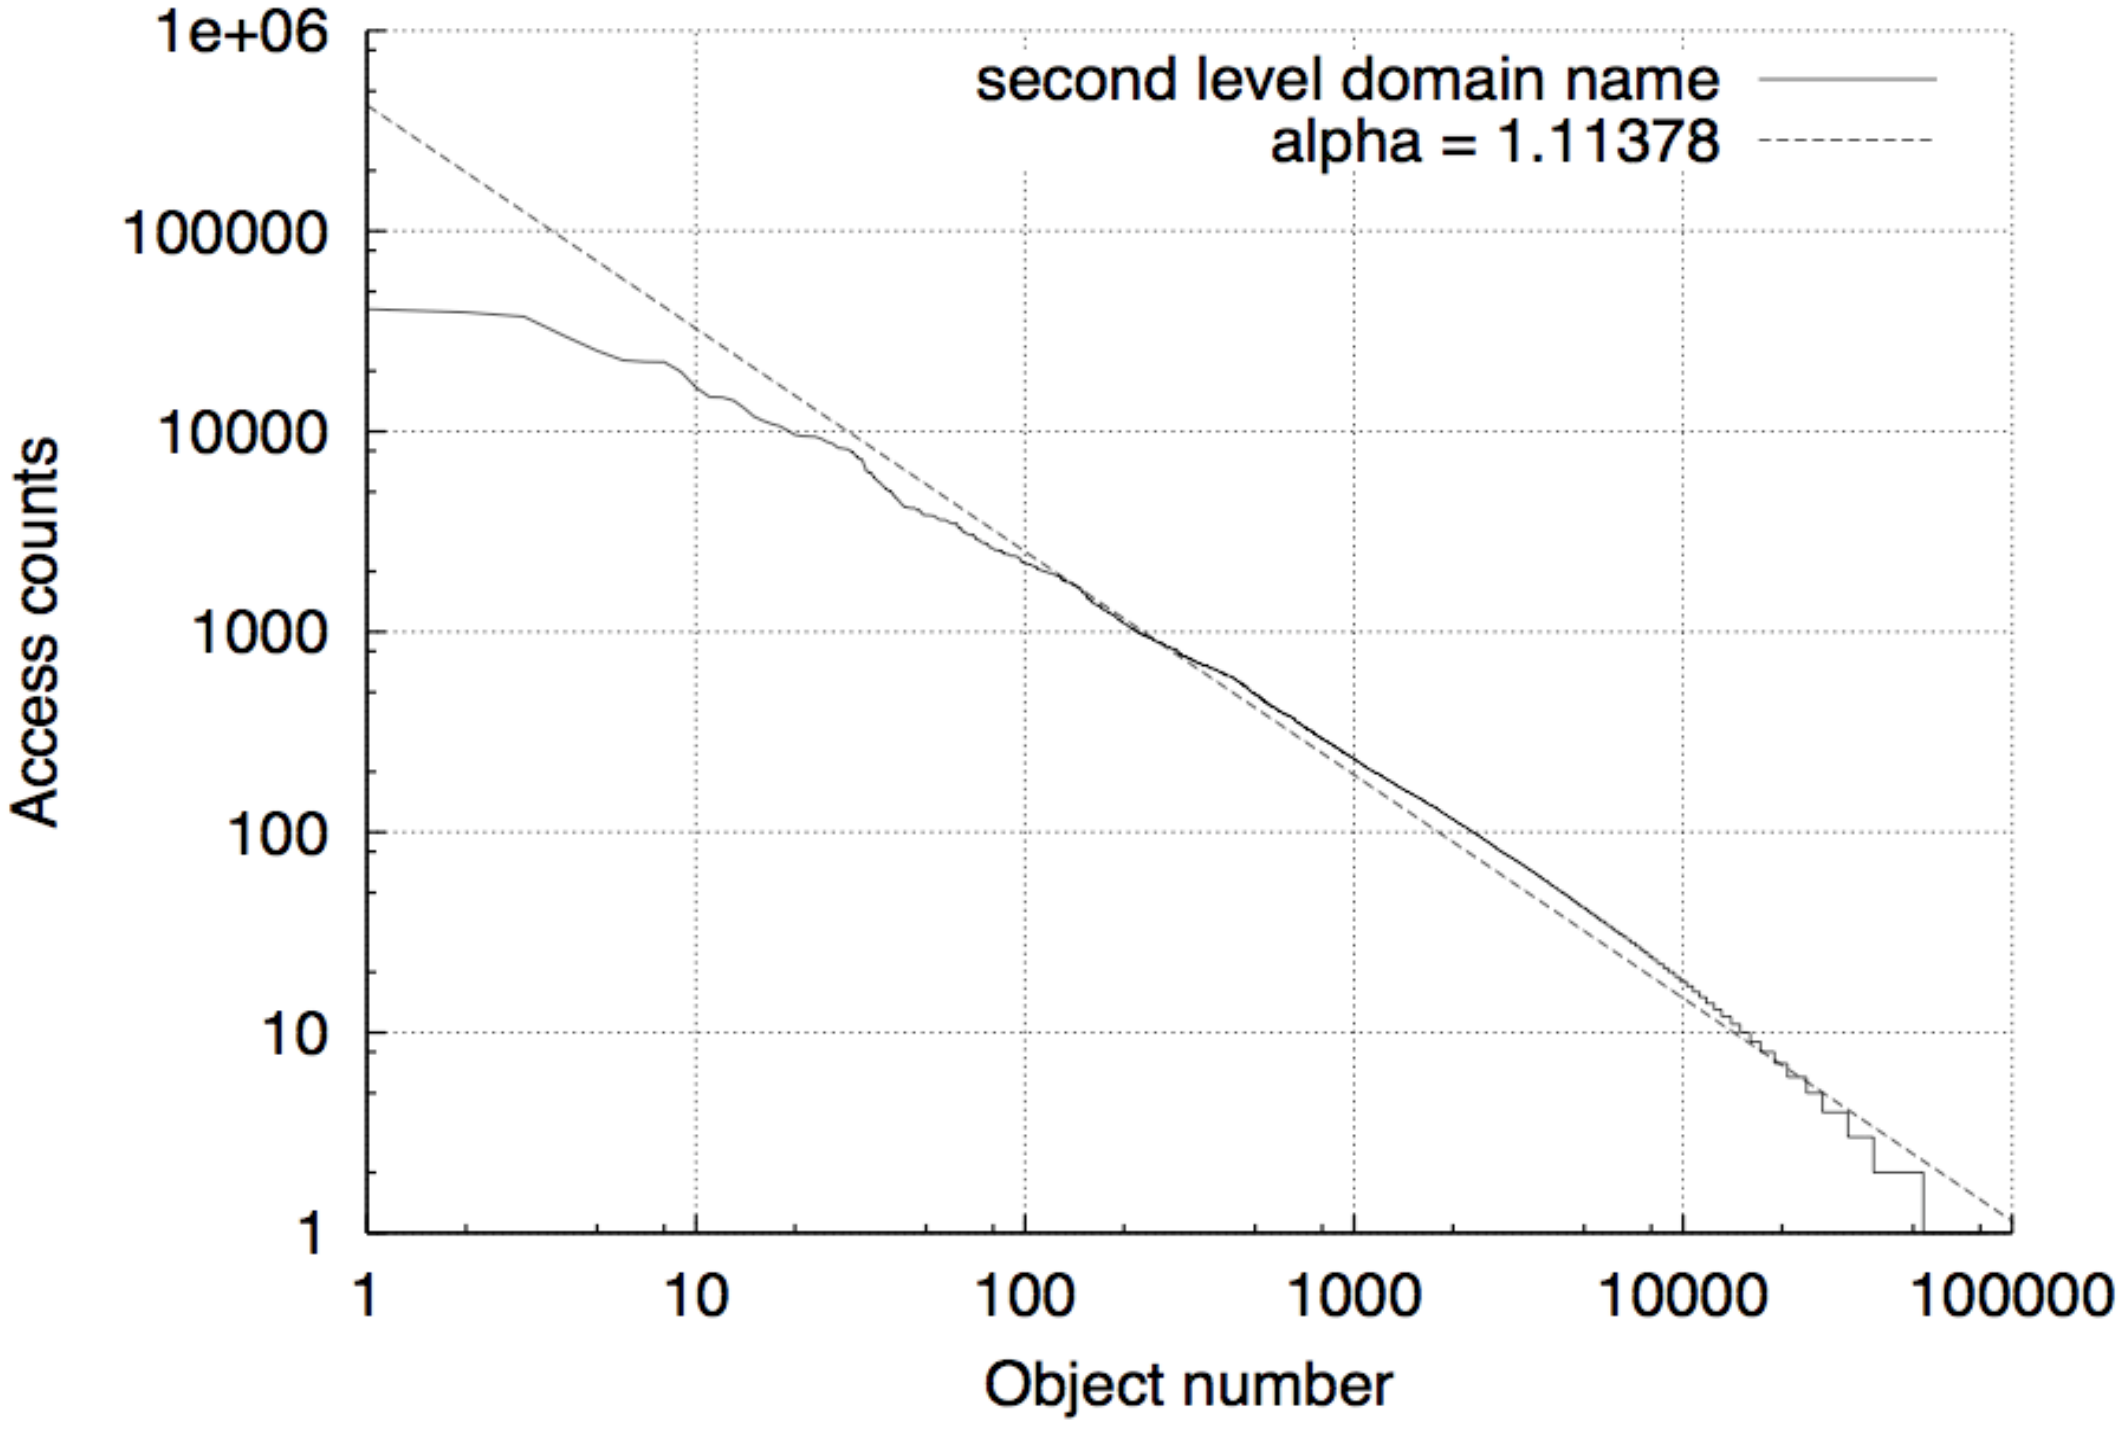
second level domain name
alpha = 1.11378
Access counts
Object number
1e+06
100000
10000
1000
100
10
1
1
10
100
1000
10000
100000

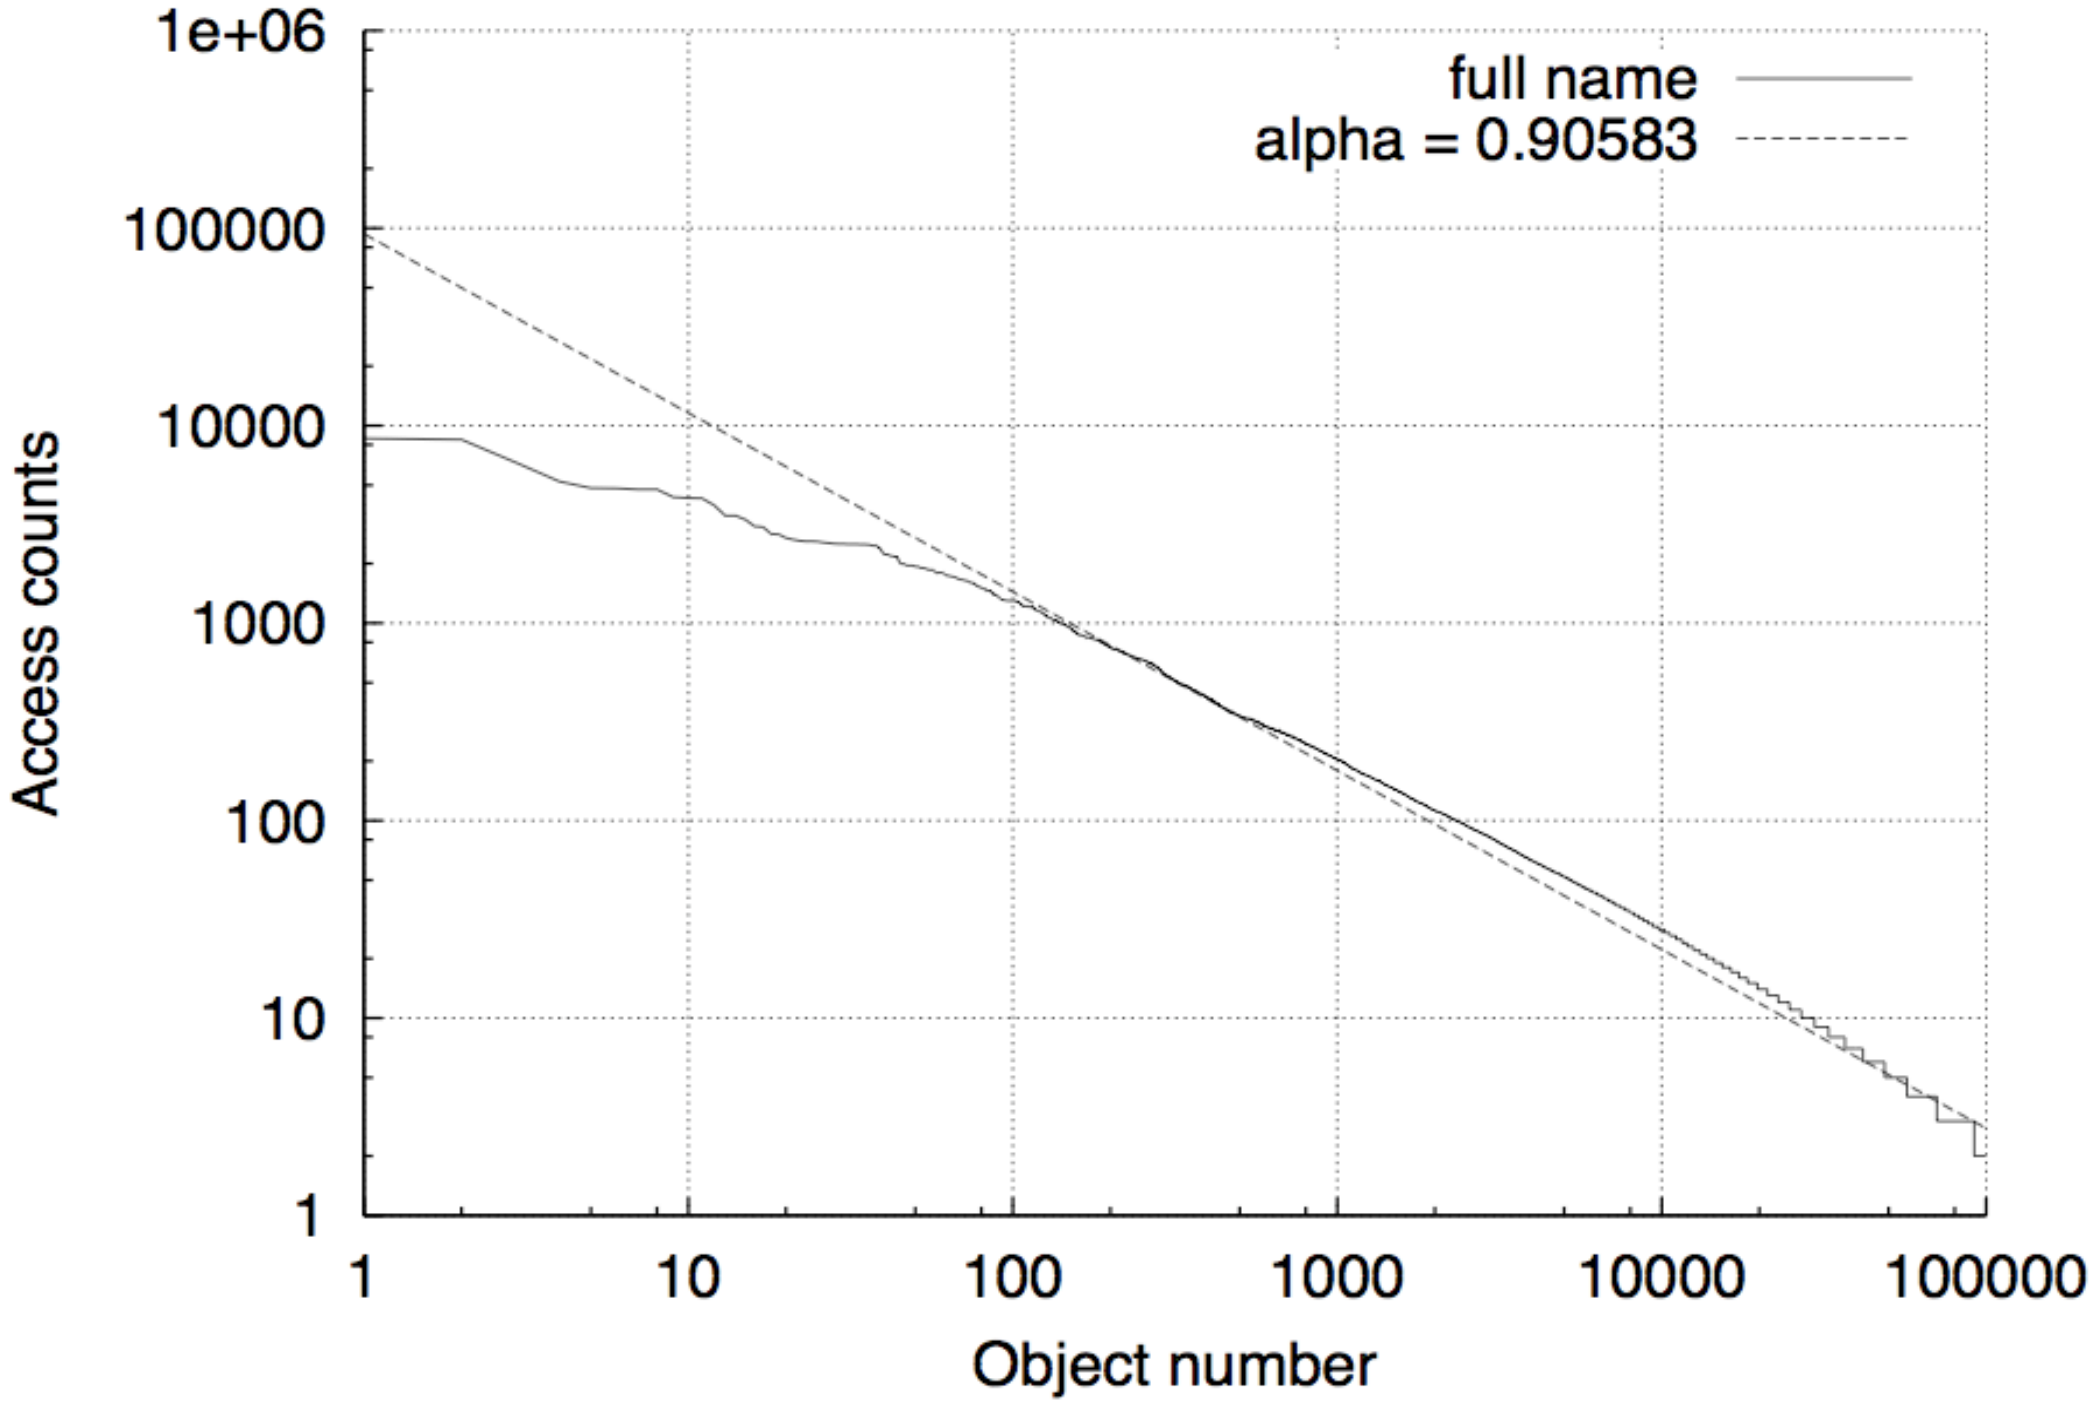
full name
alpha = 0.90583
Access counts
Object number
1e+06
100000
10000
1000
100
10
1
1
10
100
1000
10000
100000

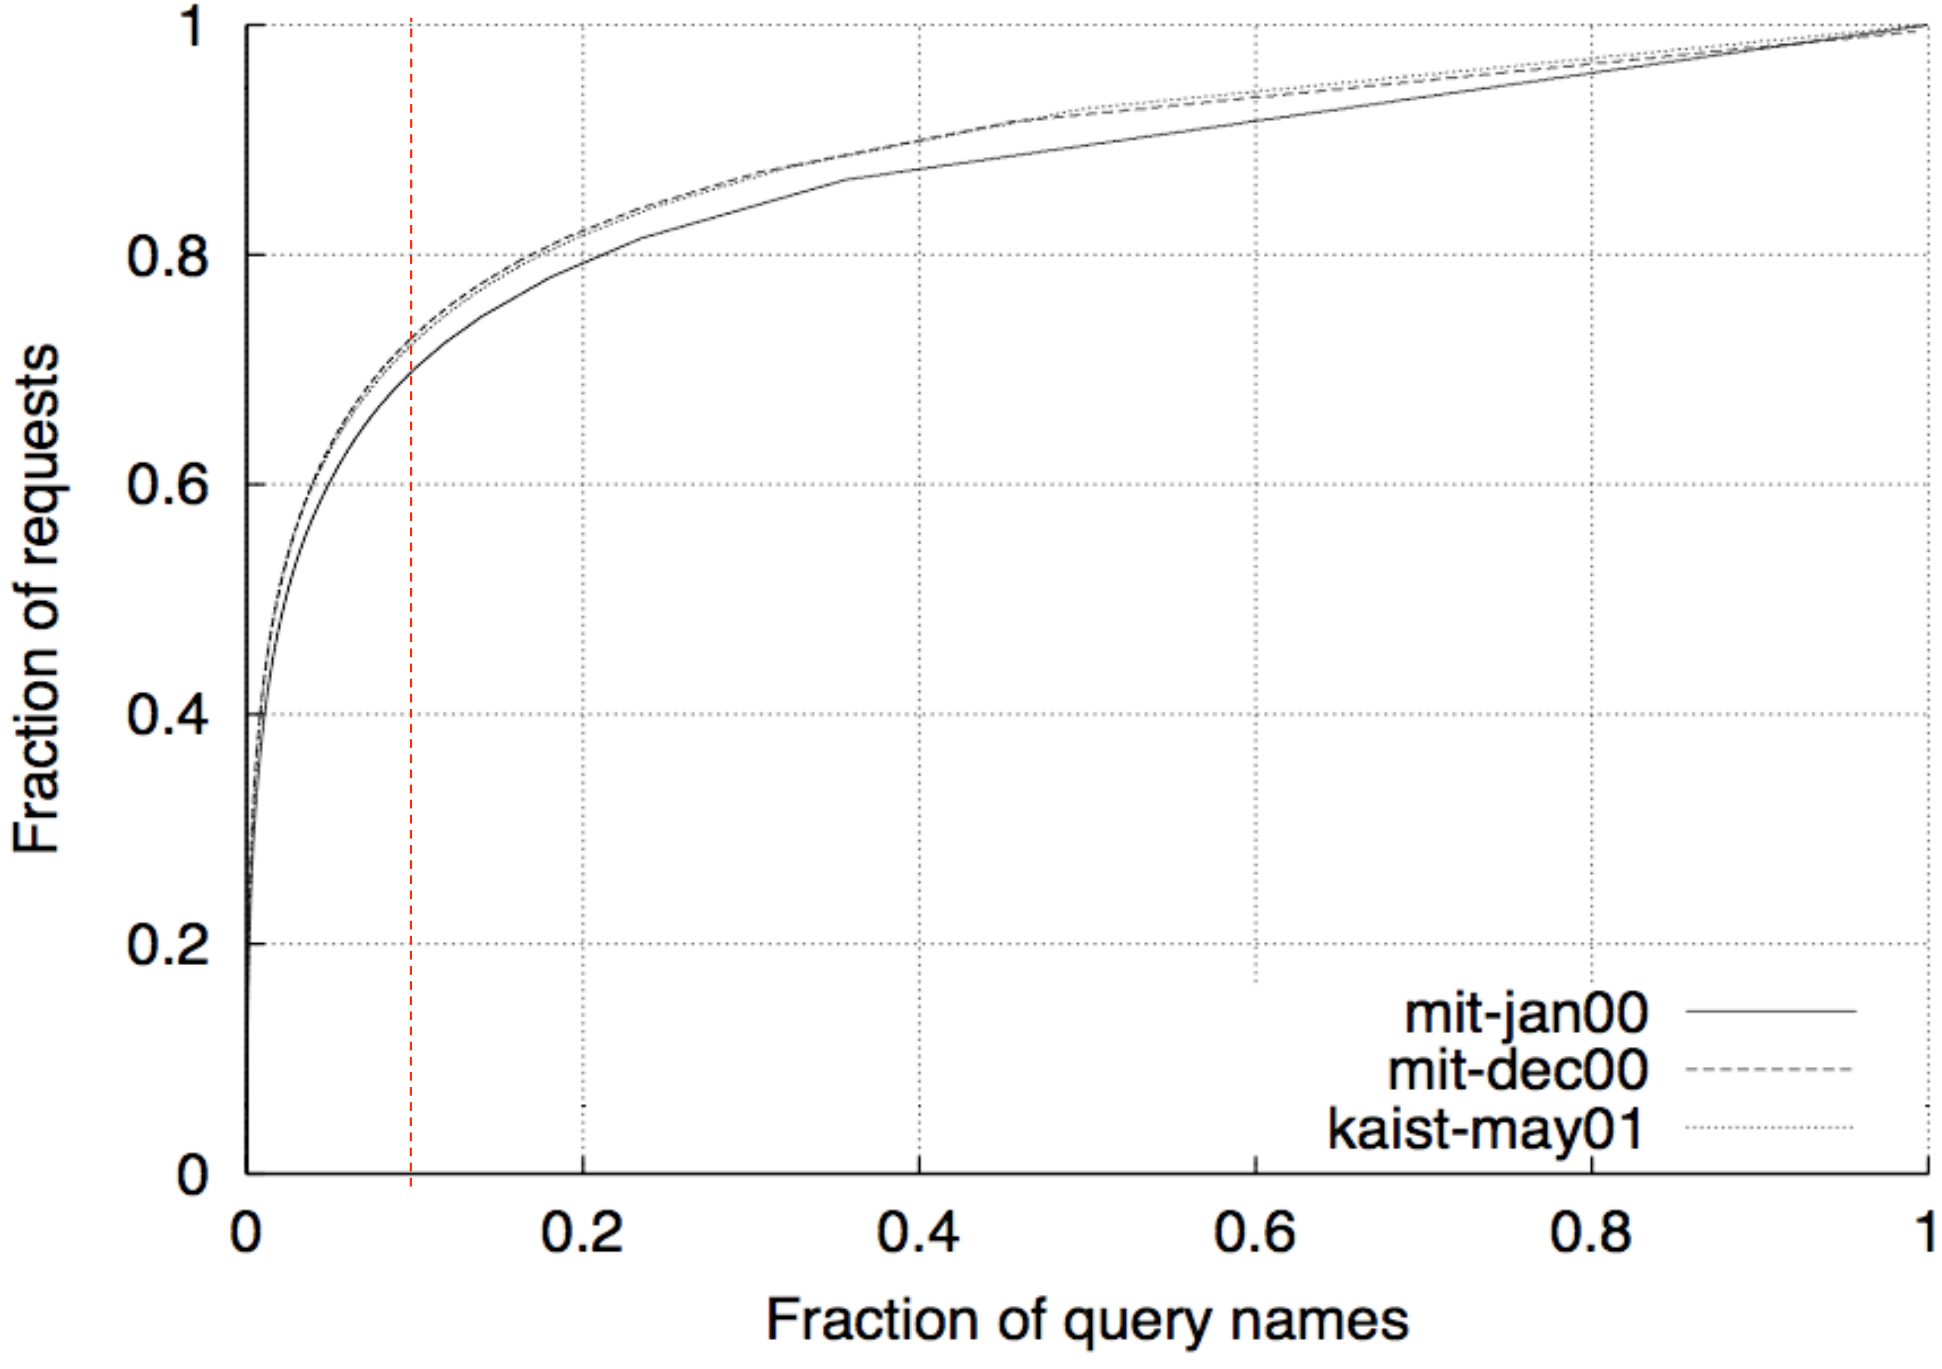
1
0.8
0.6
0.4
0.2
0
0
0.2
0.4
0.6
0.8
1
Fraction of requests
Fraction of query names
mit-jan00
mit-dec00
kaist-may01

long tail :

46%

of domain names
are accessed once

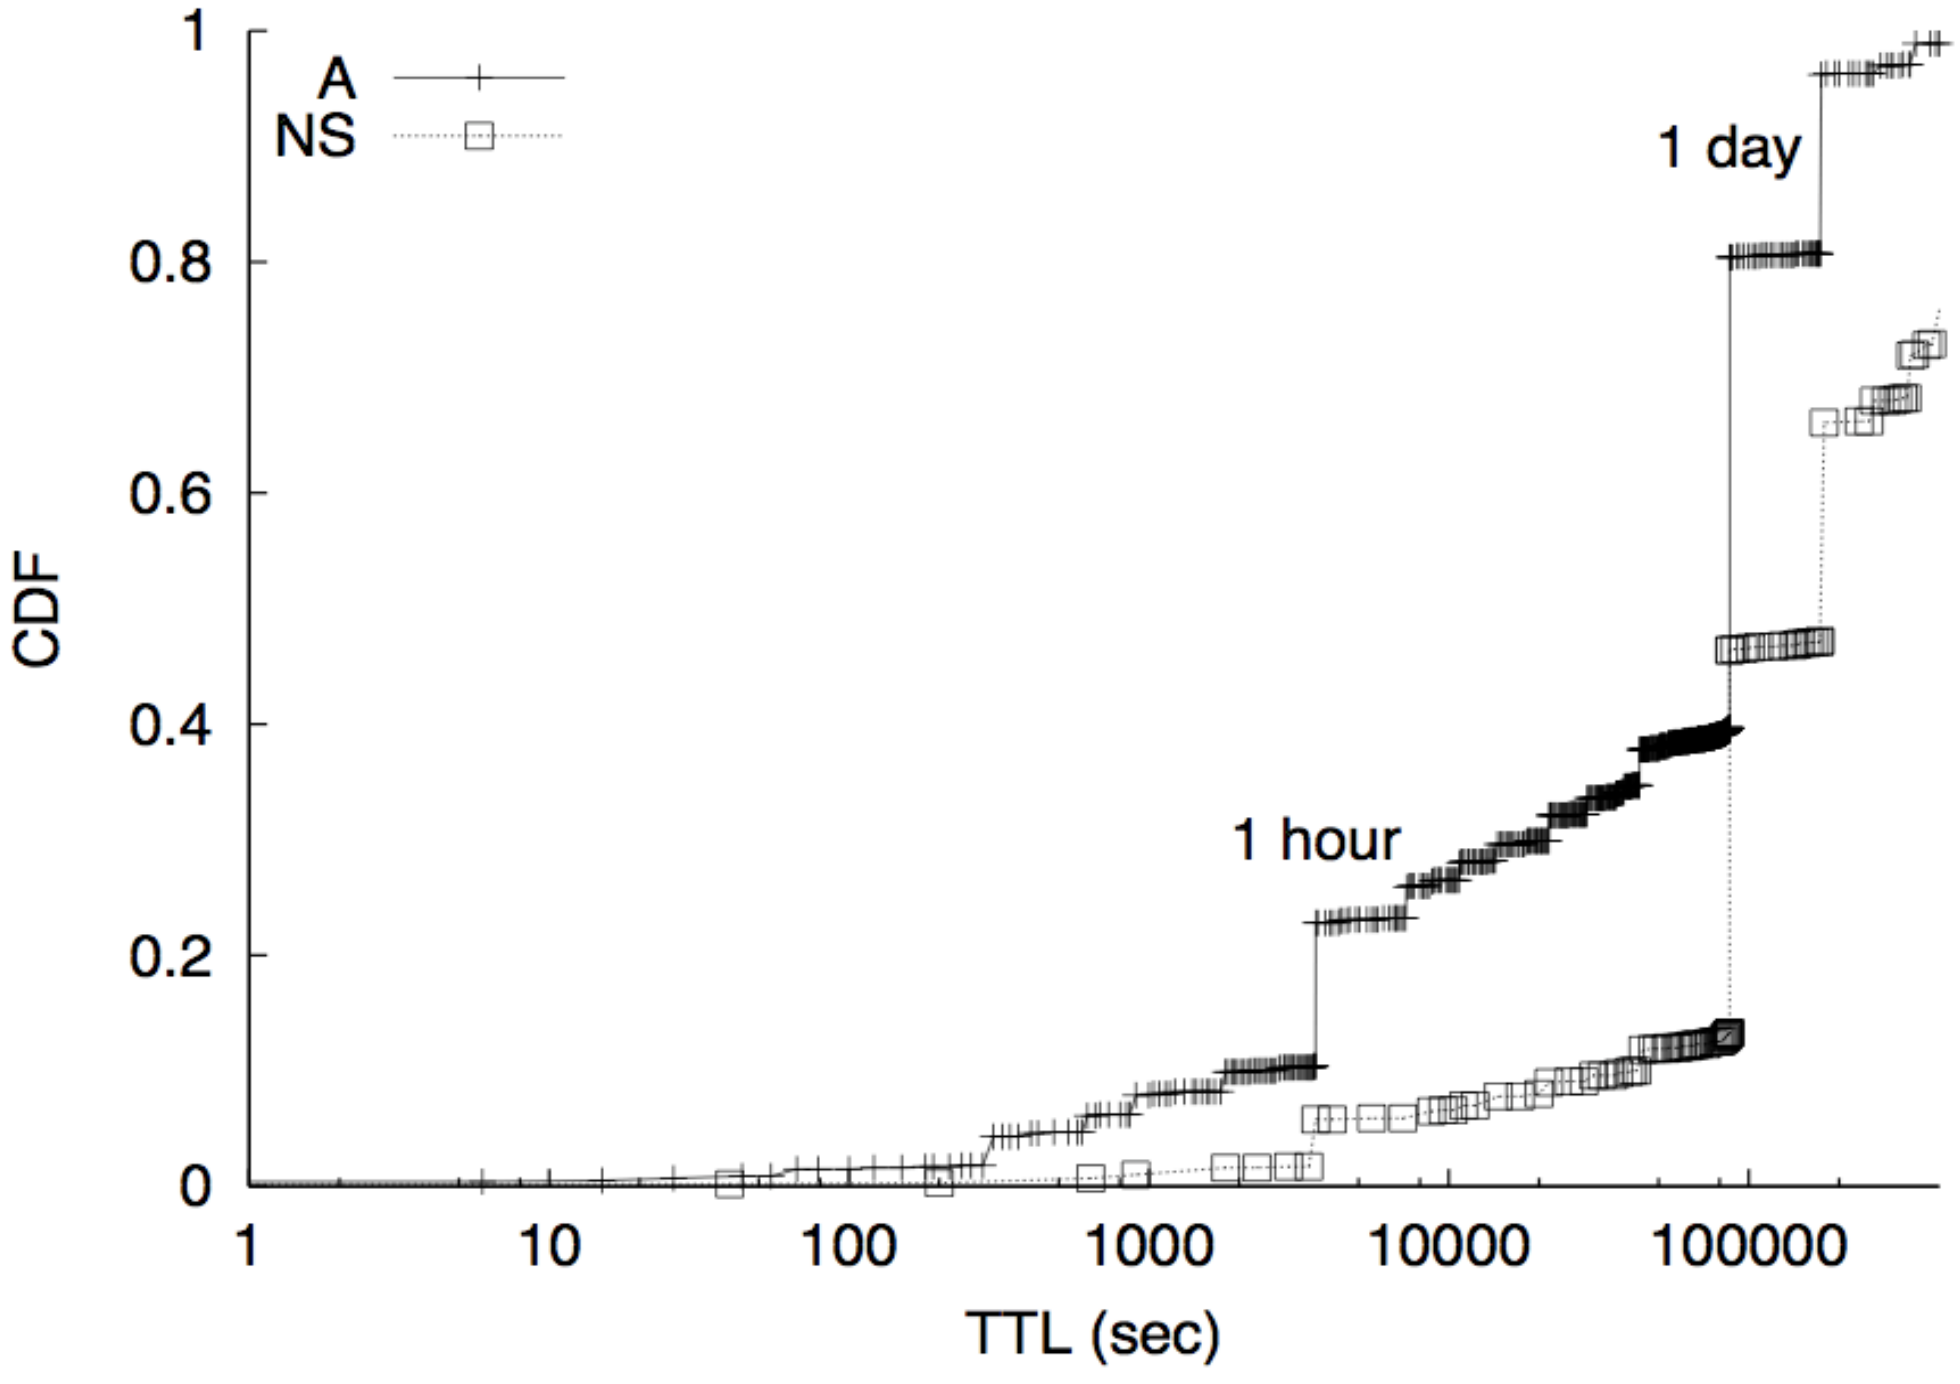
A
NS
1 day
1 hour
CDF
TTL (sec)
1
0.8
0.6
0.4
0.2
0
1
10
100
1000
10000
100000

20%

DNS responses are
from root/gTLD servers

without caching NS records:

5X

loads on root/gTLD servers

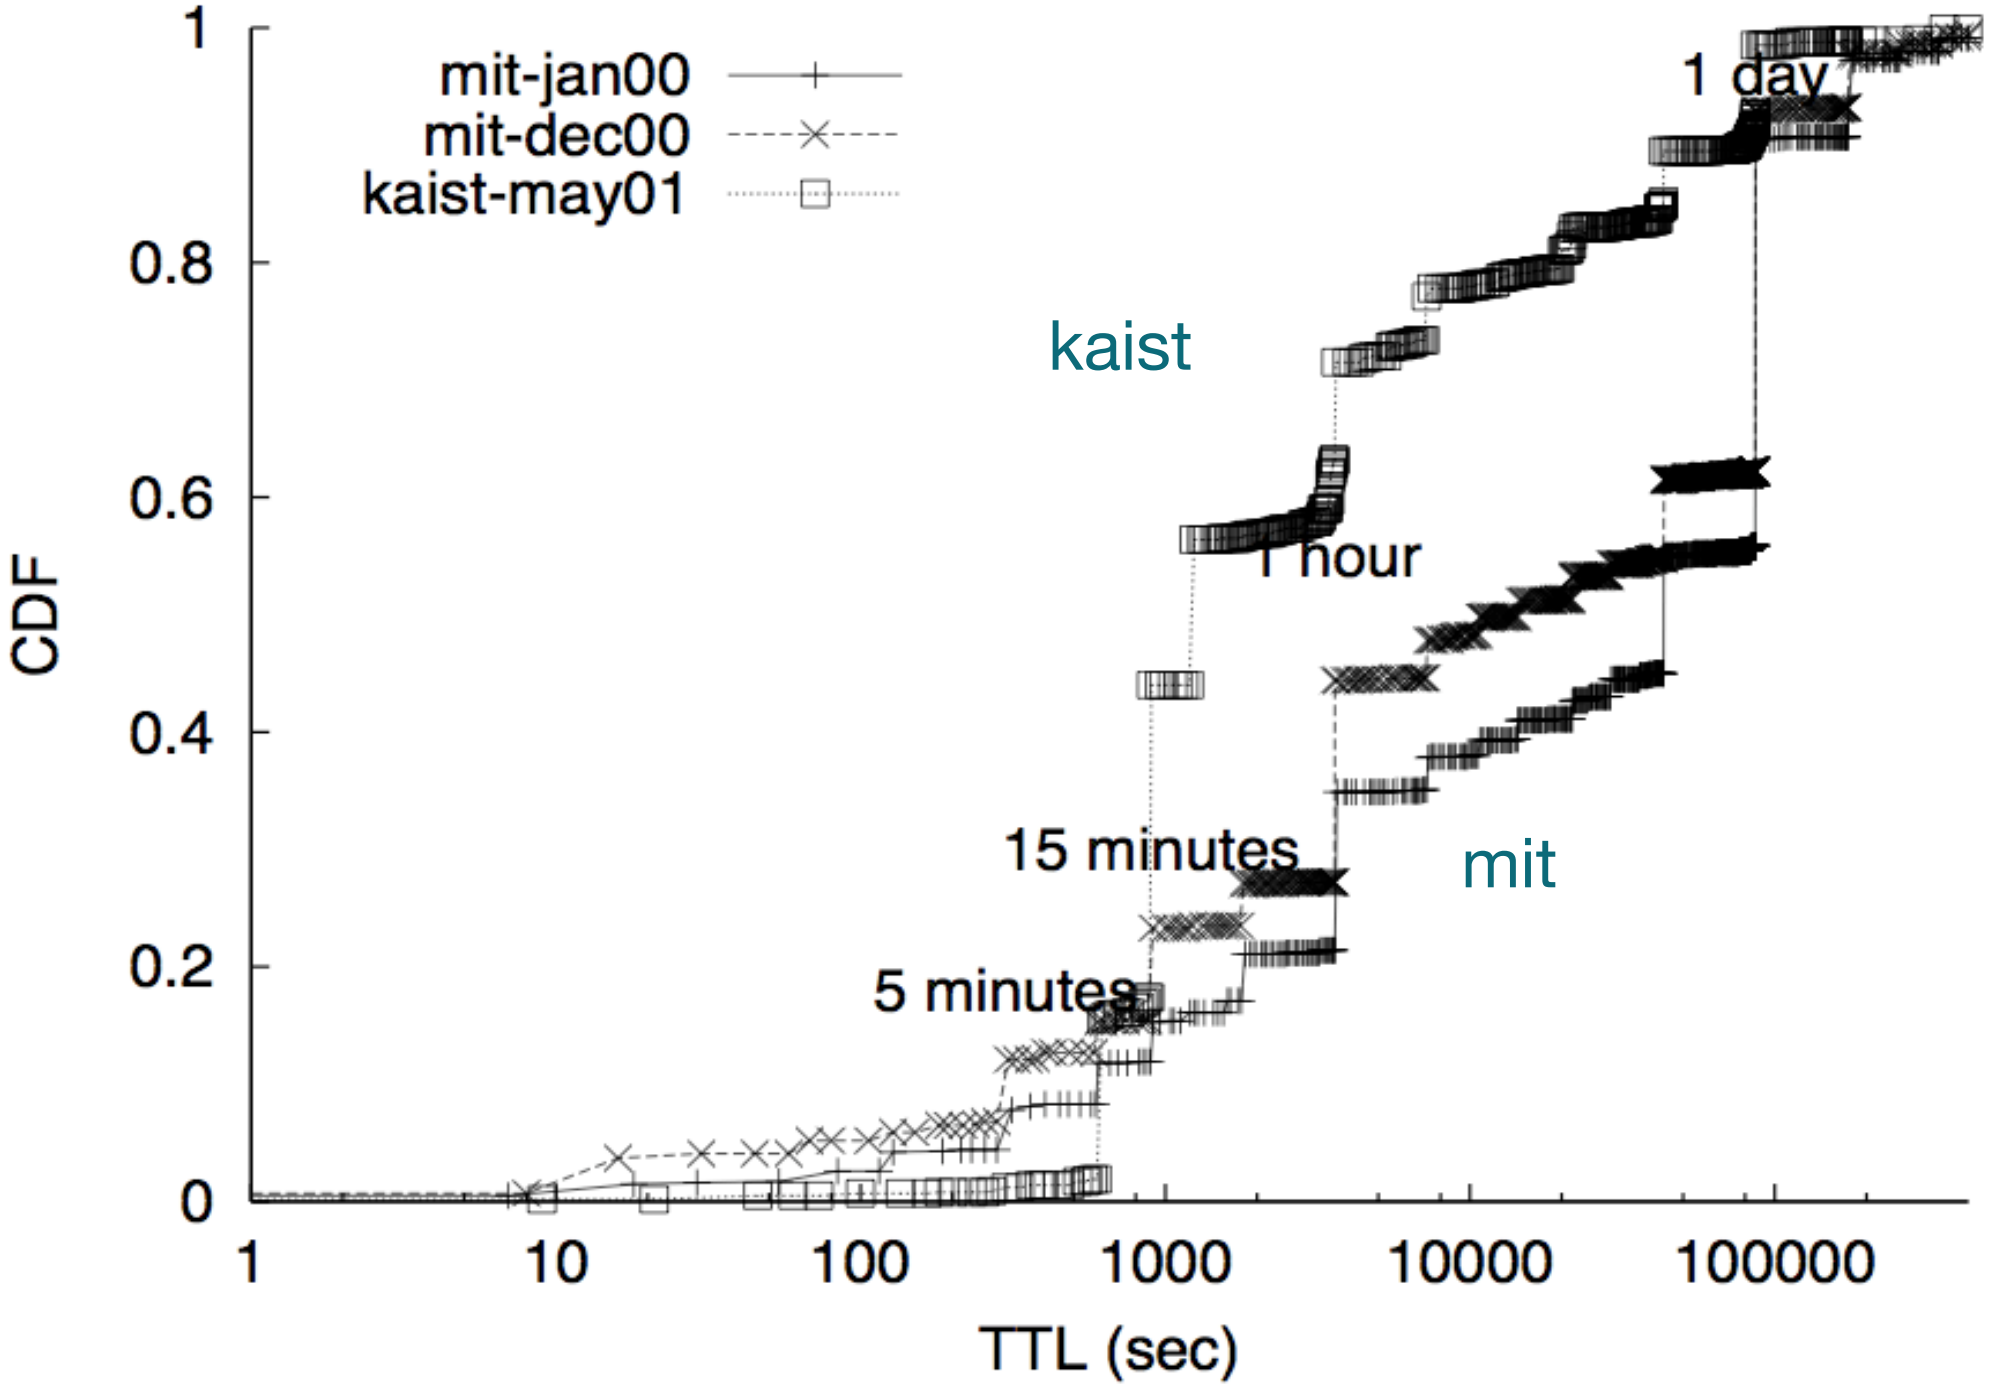
mit-jan00
mit-dec00
kaist-may01
1 day
kaist
1 hour
15 minutes
mit
5 minutes
CDF
TTL (sec)
0
0.2
0.4
0.6
0.8
1
1
10
100
1000
10000
100000

popular sites have shorter TTL

TLL reduces over time

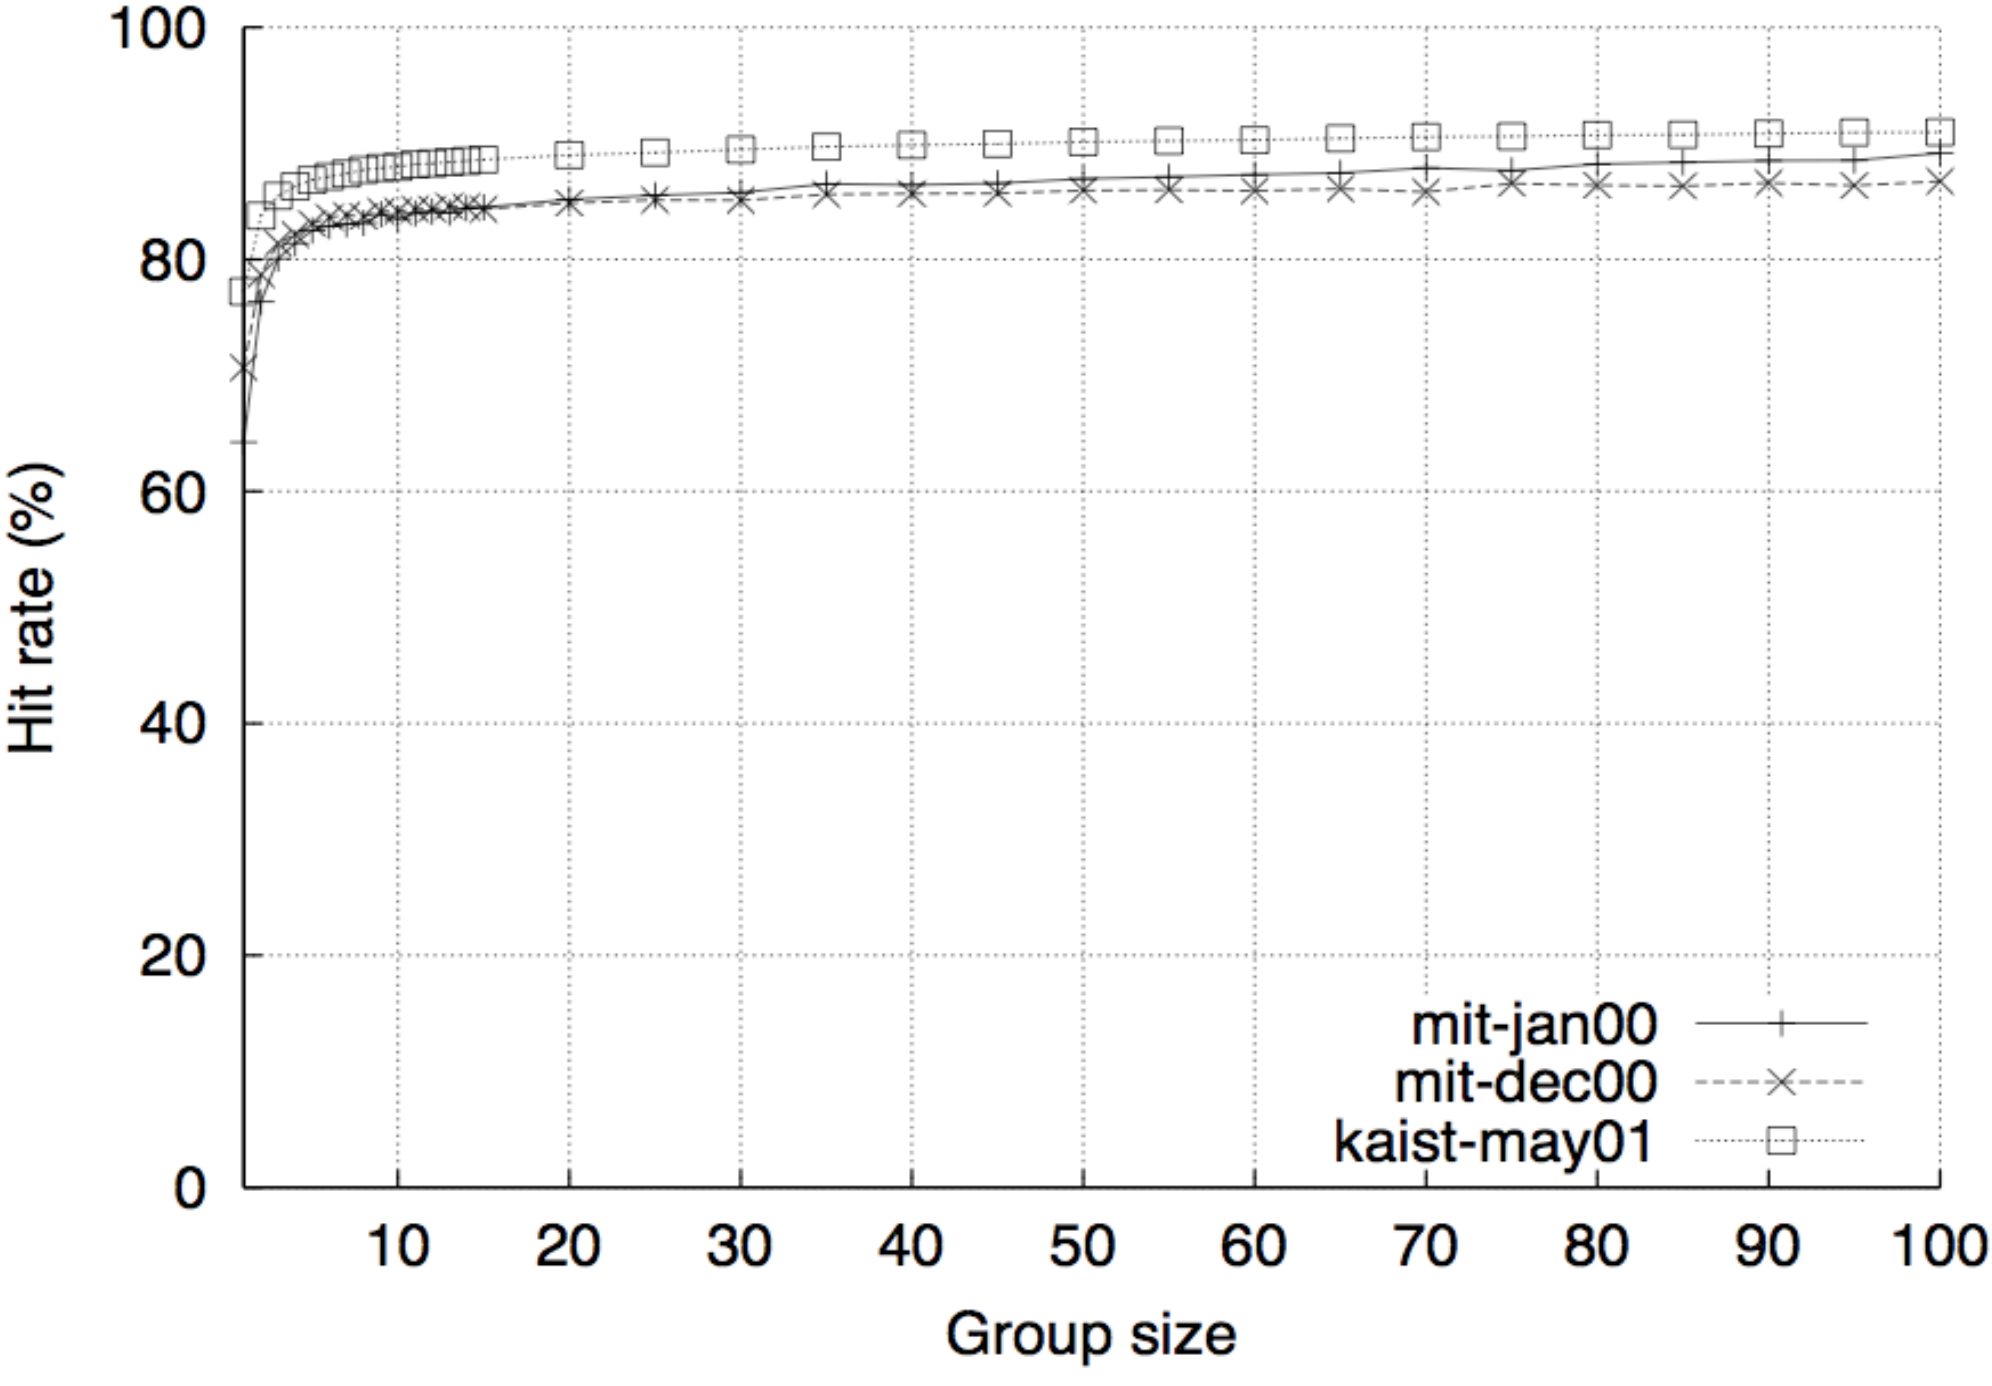
100
80
60
40
20
0
Hit rate (%)
10
20
30
40
50
60
70
80
90
100
Group size
mit-jan00
mit-dec00
kaist-may01

sharing cache does not help

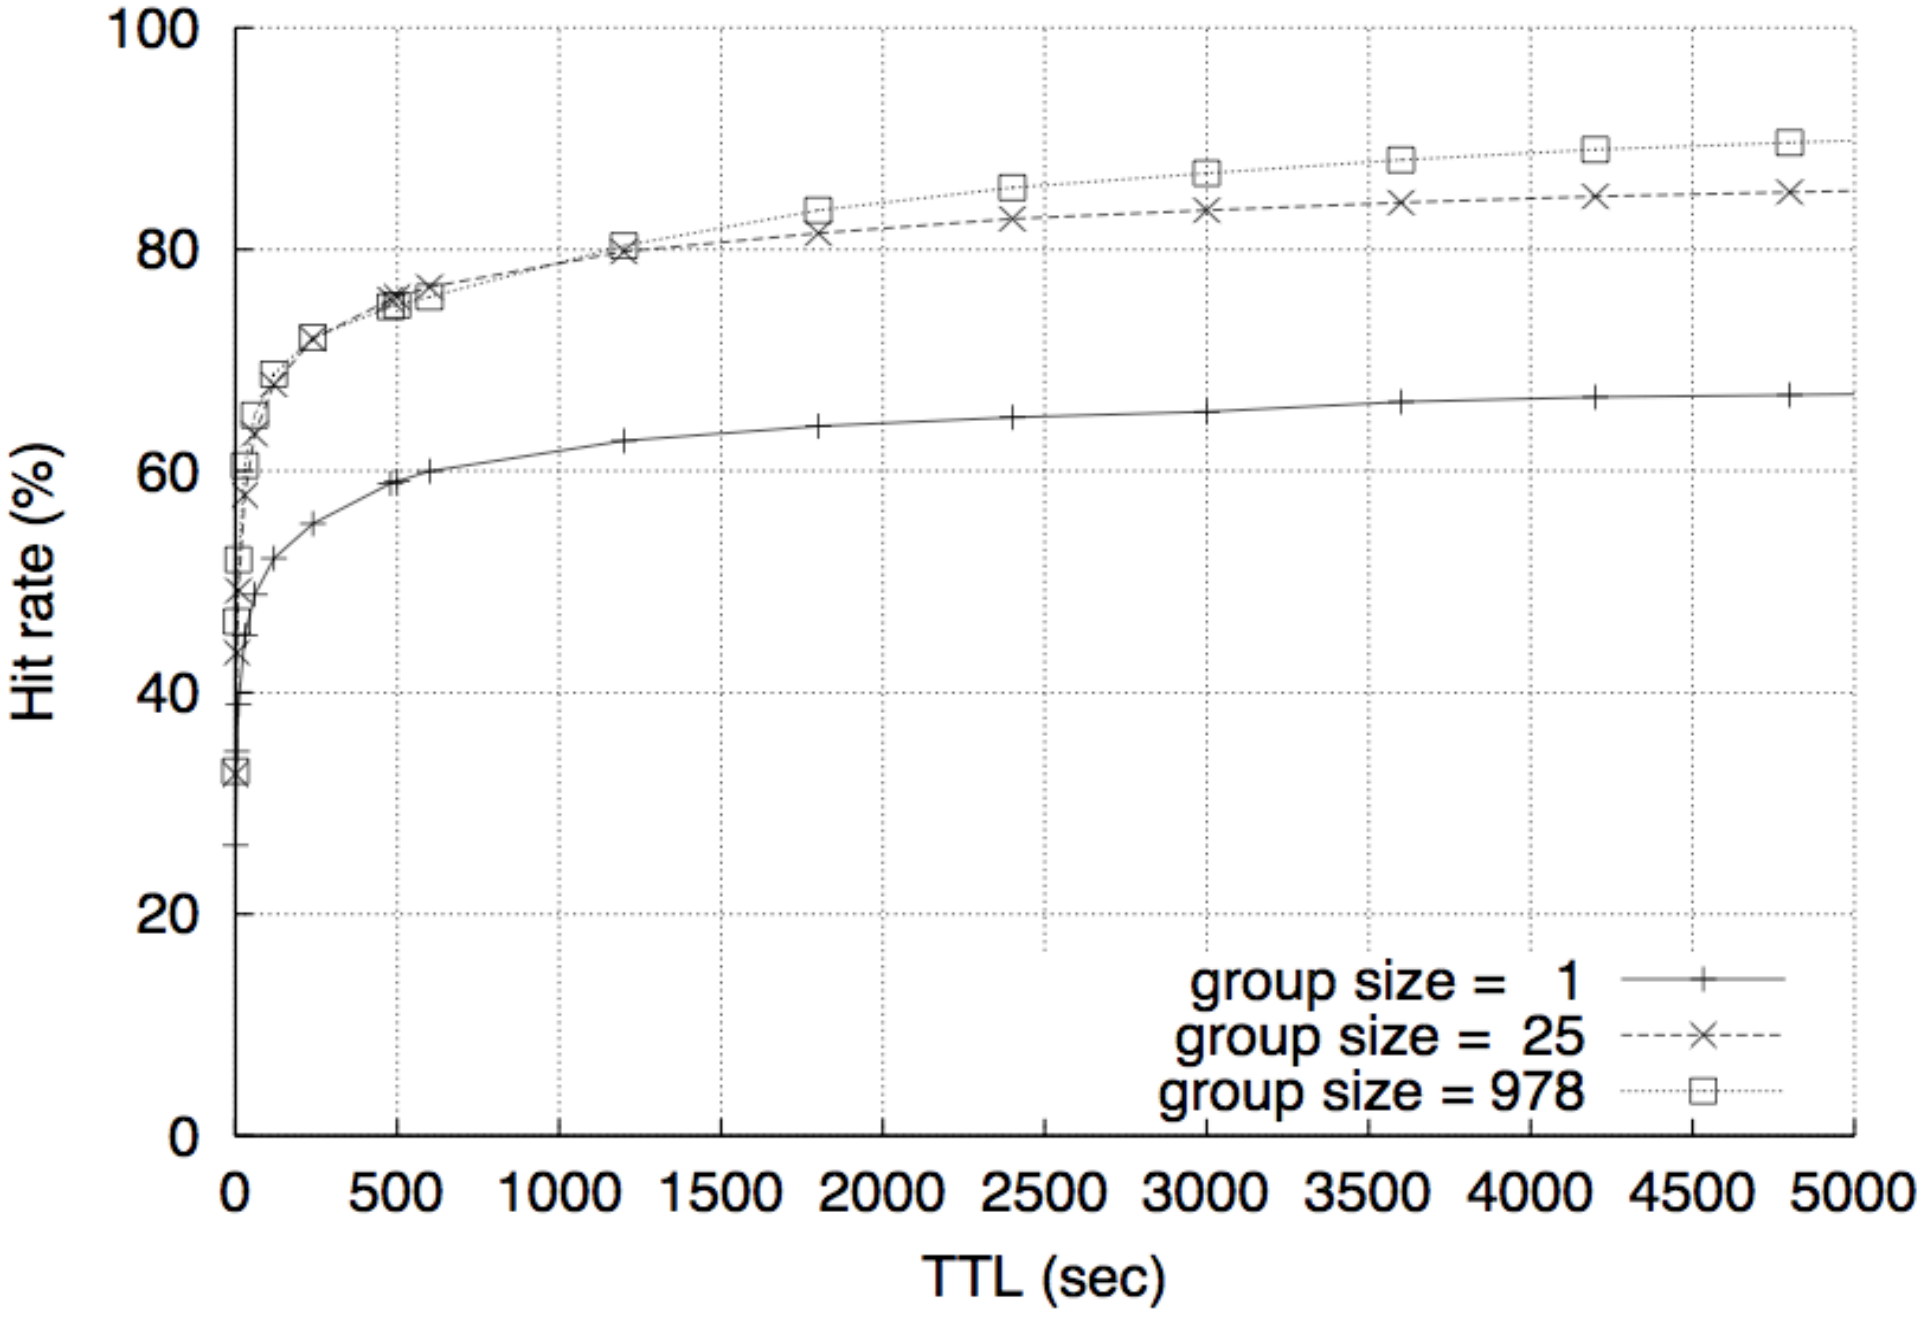
100
80
60
40
20
0
Hit rate (%)
0
500
1000
1500
2000
2500
3000
3500
4000
4500
5000
TTL (sec)
group size = 1
group size = 25
group size = 978

large TTL does not help

“King: Estimating Latency between Arbitrary End Hosts”

Gummadi et. al.

IMC 2002

how to find alternate path
with lower latency?

need to find latency between pair of Internet hosts

assume name server
is close to end host

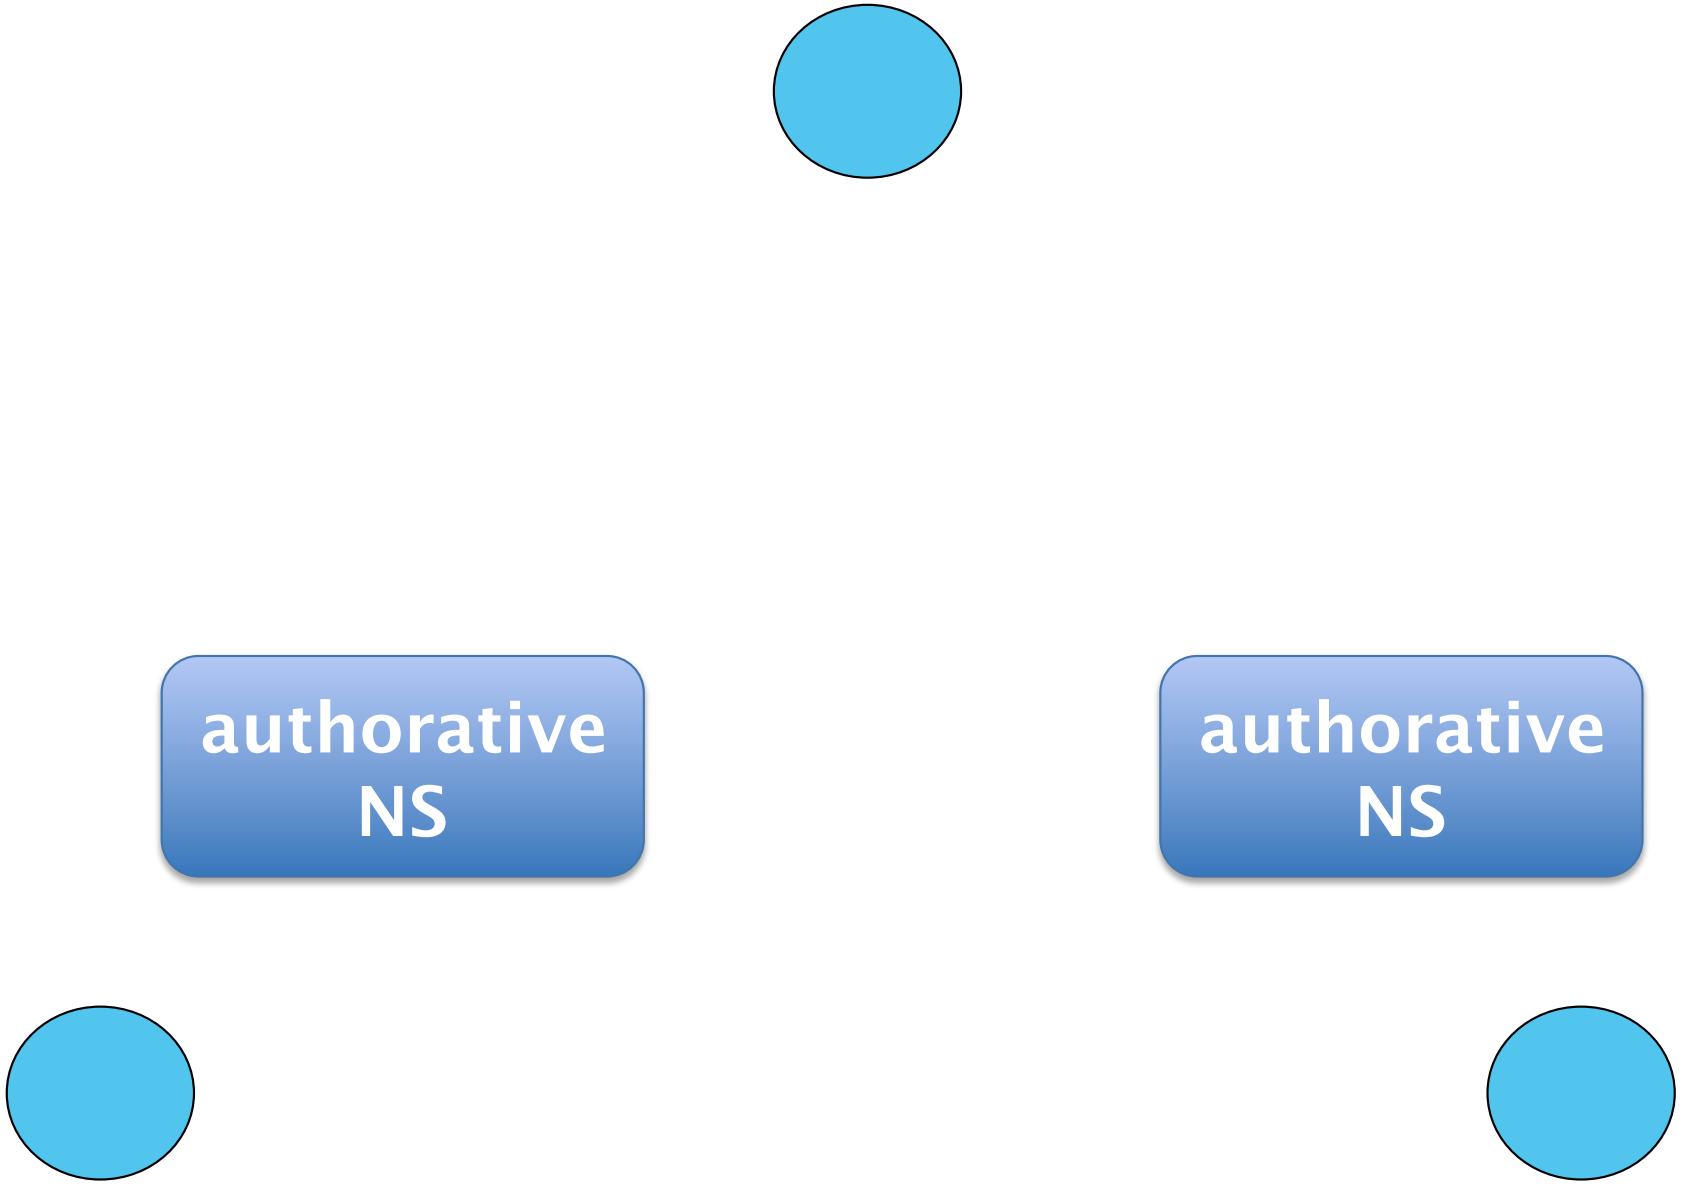
authorative NS
authorative NS

how to find a name server
close to an end host?

look at domain name, IP prefix

how likely does a name server allow recursive query?

> 72% of name server tested

how accurate is the estimated latency?

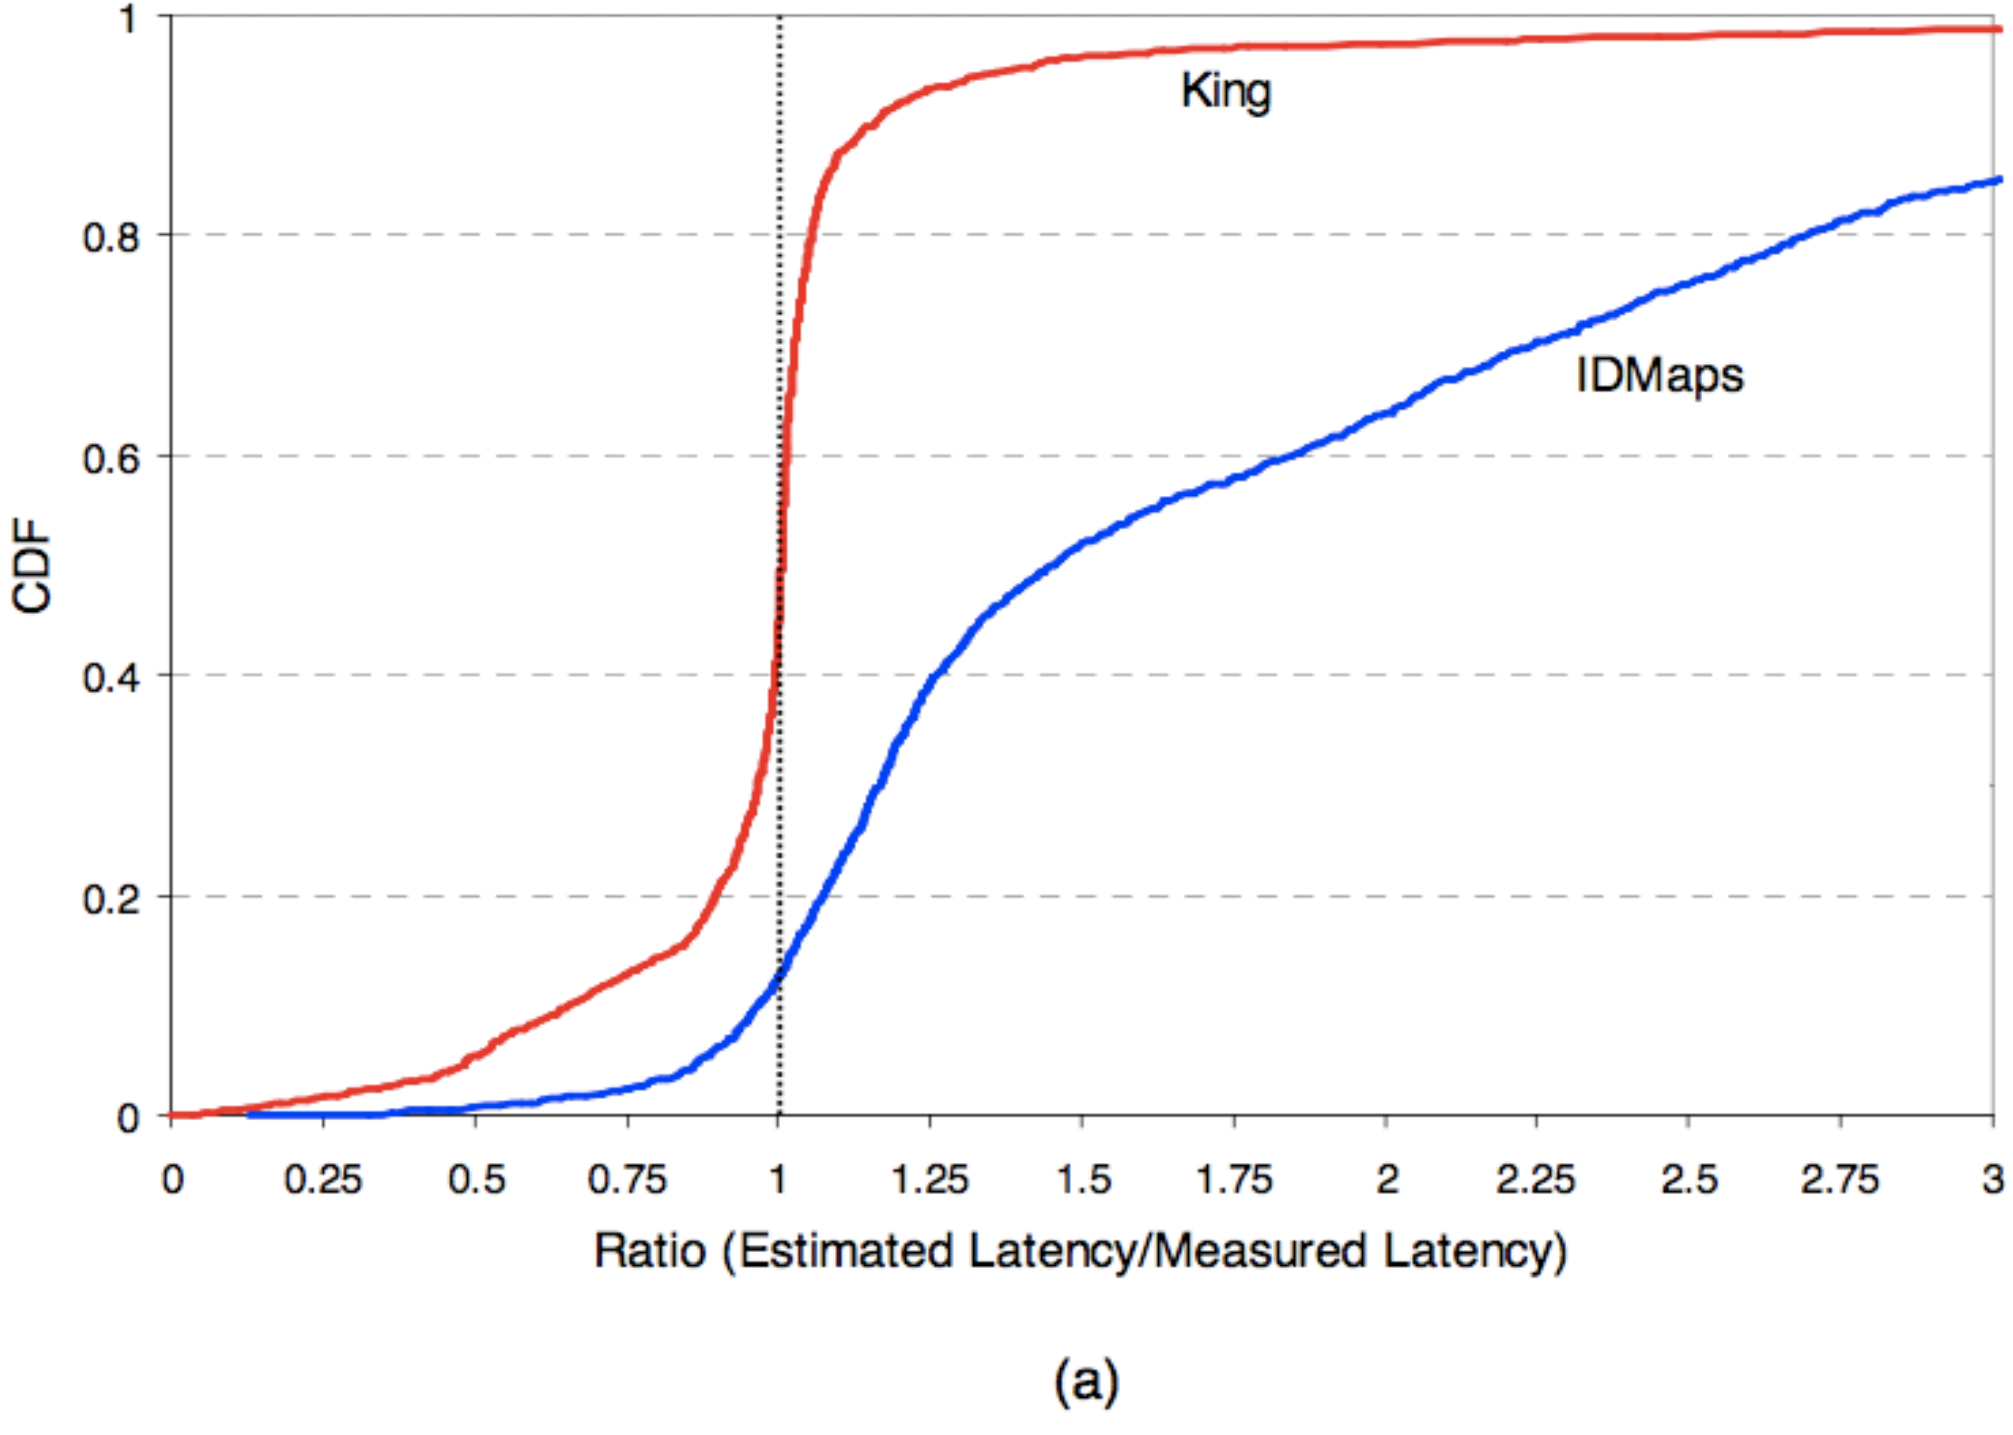
1
0.8
0.6
0.4
0.2
0
CDF
King
IDMaps
0
0.25
0.5
0.75
1
1.25
1.5
1.75
2
2.25
2.5
2.75
3
Ratio (Estimated Latency/Measured Latency)

(a)

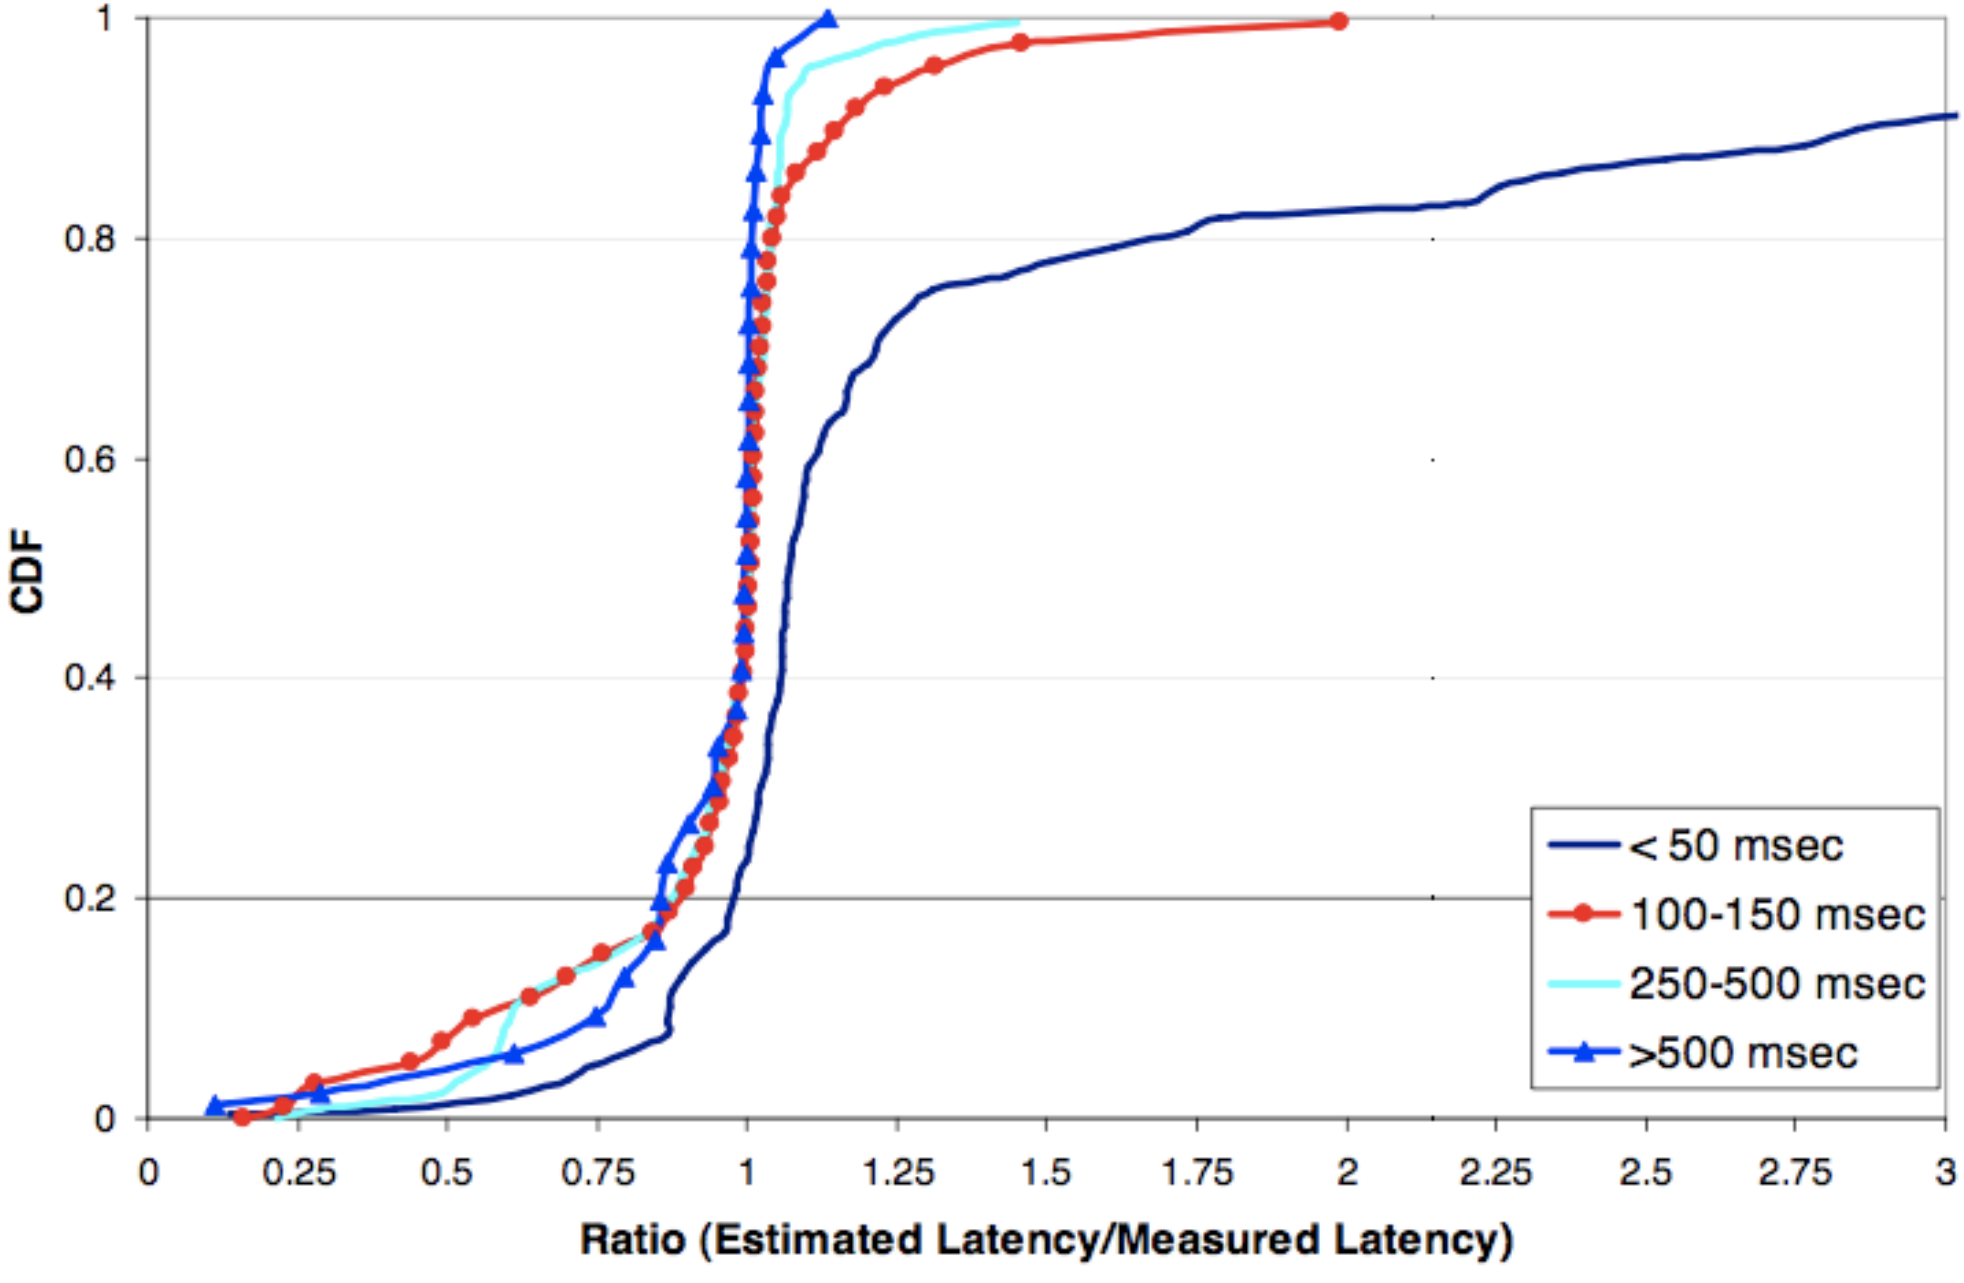
1
0.8
0.6
0.4
0.2
0
CDF
0
0.25
0.5
0.75
1
1.25
1.5
1.75
2
2.25
2.5
2.75
3
Ratio (Estimated Latency/Measured Latency)
< 50 msec
100-150 msec
250-500 msec
>500 msec

how close, really, is an end host to its authorative NS?

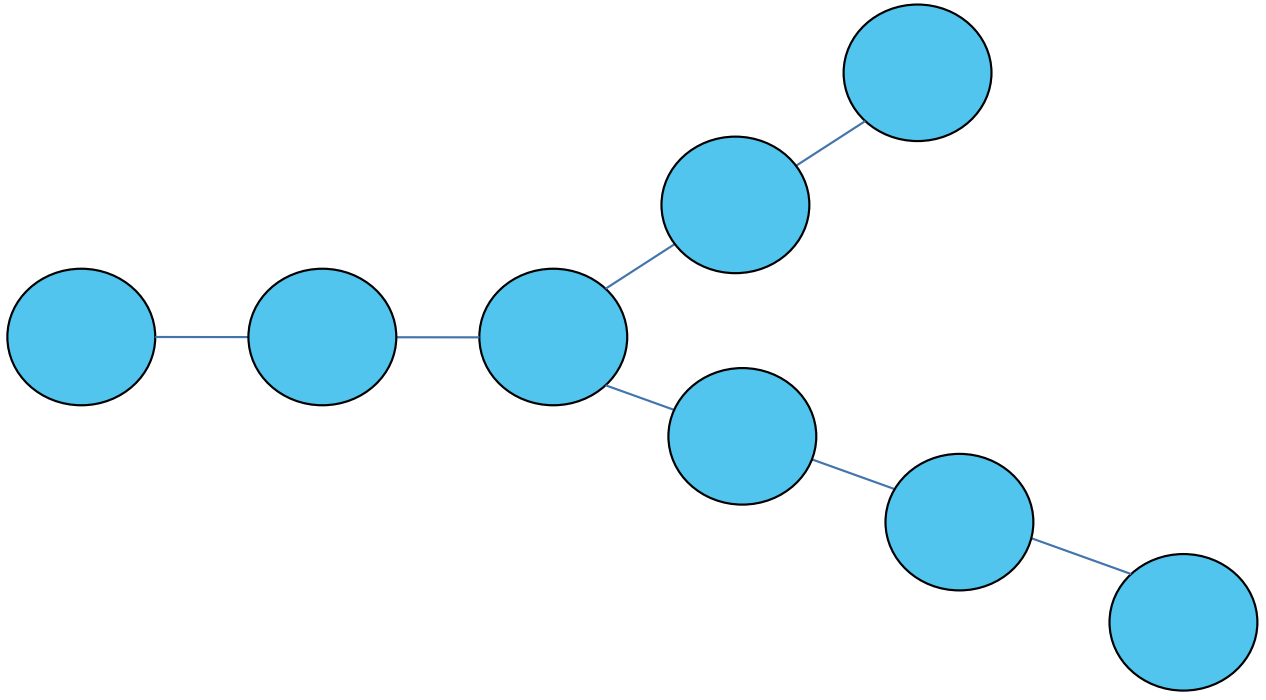

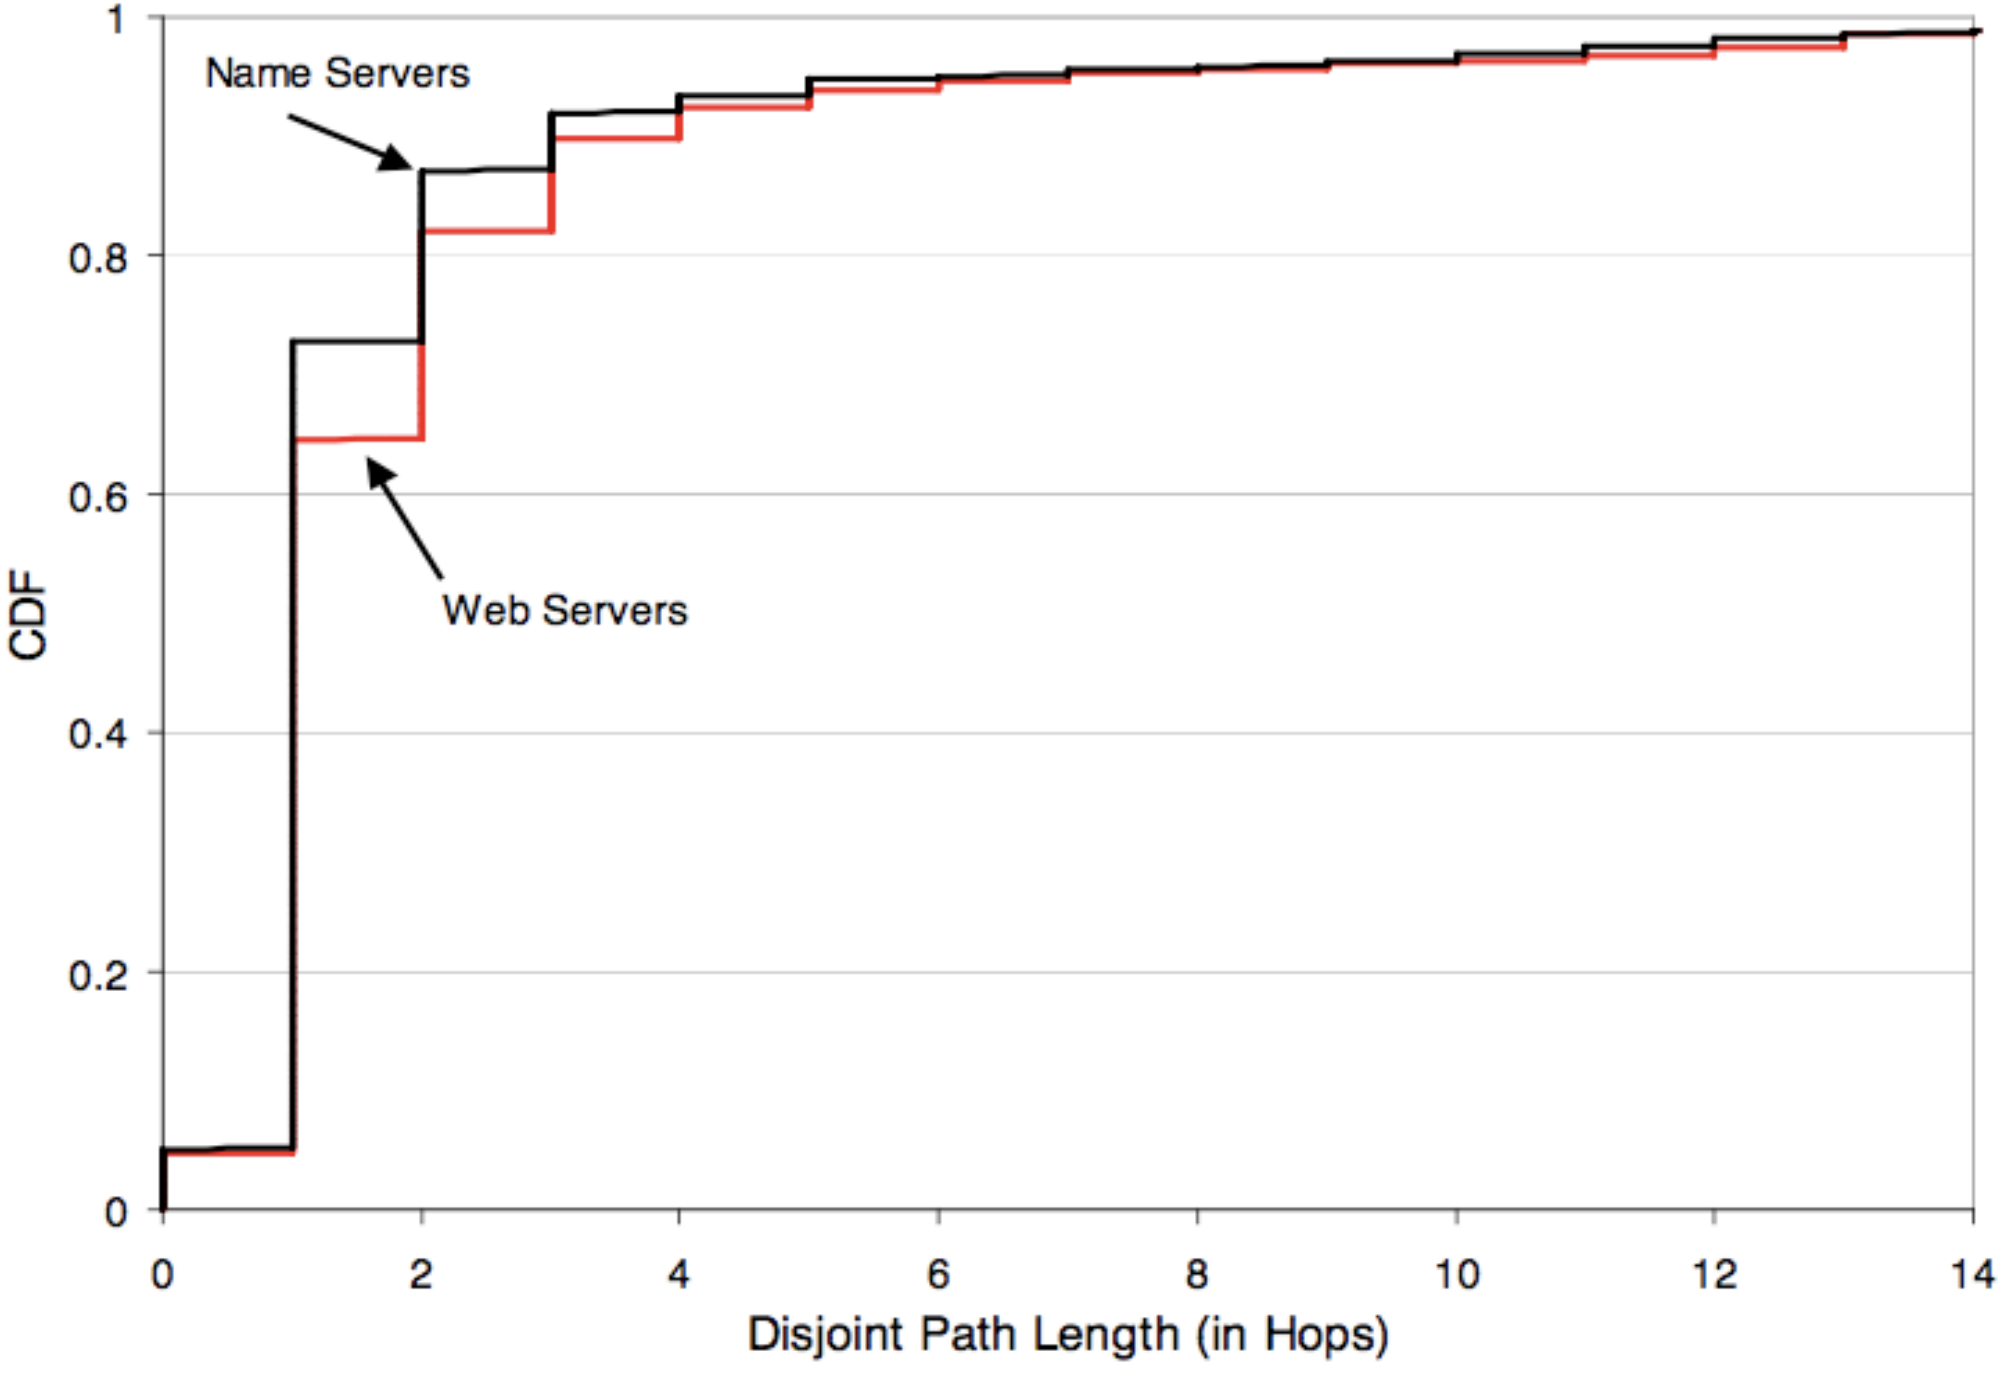
Name Servers
Web Servers
CDF
1
0.8
0.6
0.4
0.2
0
0
2
4
6
8
10
12
14
Disjoint Path Length (in Hops)

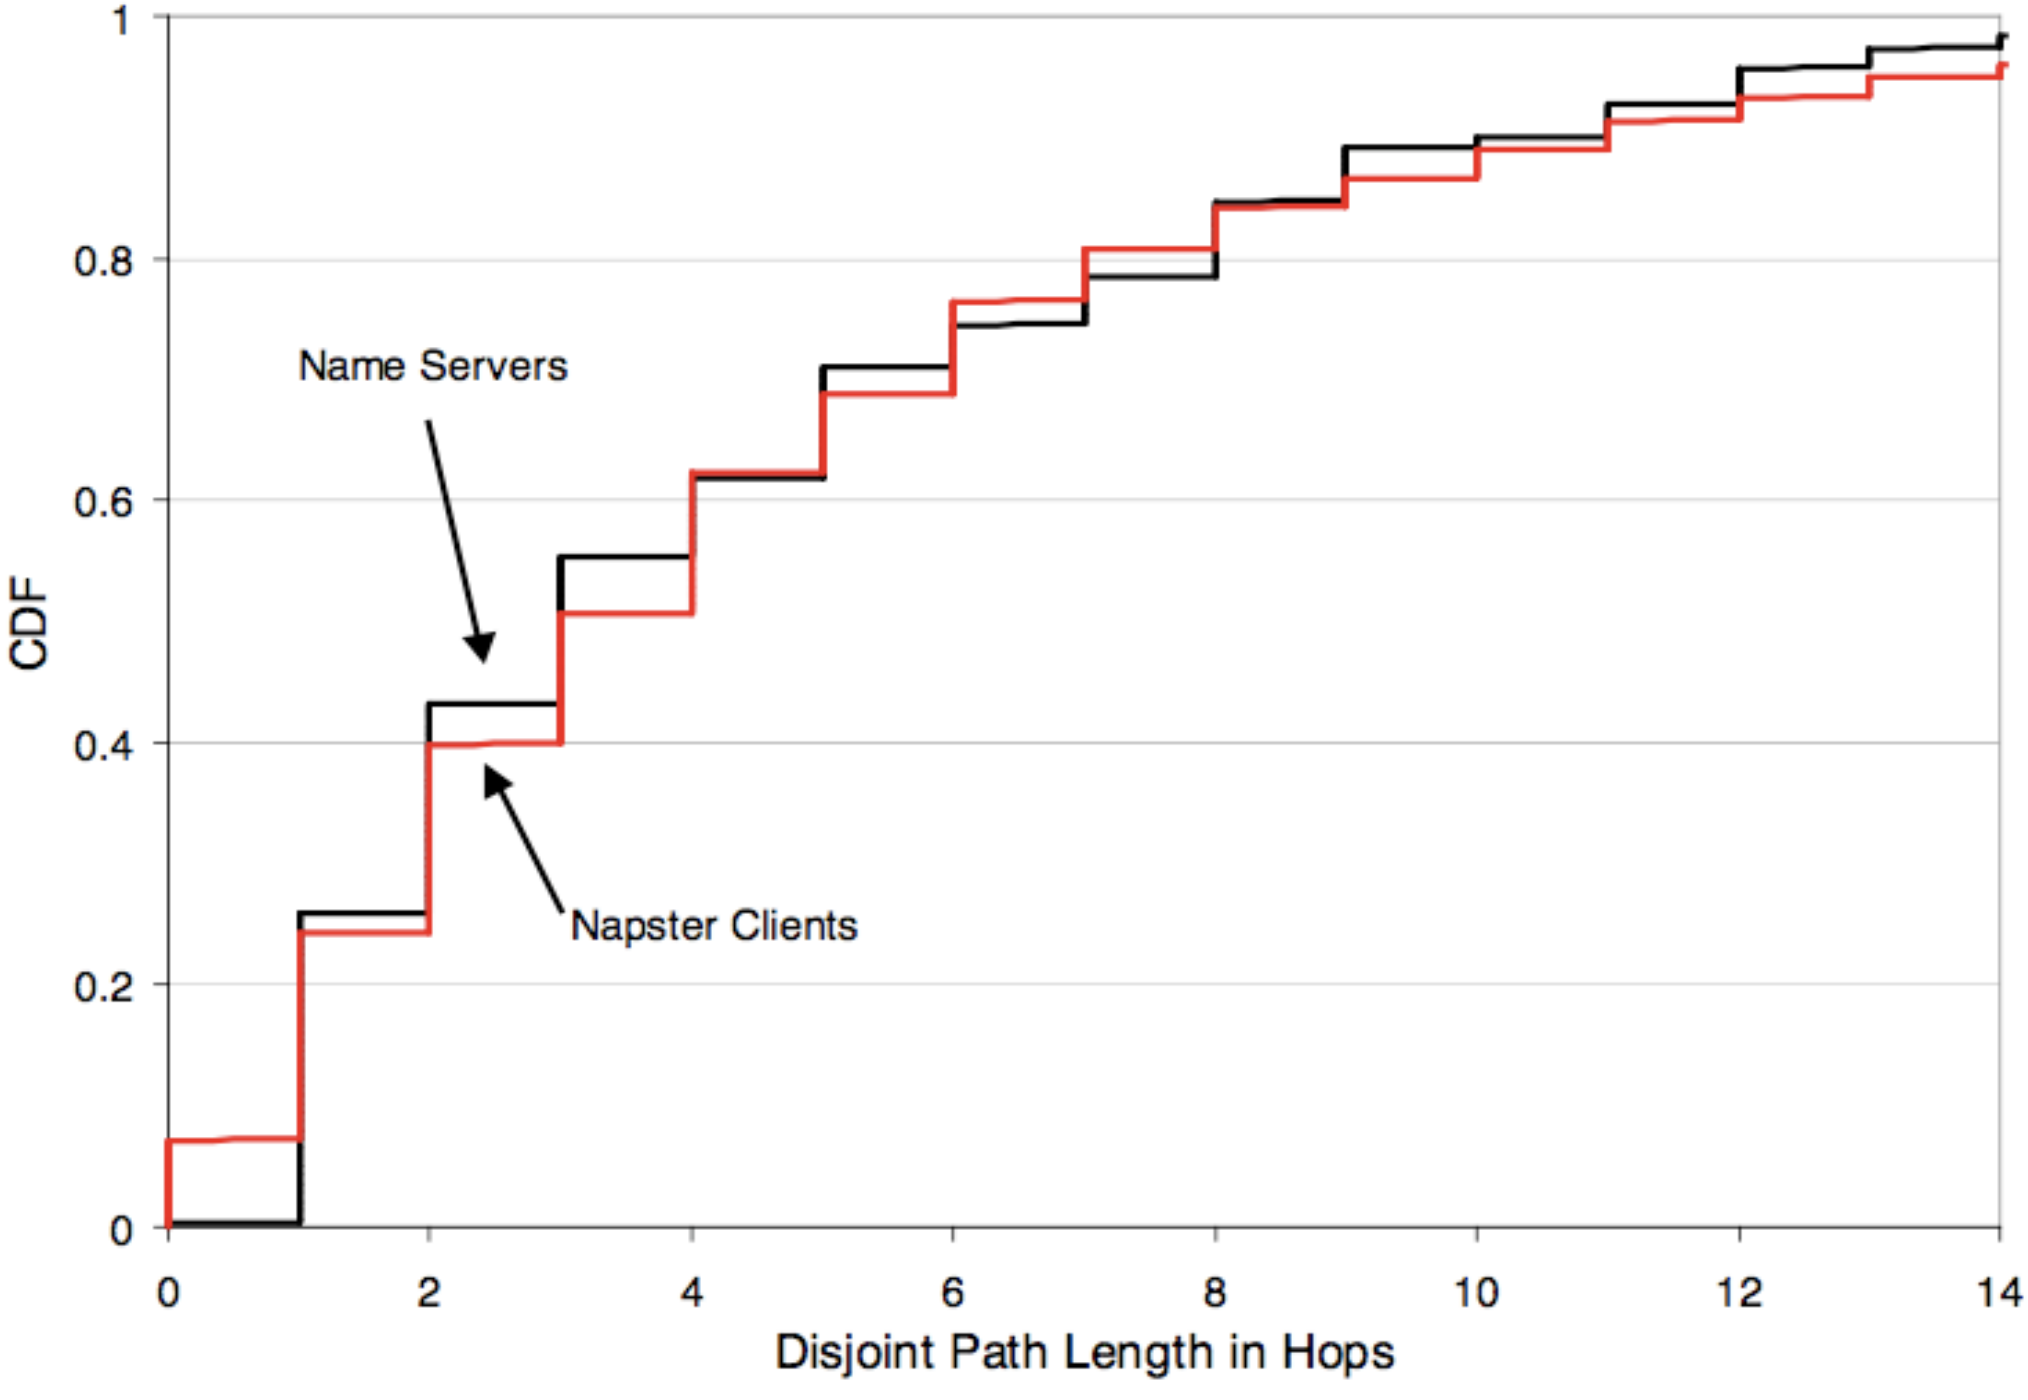
1
0.8
0.6
0.4
0.2
0
CDF
Name Servers
Napster Clients
0
2
4
6
8
10
12
14
Disjoint Path Length in Hops

many diverged paths consist of routers that are physically adjacent

75-80% of the disjoint paths have latency less than 10ms in the Napster trace

Summary:

Measurement study of DNS

Use DNS to do measurement