# Recursive Resolver Delegation Selection

Kyle Schomp

OARC 30
2019-05-12

Akamai
Experience the Edge

# How do recursive resolvers choose among delegations?

```
$dig @a.gtld-servers.net edgekey.net +noall +auth

; <<>> DiG 9.10.3-P3 <<>> @a.gtld-servers.net edgekey.net +noall +auth
; (2 servers found)
;; global options: +cmd
edgekey.net.            172800      IN          NS      ns1-66.akam.net.
edgekey.net.            172800      IN          NS      usw6.akam.net.
edgekey.net.            172800      IN          NS      adns1.akam.net.
edgekey.net.            172800      IN          NS      ns4-66.akam.net.
edgekey.net.            172800      IN          NS      ns7-65.akam.net.
edgekey.net.            172800      IN          NS      ns5-66.akam.net.
edgekey.net.            172800      IN          NS      a6-65.akam.net.
edgekey.net.            172800      IN          NS      a12-65.akam.net.
edgekey.net.            172800      IN          NS      a5-65.akam.net.
edgekey.net.            172800      IN          NS      a16-65.akam.net.
edgekey.net.            172800      IN          NS      a18-65.akam.net.
edgekey.net.            172800      IN          NS      a28-65.akam.net.
edgekey.net.            172800      IN          NS      a13-65.akam.net.
```

13 NS records with accompanying A/AAAA records

# How do recursive resolvers choose among delegations?

| | |
|---|---|
| Specific resolver software in the lab | Yu, Yingdi, et al. "Authority server selection in DNS caching resolvers." *ACM SIGCOMM Computer Communication Review* 42.2 (2012): 80-86. |
| Probing resolvers on the Internet | Müller, Moritz, et al. "Recursives in the wild: engineering authoritative DNS servers." *Proceedings of the 2017 Internet Measurement Conference*. ACM, 2017. |
| Resolvers on the Internet with real traffic | This talk… |

# Reasons why we want to know

1. Informs decisions made in authoritative nameserver deployments

2. Knowing the limitations in common behavior among recursive resolvers can motivate improvements in that behavior

# Dataset

- 10min of Akamai's authoritative DNS servers' logs
  - Queries for CDN domain `edgekey.net`
  - IPv4 traffic only
  - Assume latency stable over short interval
- Repeated experiments
  - Different times
  - Different CDN domains
  - Similar findings
- Ping each source IP address in logs of authoritative DNS server

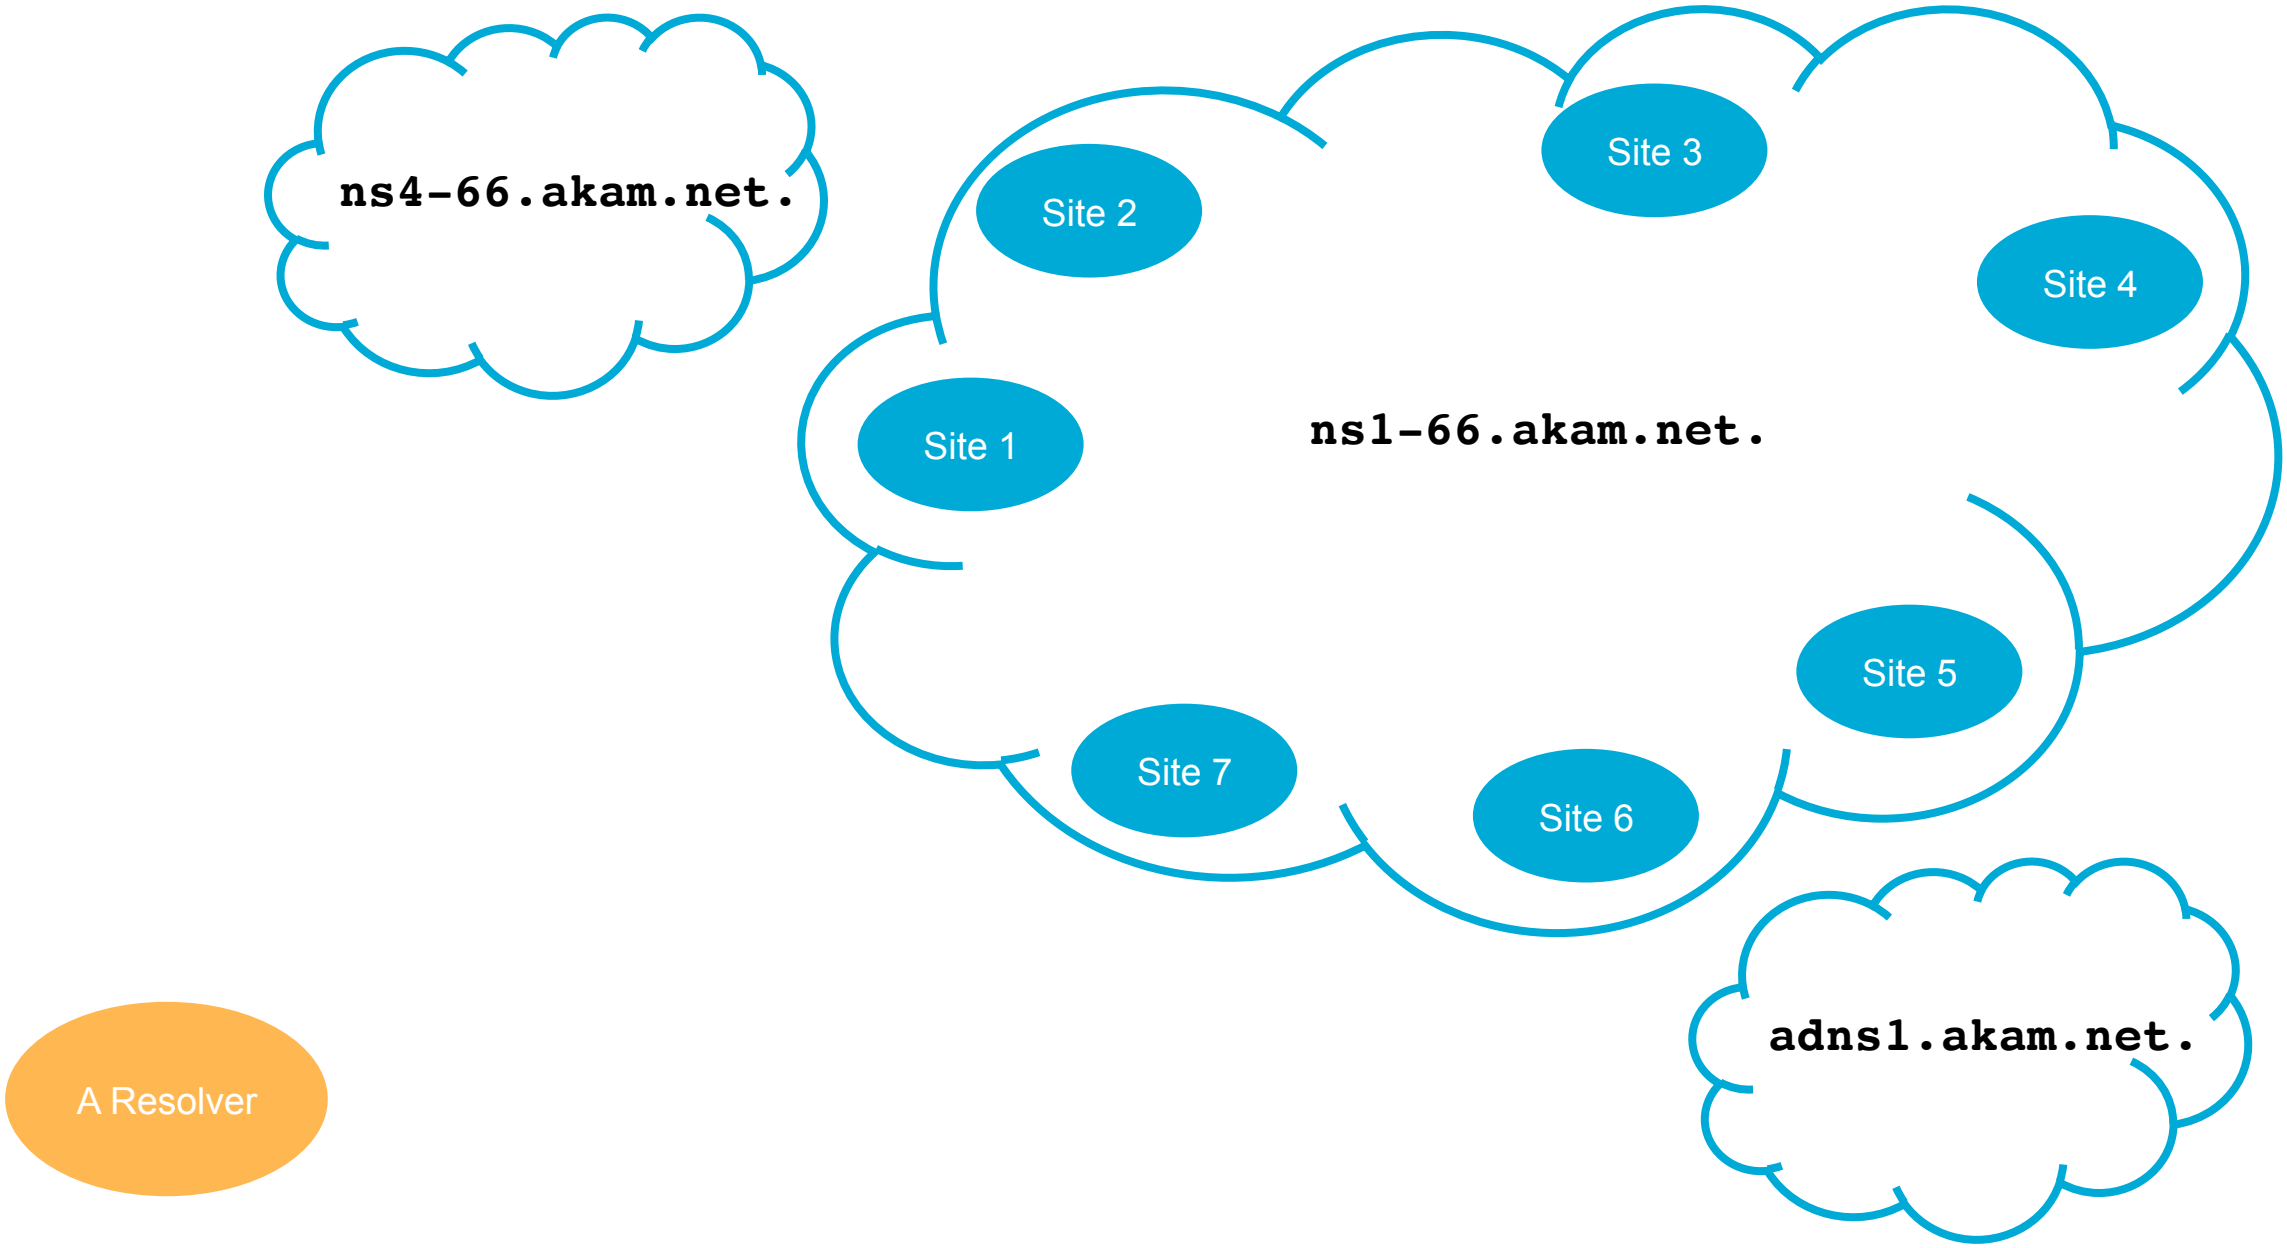


 

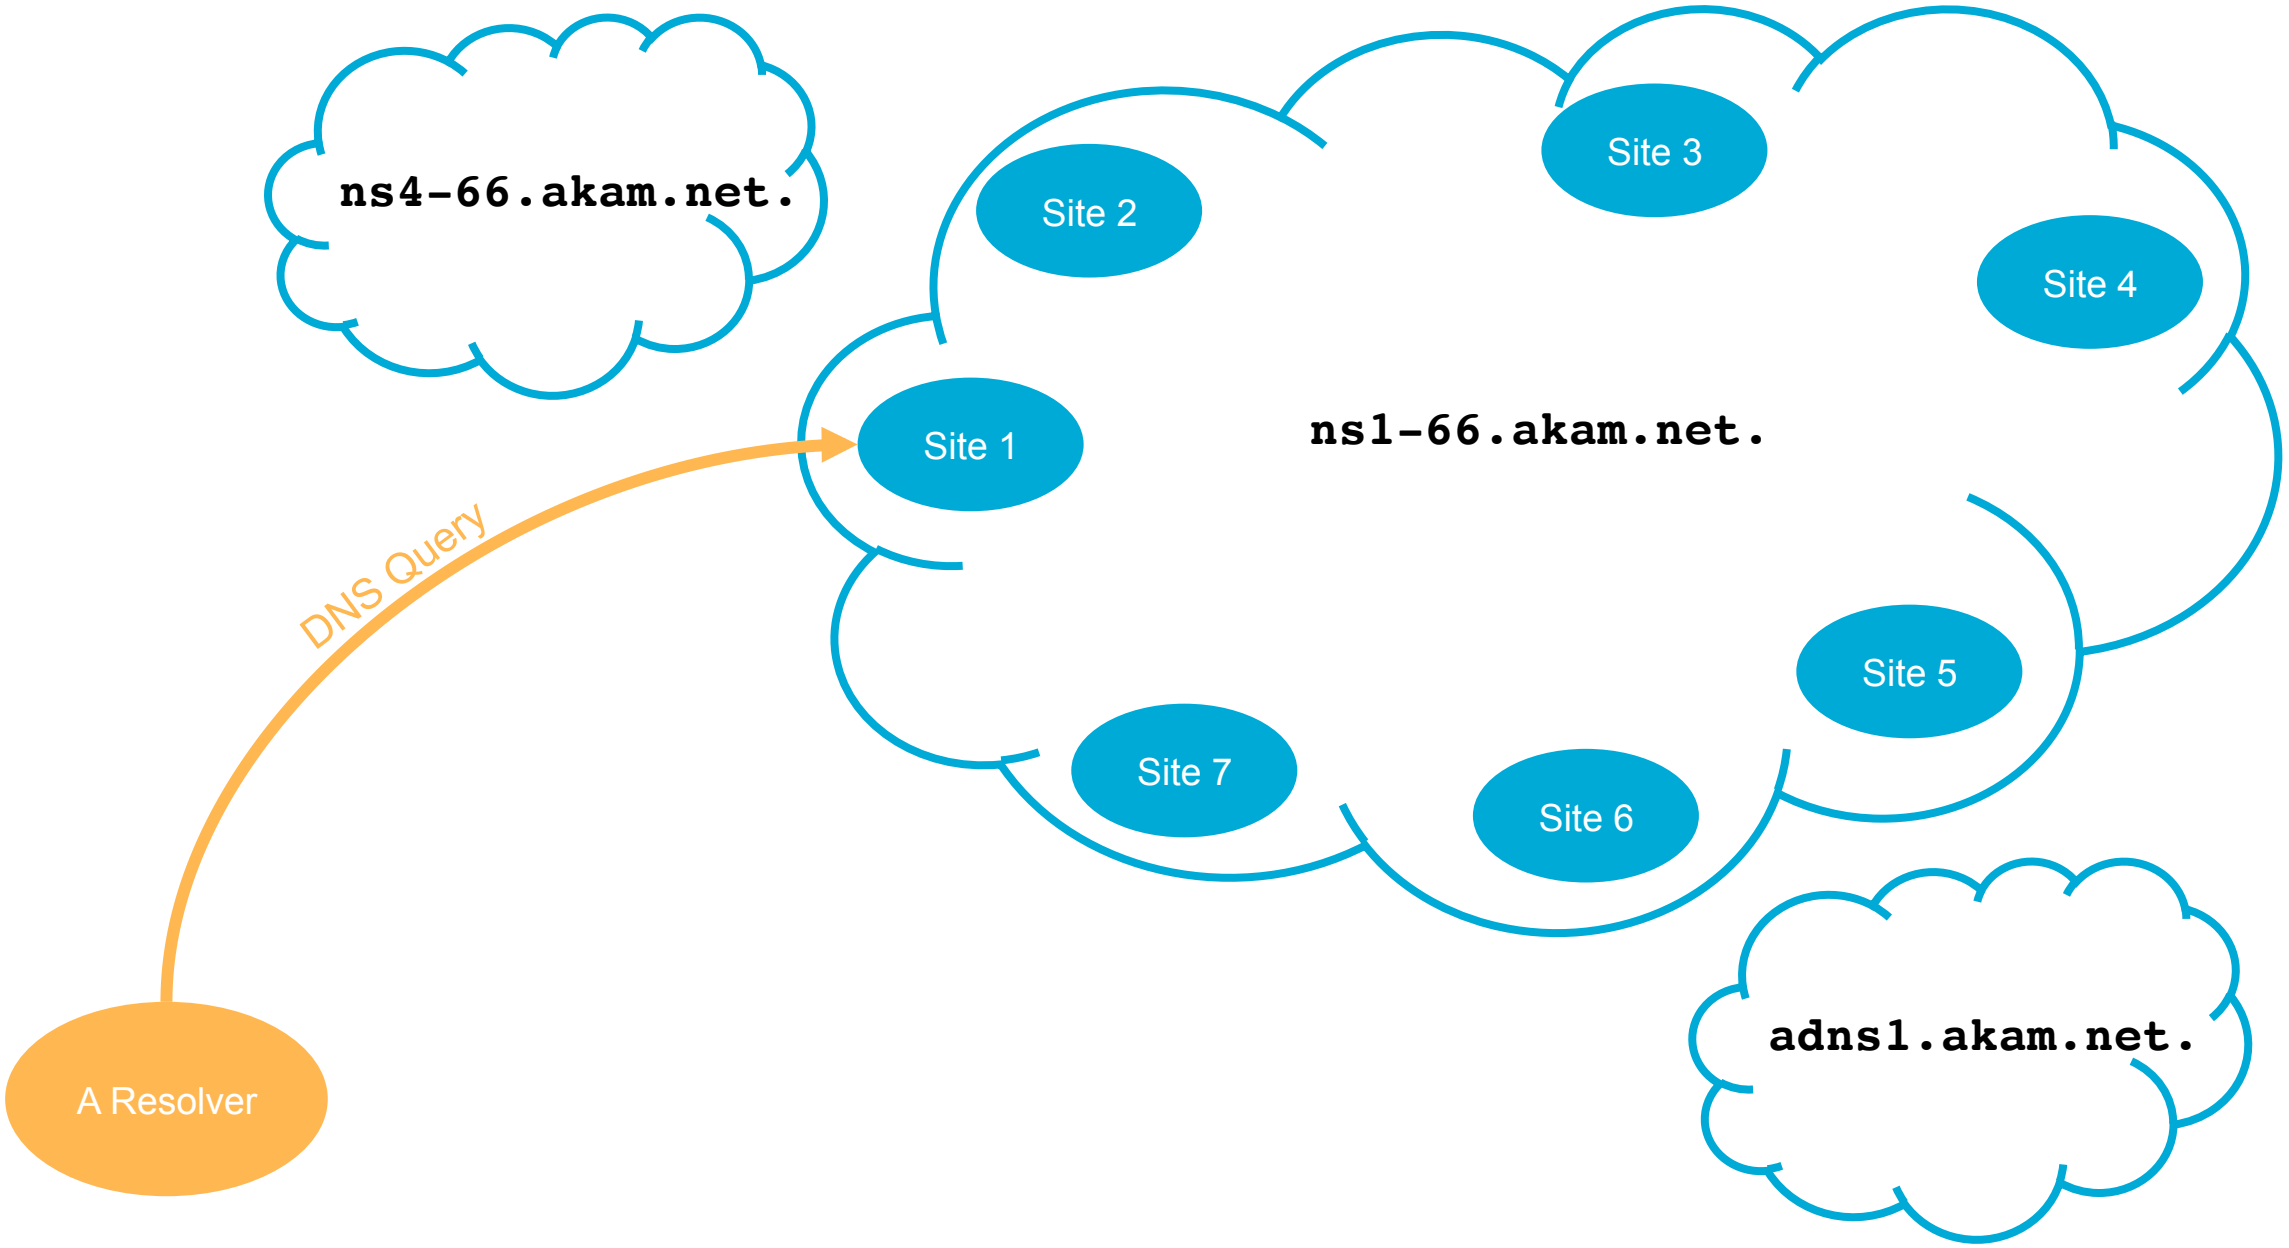


 

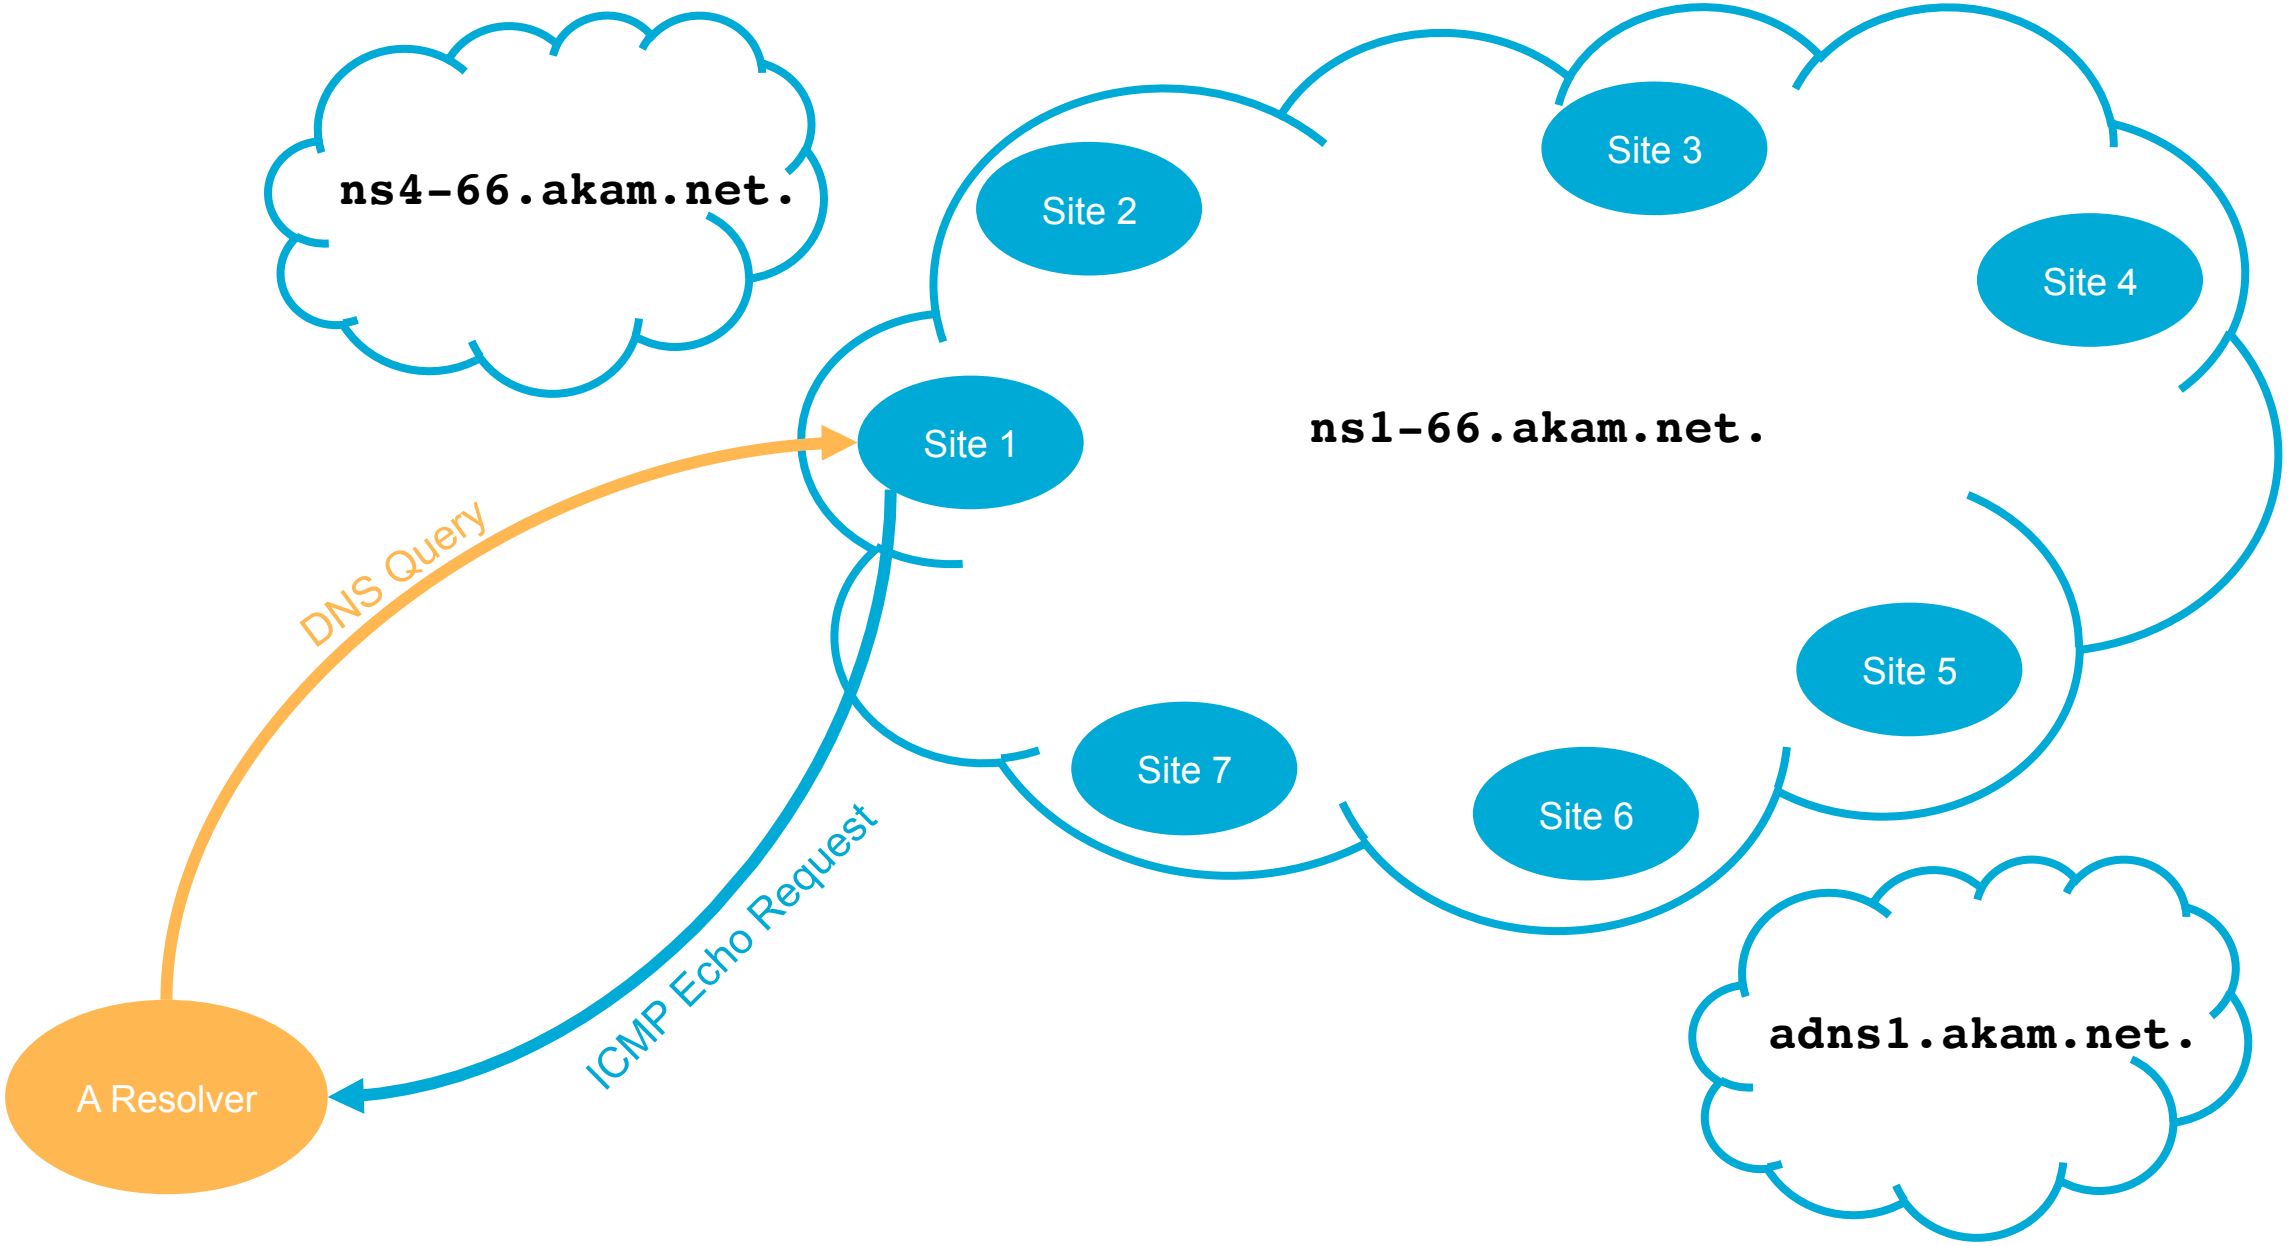


 

# Dataset

- 890k source IP addresses
  - ~89k with ≥ 90 DNS queries
- 66% of resolvers with ≥ 90 DNS queries responded to ping

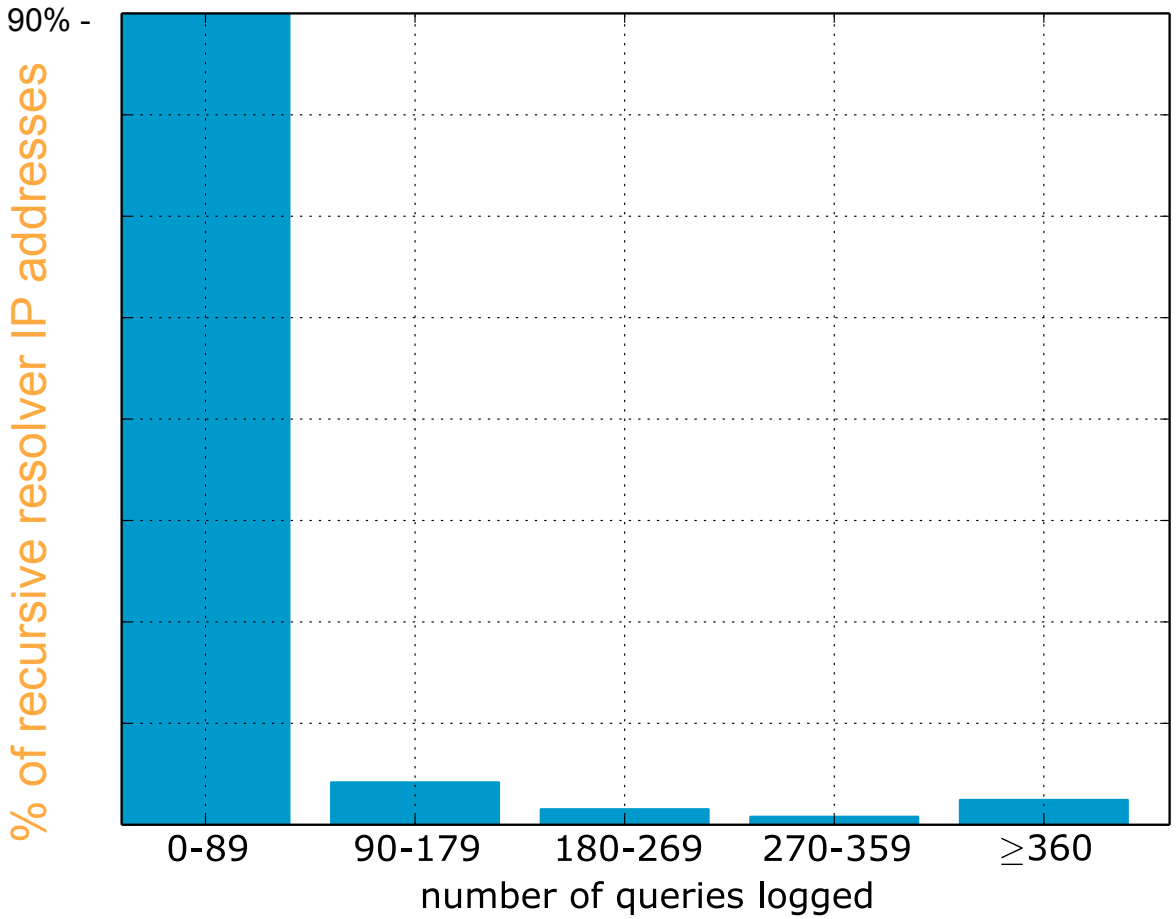

# Dataset

- 890k source IP addresses
  - ~89k with ≥ 90 DNS queries
- 66% of resolvers with ≥ 90 DNS queries responded to ping

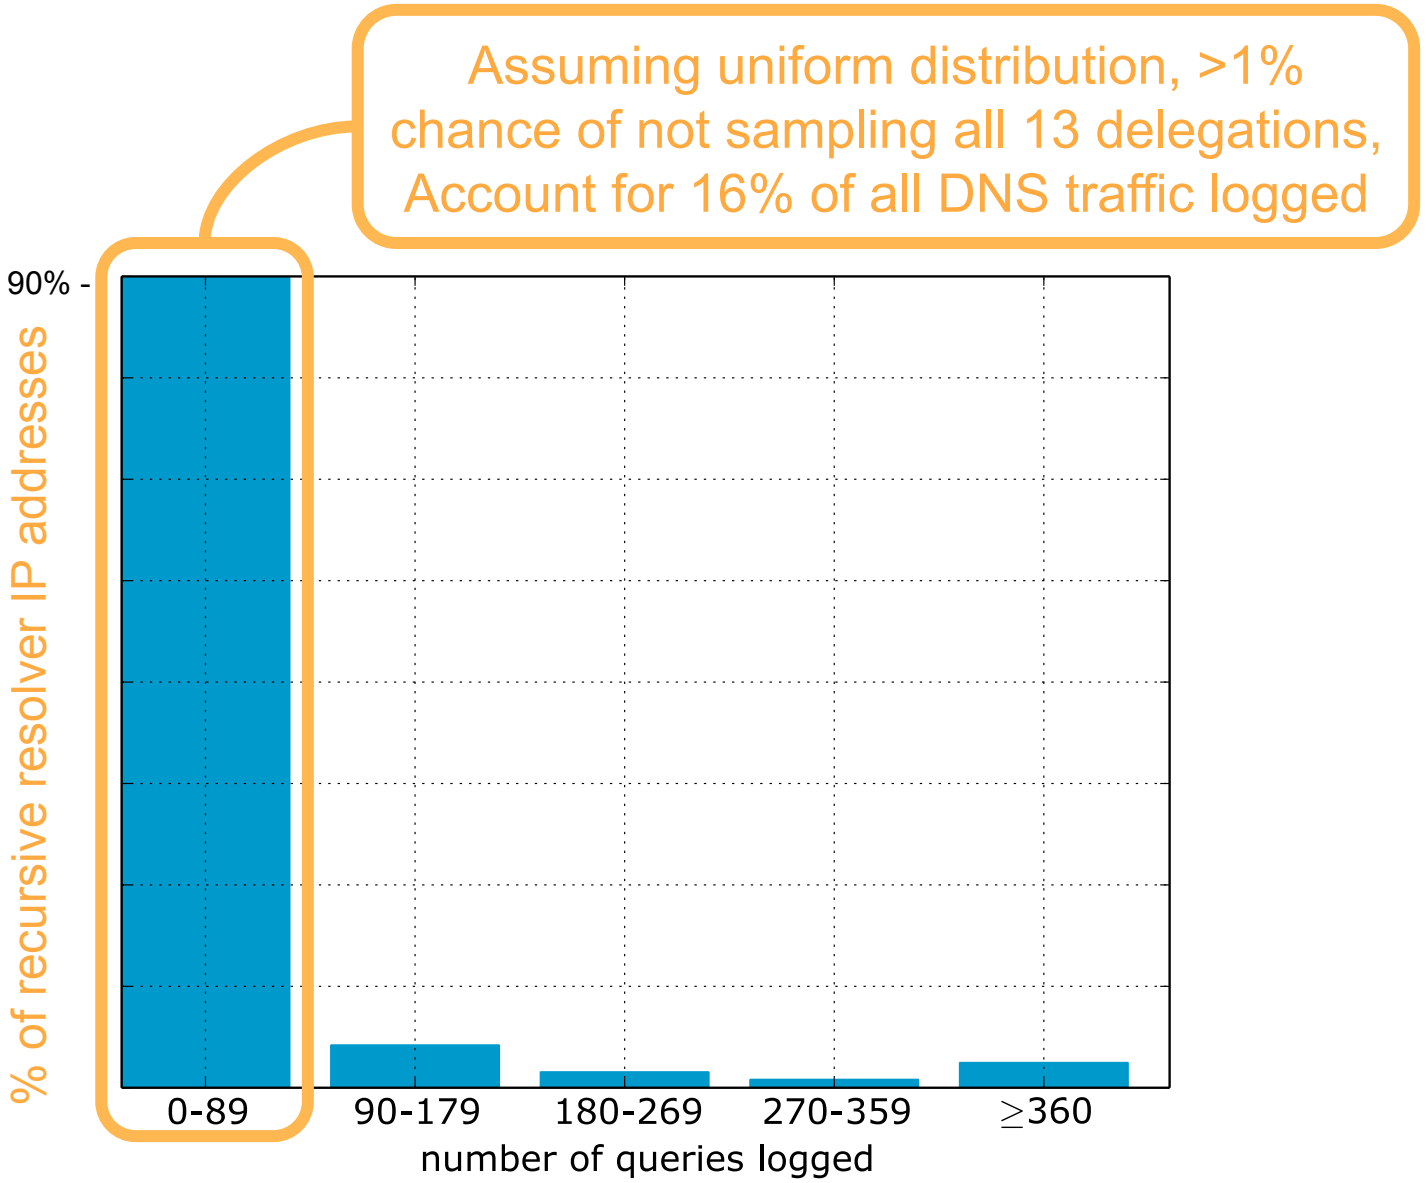


 

# Dataset

- 890k source IP addresses
  - ~89k with ≥ 90 DNS queries
- 66% of resolvers with ≥ 90 DNS queries responded to ping

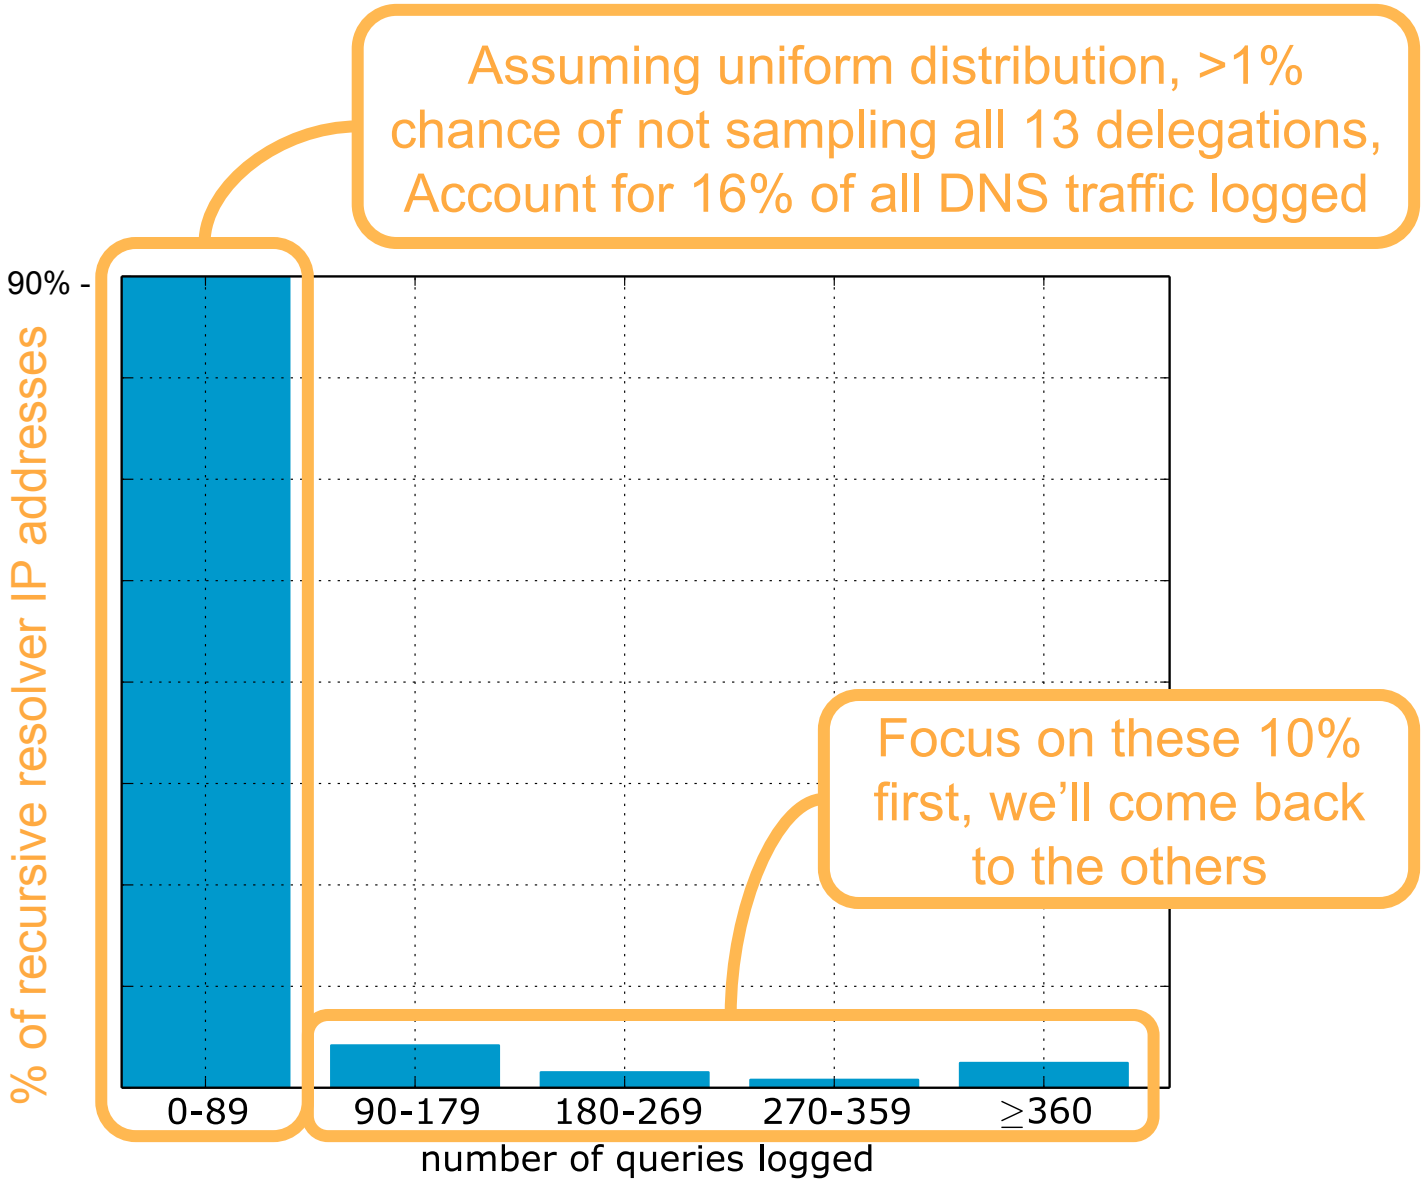


 

# # of delegations used

1. 25% of resolvers query all delegations
2. No obvious limit on the number of delegations used

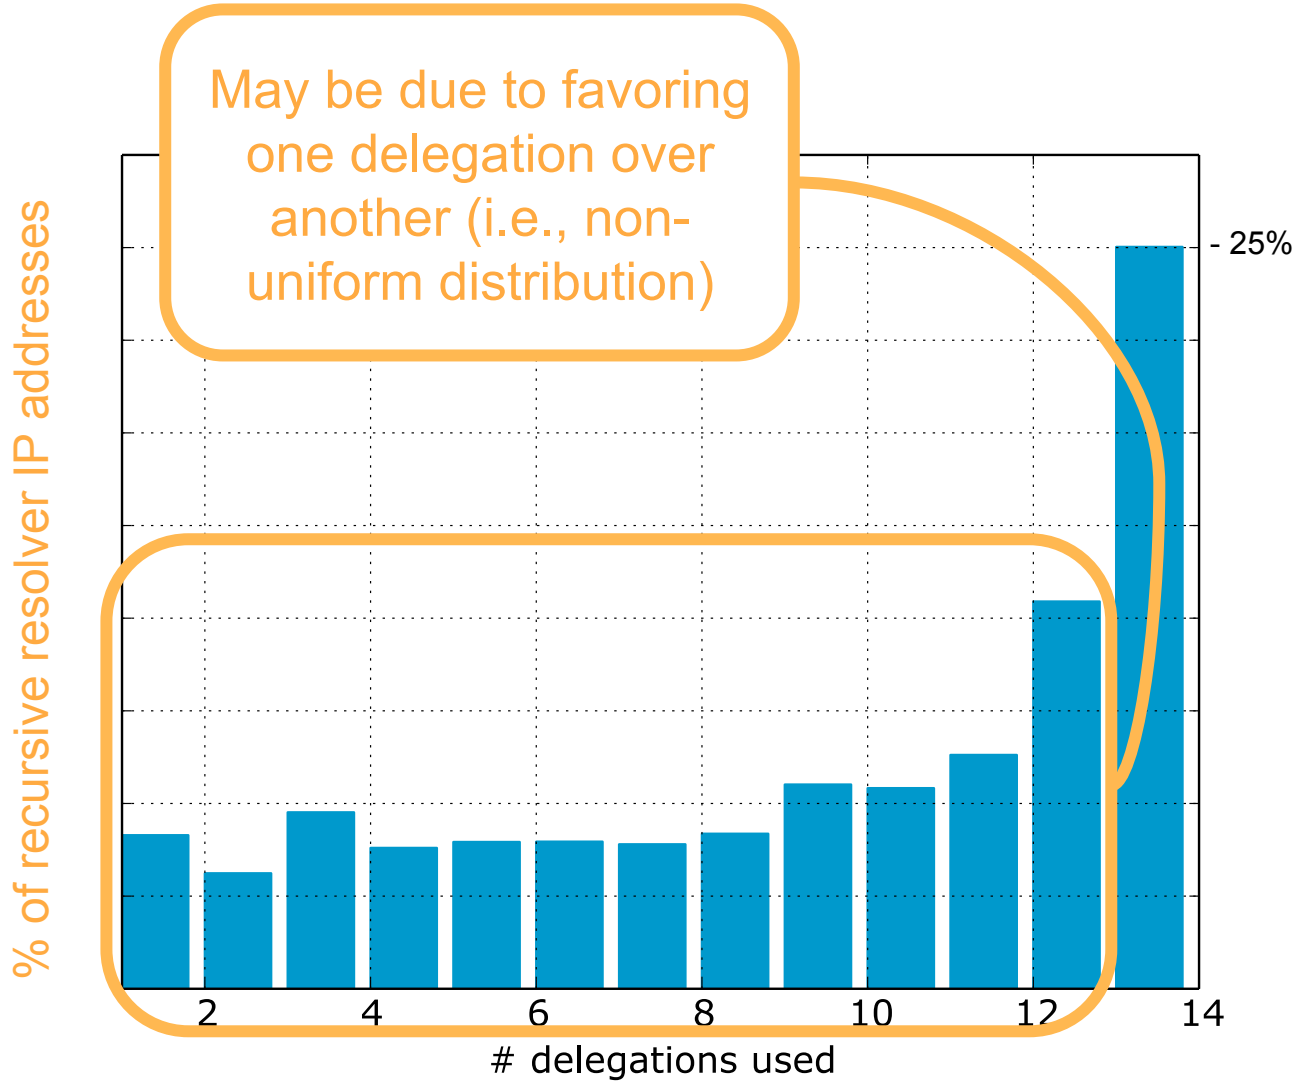


 

# Uniform distribution of queries among delegations

- Use $\chi^2$ test for uniformity in queries per delegation
  - ~1.7% resolvers potentially uniform
- Assume that non-queried delegations are excluded and that resolver uniformly selects among queried subset
  - ~6.7% resolvers potentially uniform
- Bounded, real answer likely between

 

# Resolvers using a single delegation for nearly all traffic

- Choice may be random or by latency
- 5% only ever queried a single delegation
  - Cannot tell whether using lowest latency delegation
- Nearly ¾ of others use the fastest delegation for nearly all traffic

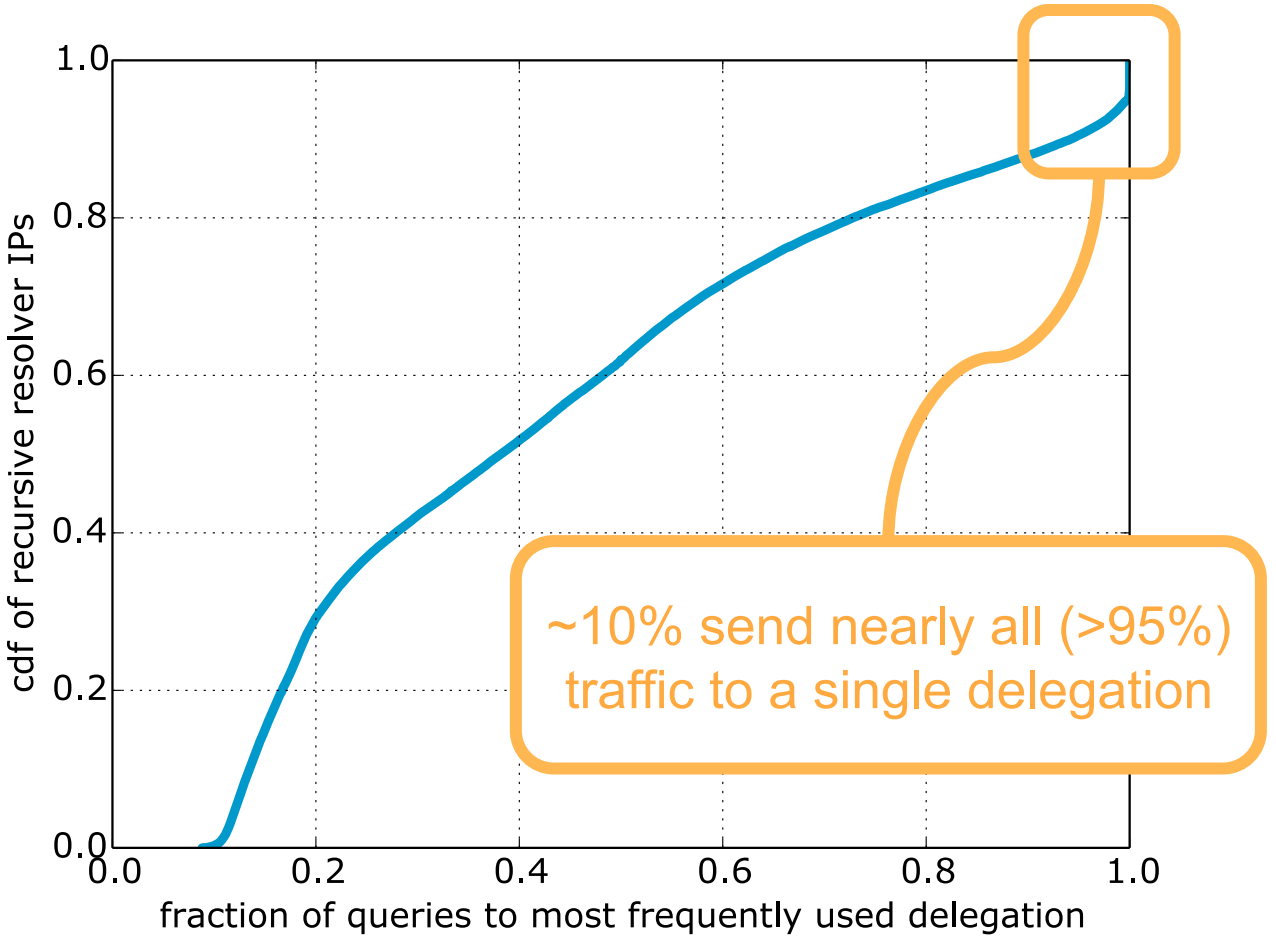

# What about the other ~83% of resolvers?

- Neither uniform nor single delegation
- Uneven distribution of queries among the 13 delegations
- Measure of unevenness – Shannon entropy

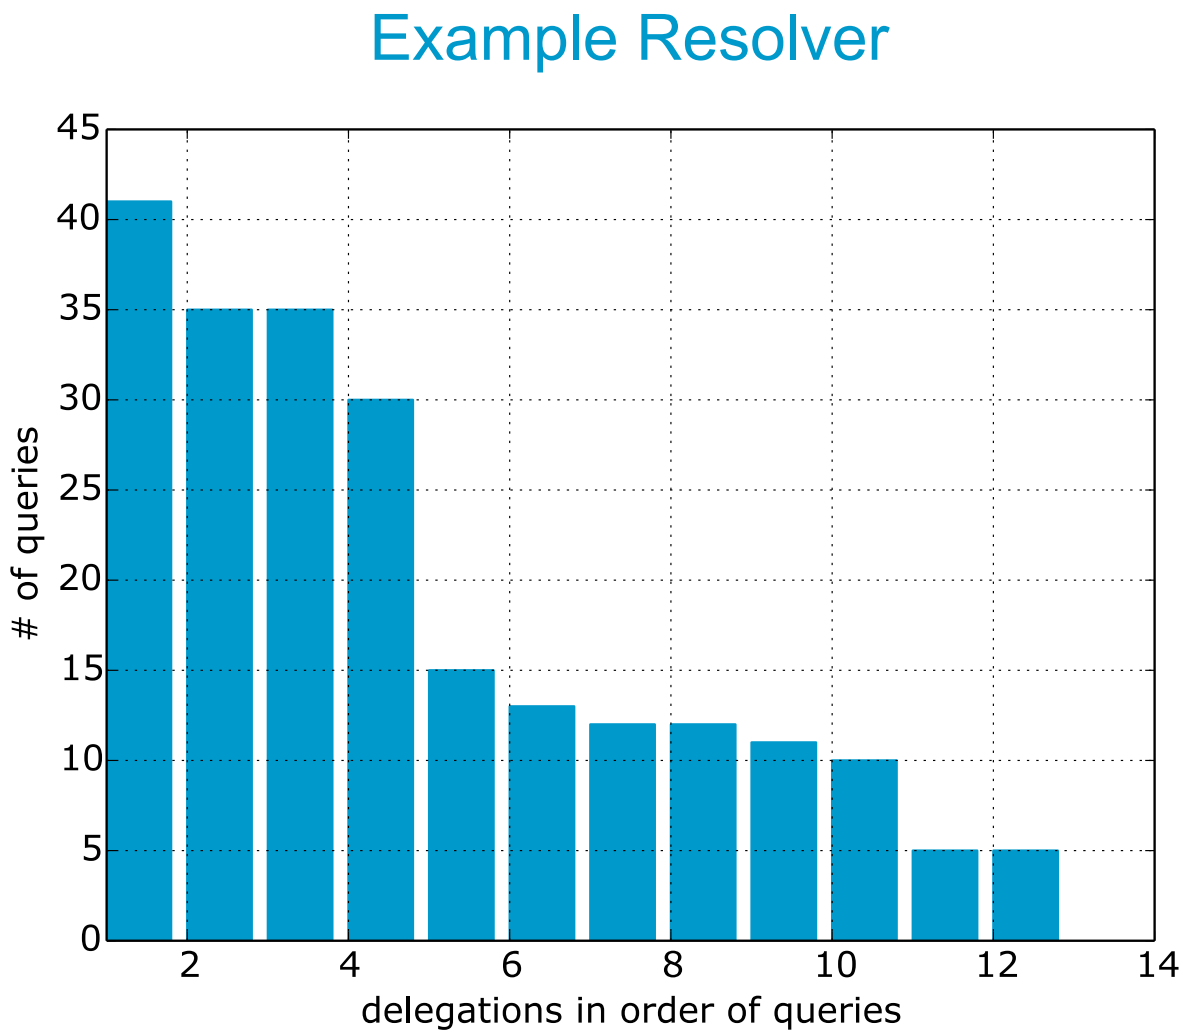


 

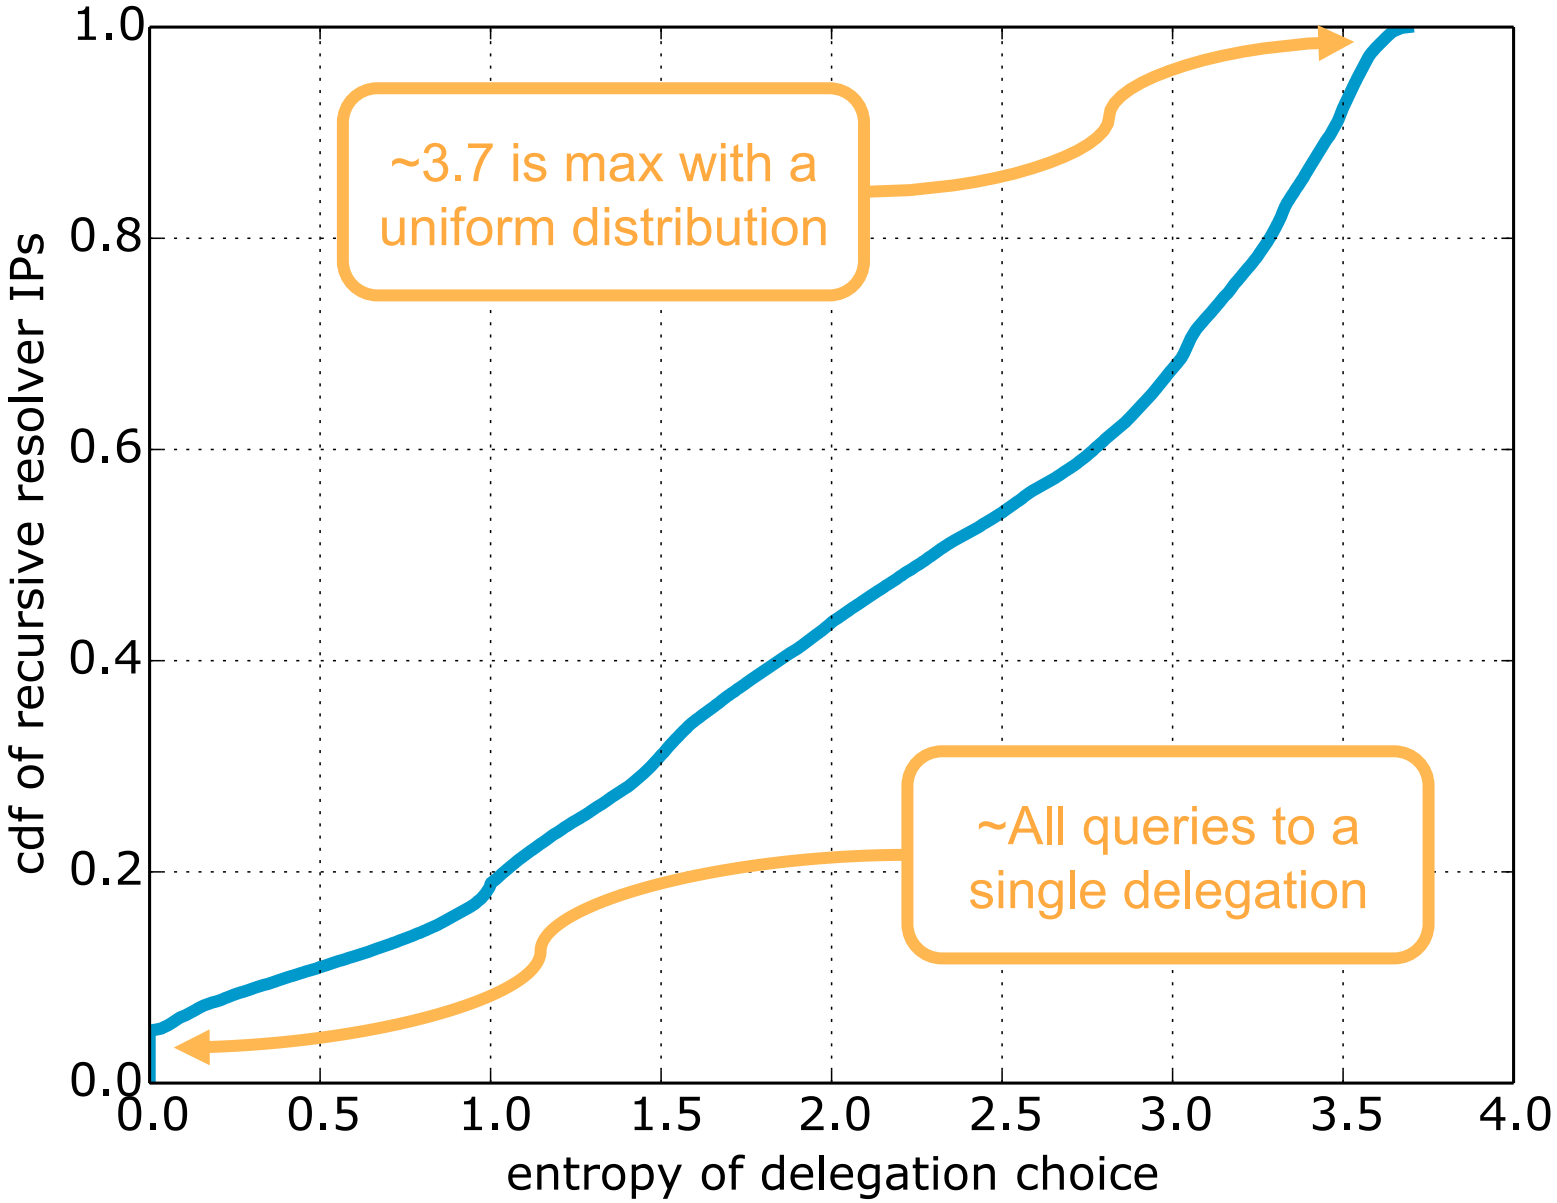


- Wide variation in distribution
- Clear preference for some delegations over others but degree of preference varies

 

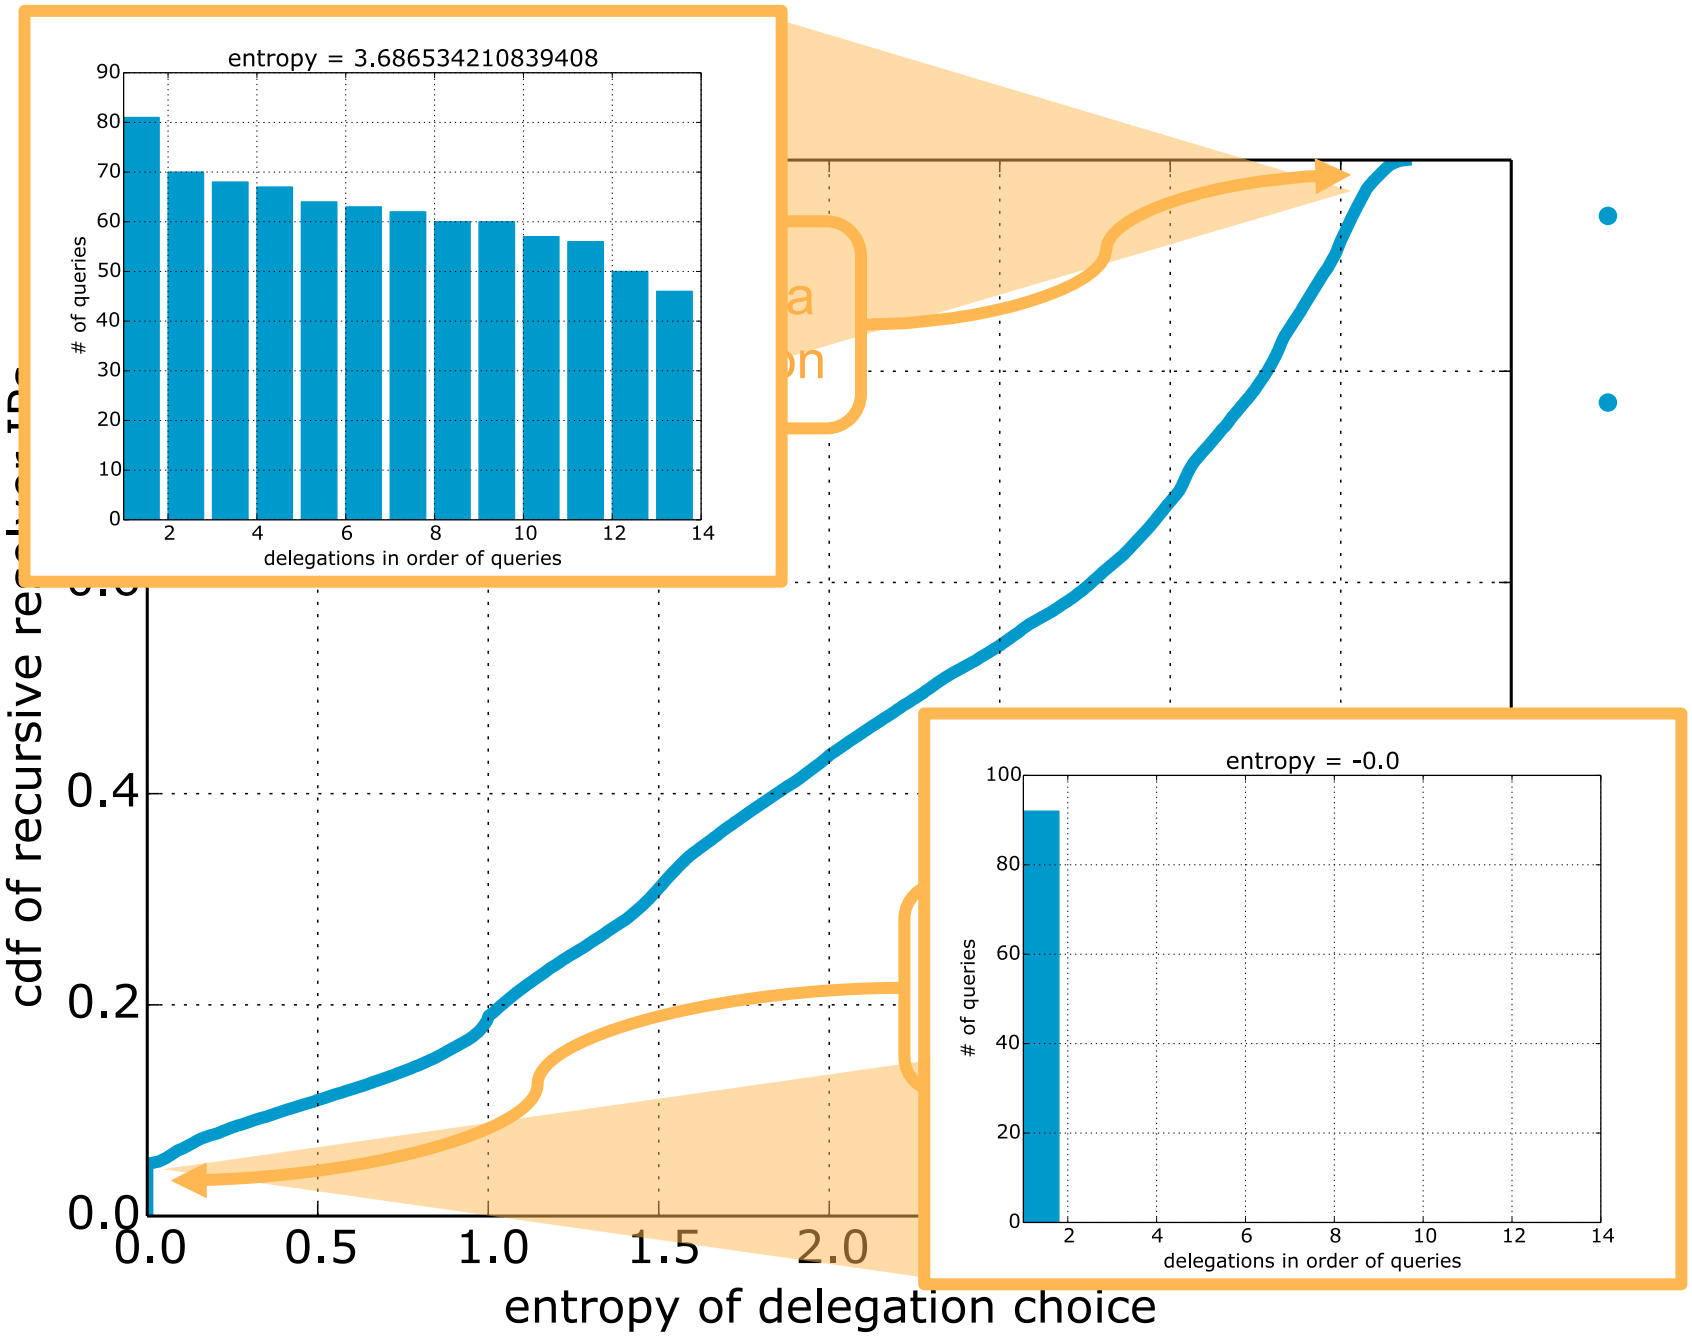


- Wide variation in distribution
- Clear preference for some delegations over others but degree of preference varies

 

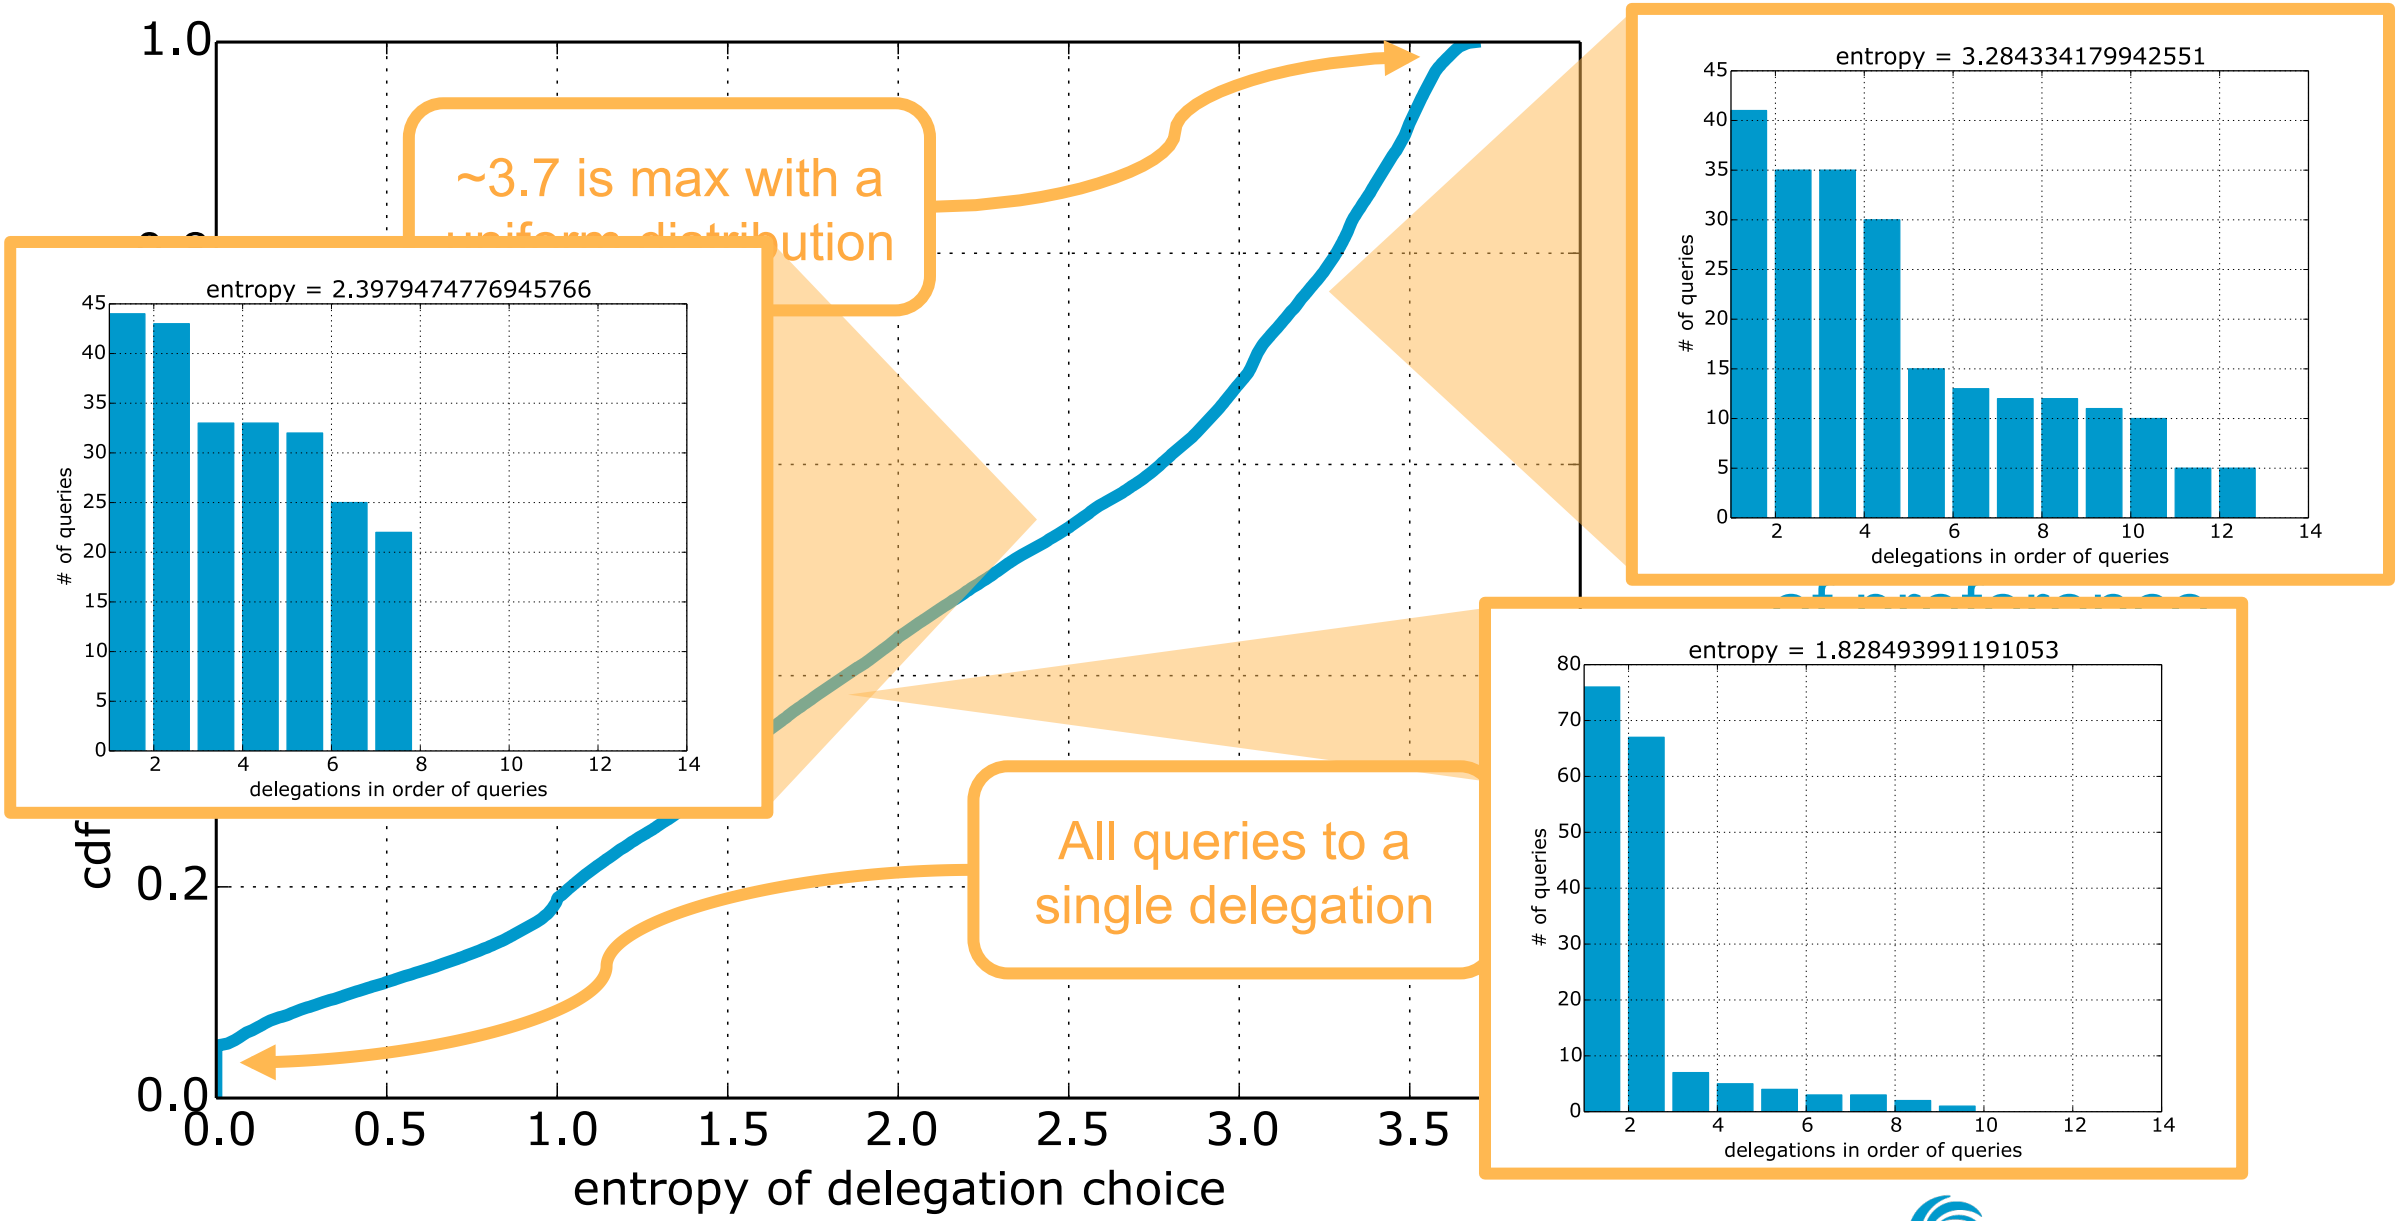


 

# Delegation usage dependent upon latency

- Previous research has shown that some recursive resolver software selects delegation by weights inversely proportional to estimated RTT
  - $W \sim 1/RTT$
- Relationship may not be linear
- Discovering RTT
  - Each DNS query is an opportunity to measure RTT
  - Un-queried delegations have unknown RTT

 

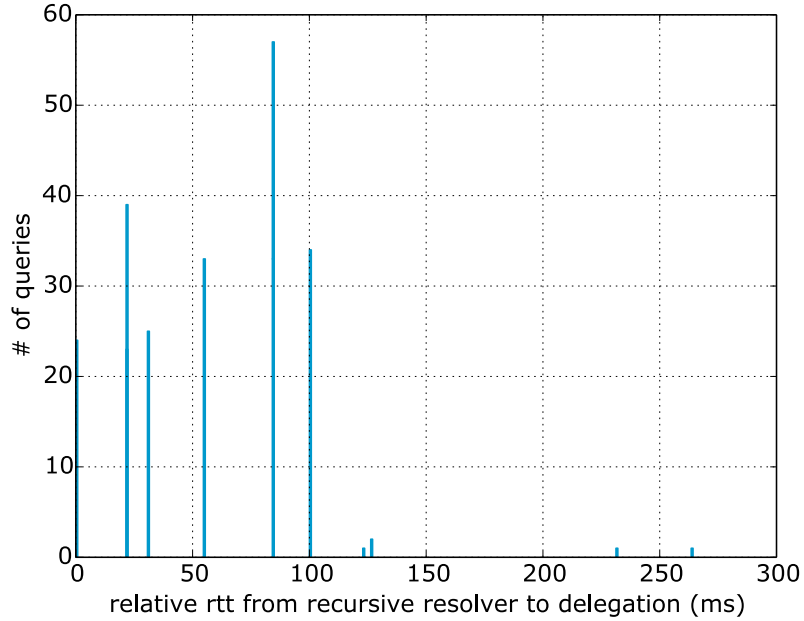


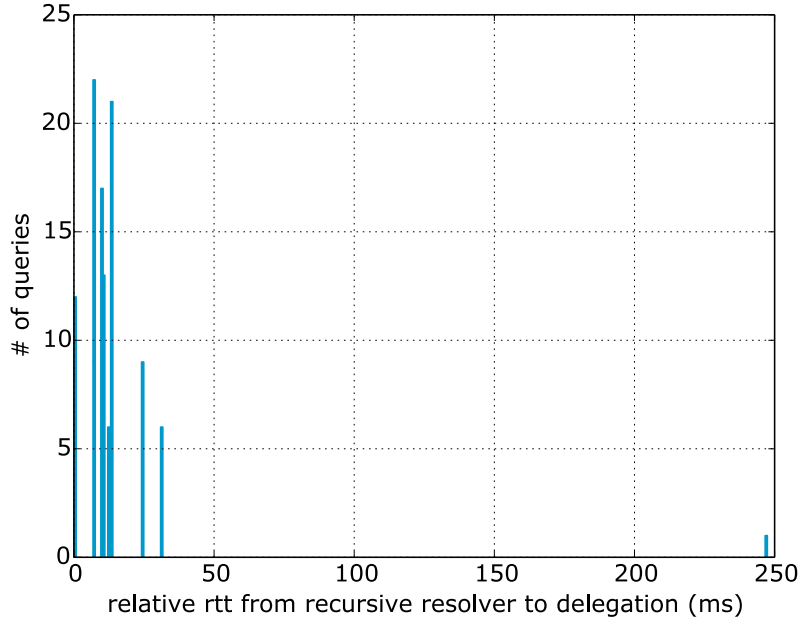


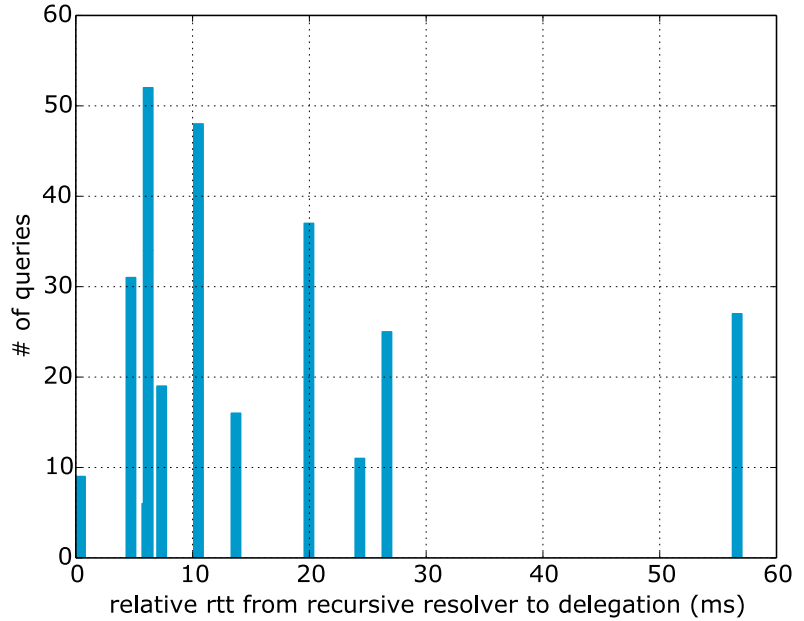


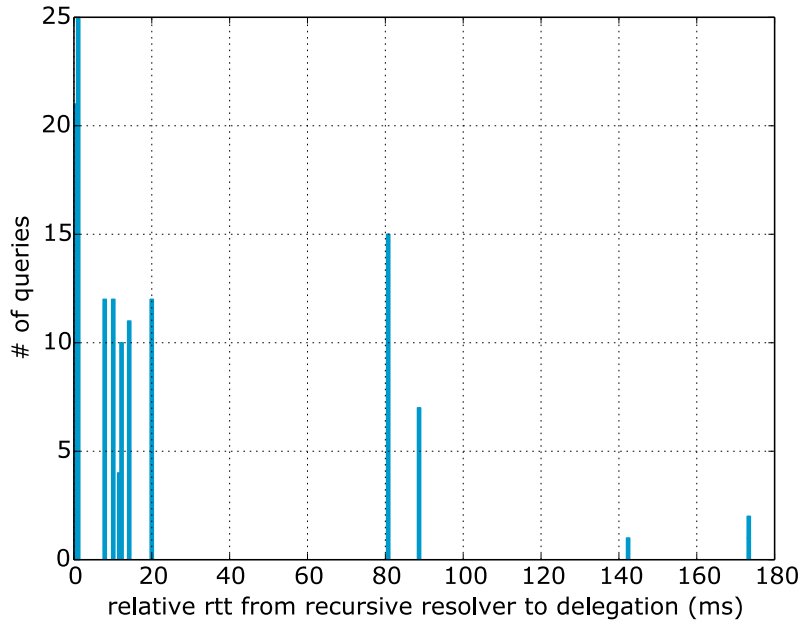

Calculate fraction of queries to fastest ½ delegations

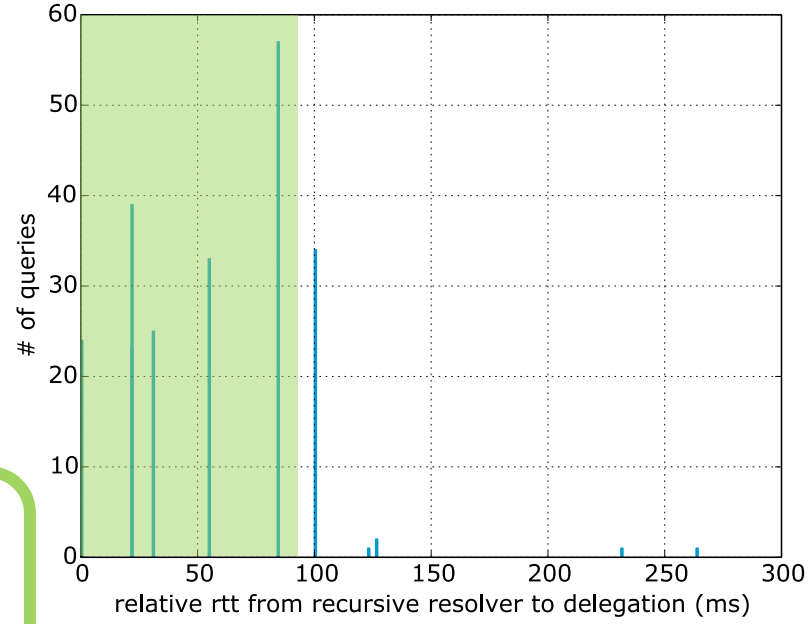


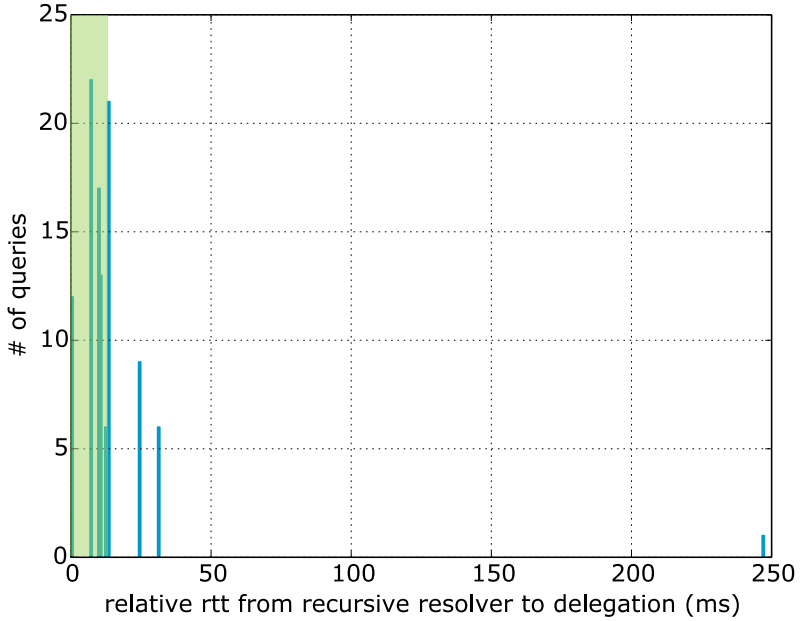


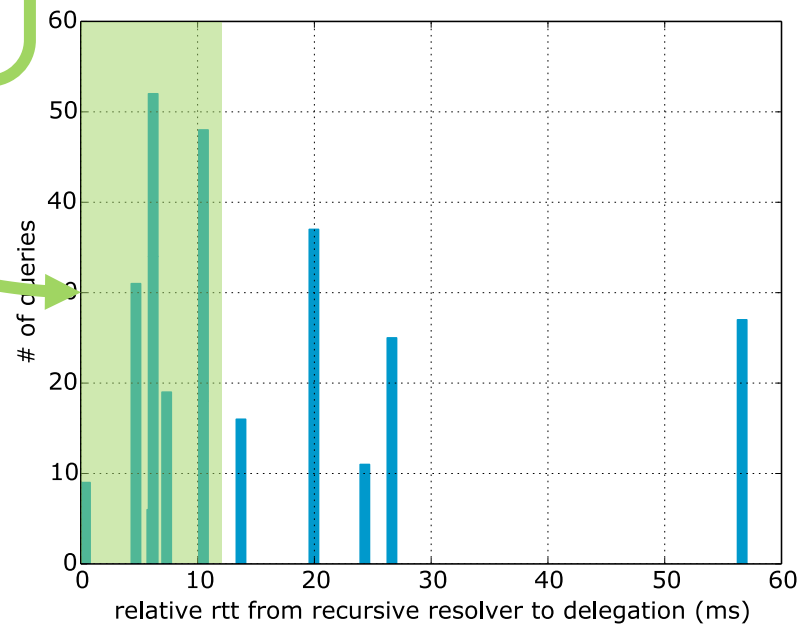


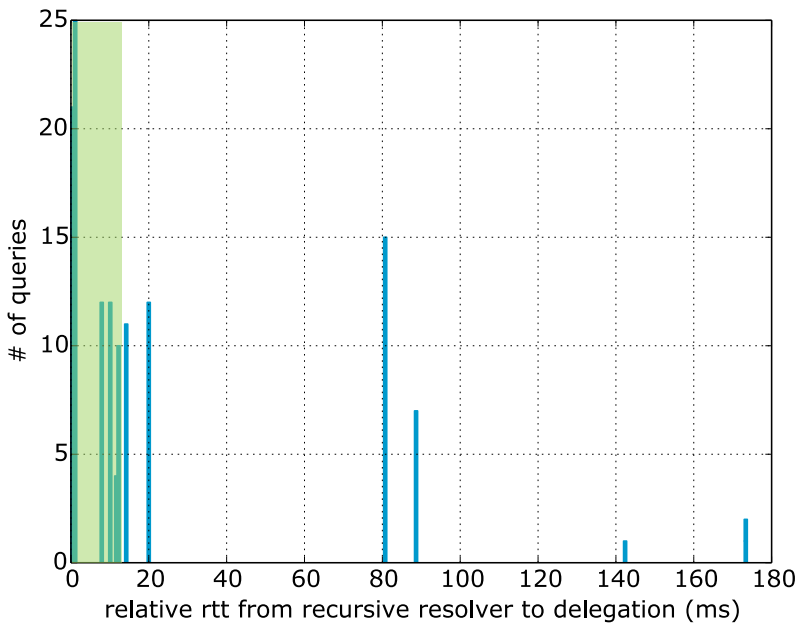


 

# Low latency preference

- Most resolvers show a preference for faster delegations

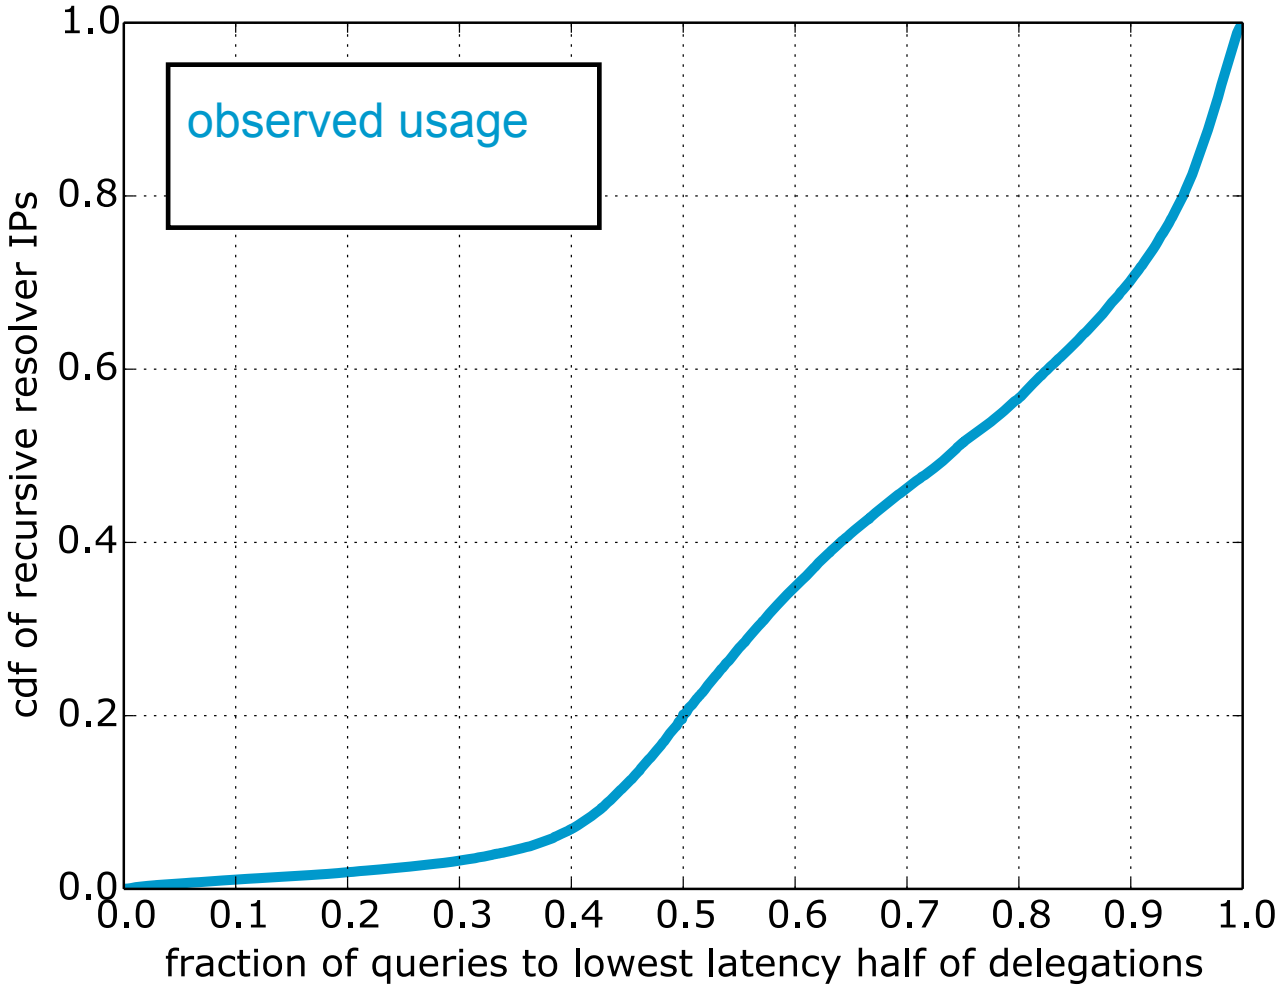


 

# Low latency preference

- Most resolvers show a preference for faster delegations

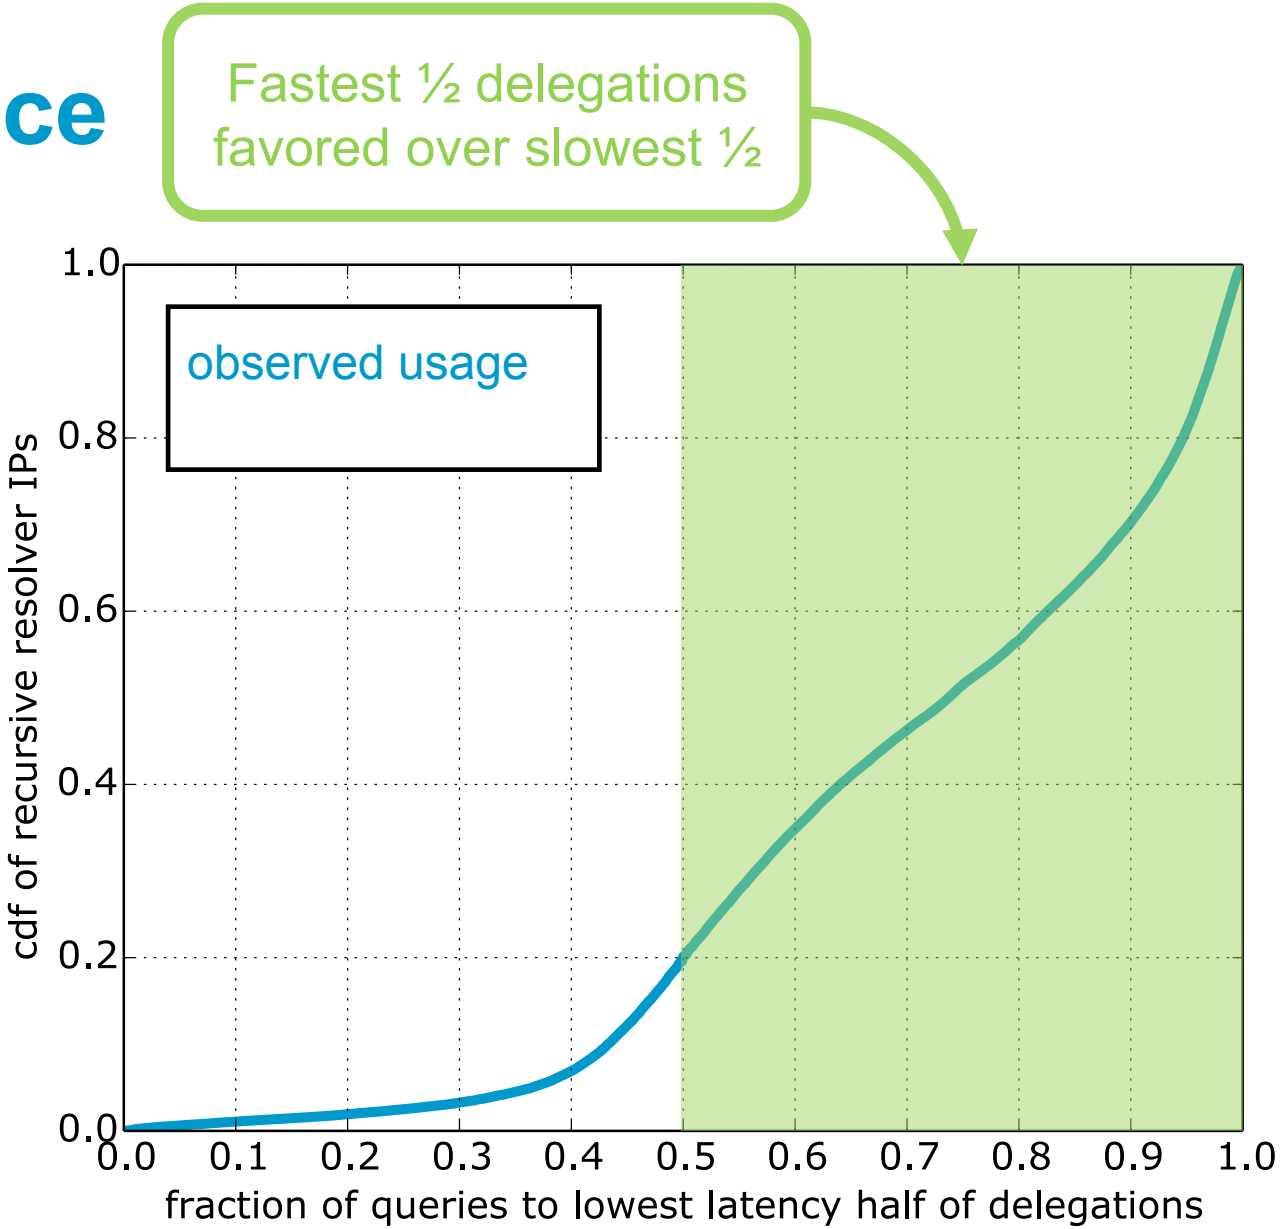


 

# Low latency preference

- Most resolvers show a preference for faster delegations

Fastest ½ delegations favored over slowest ½

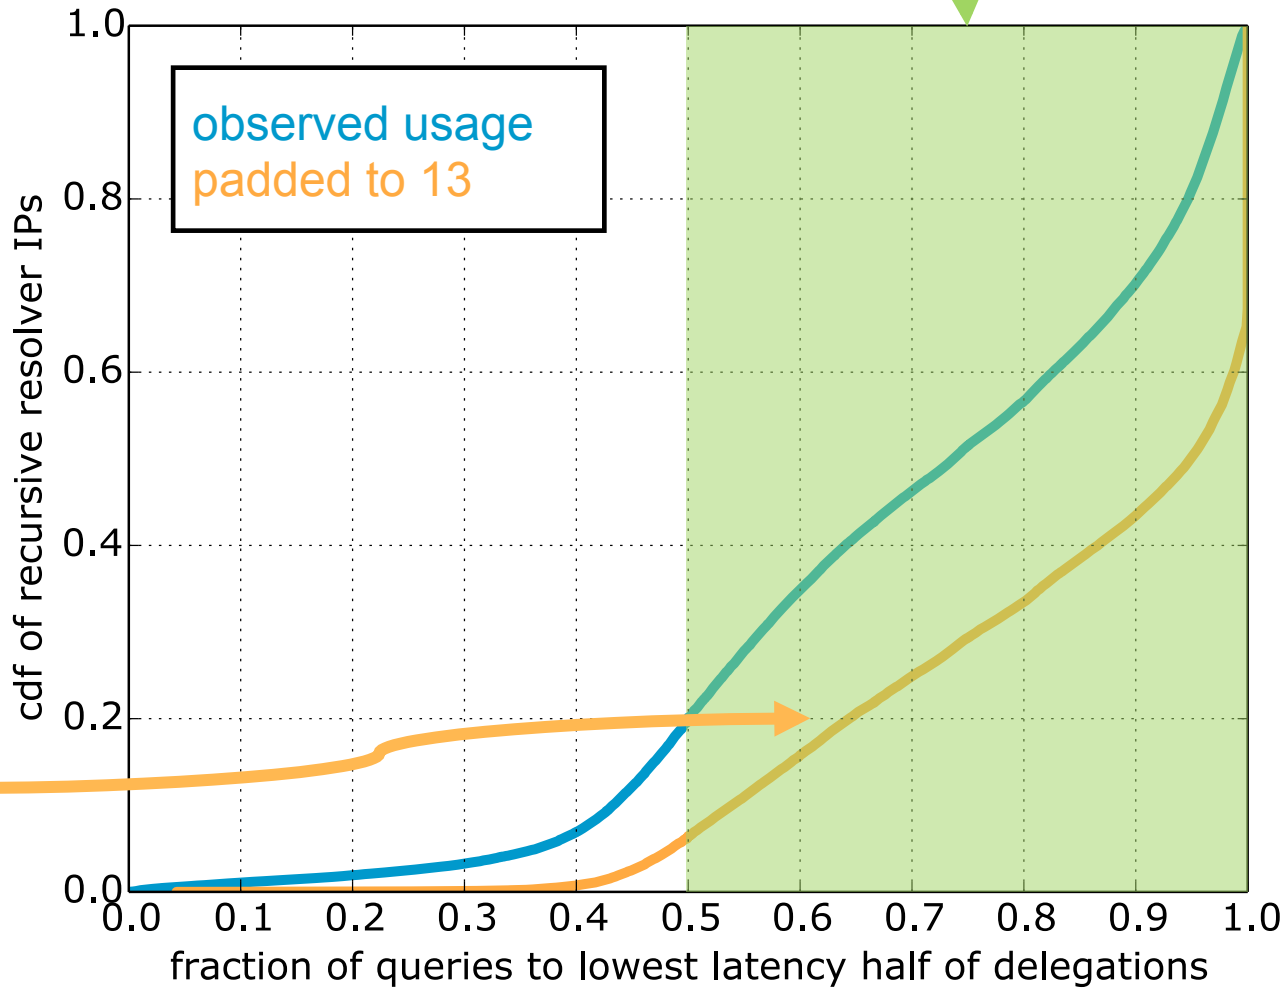


Assumes unused delegations were previously observed to be high latency

 

# Delay added to resolutions

- Not using the fastest delegation increases resolution time
- Upper bound on impact since some queries may be prefetching

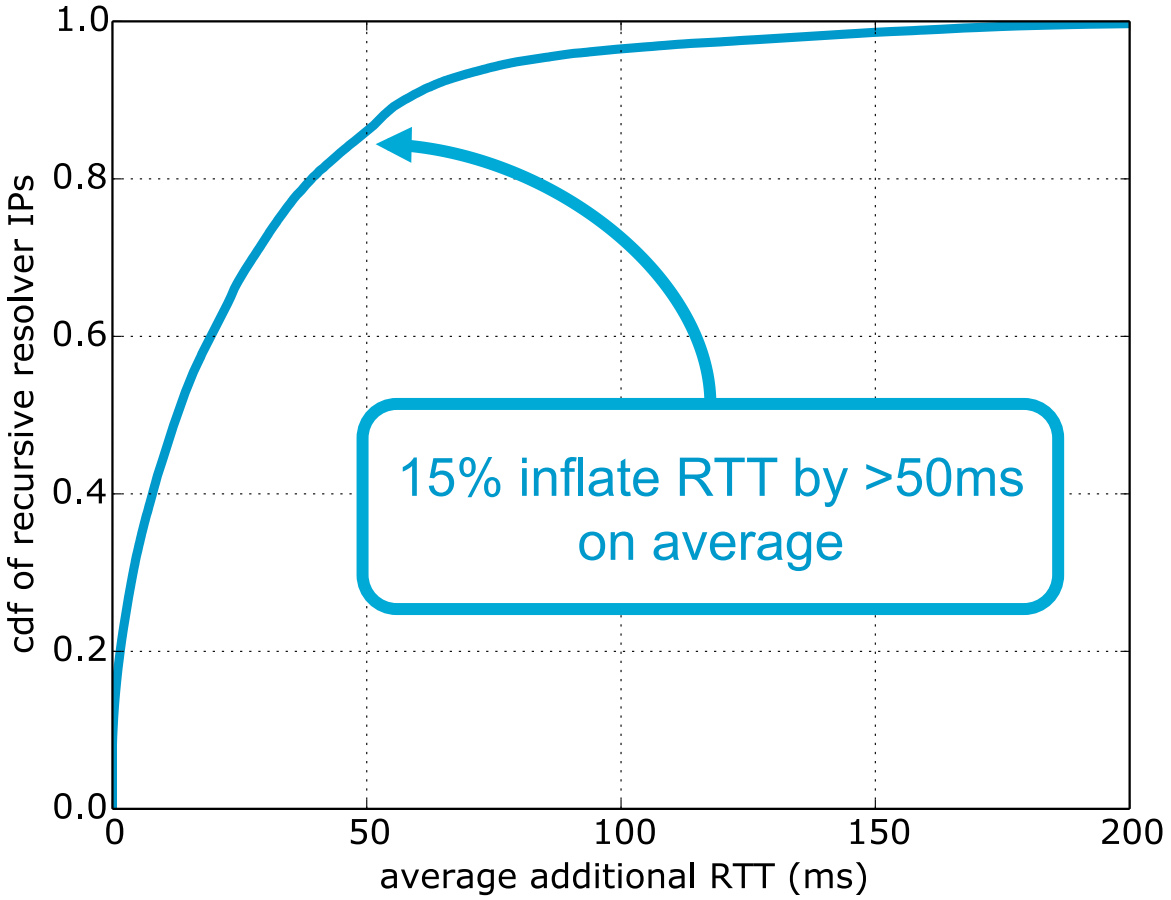

# Delay added to resolutions

- Not using the fastest delegation increases resolution time
- Upper bound on impact since some queries may be prefetching

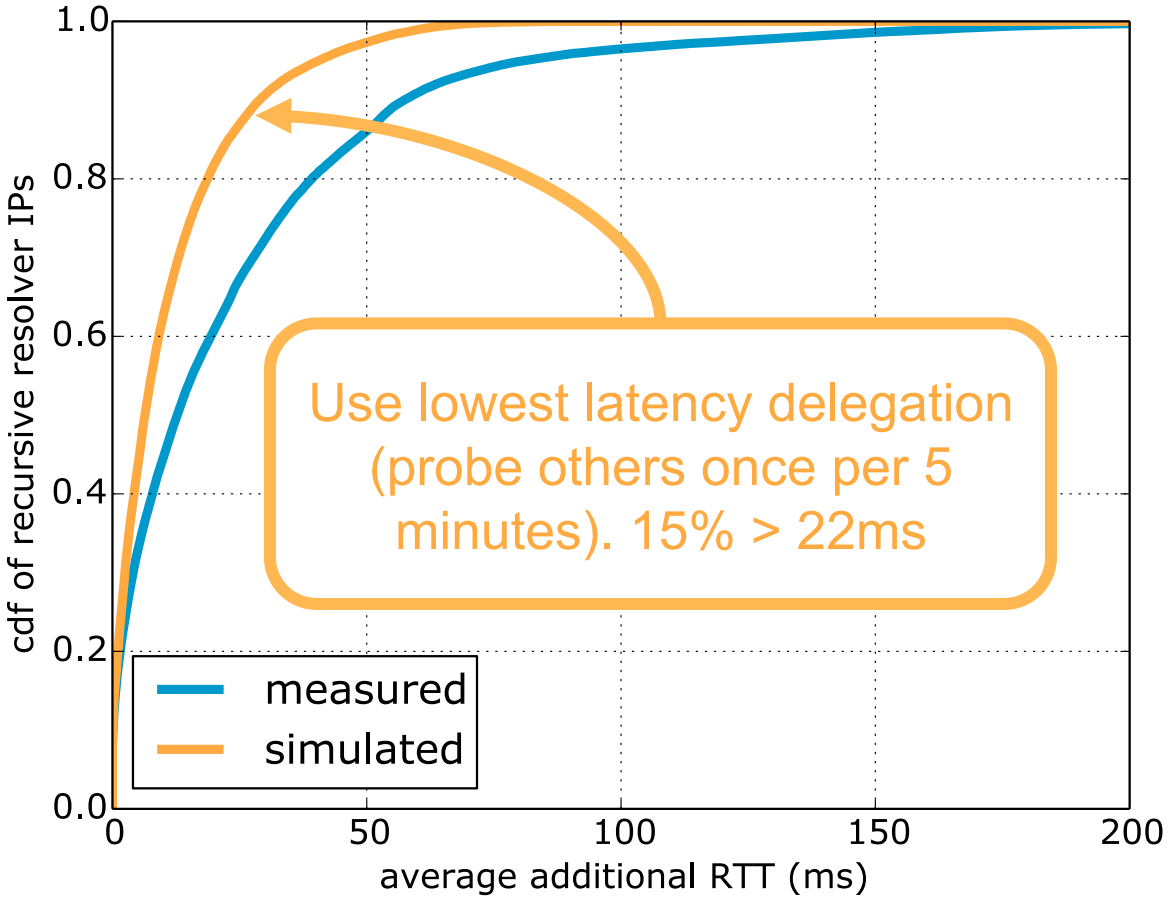

# Preventing Cache Poisoning

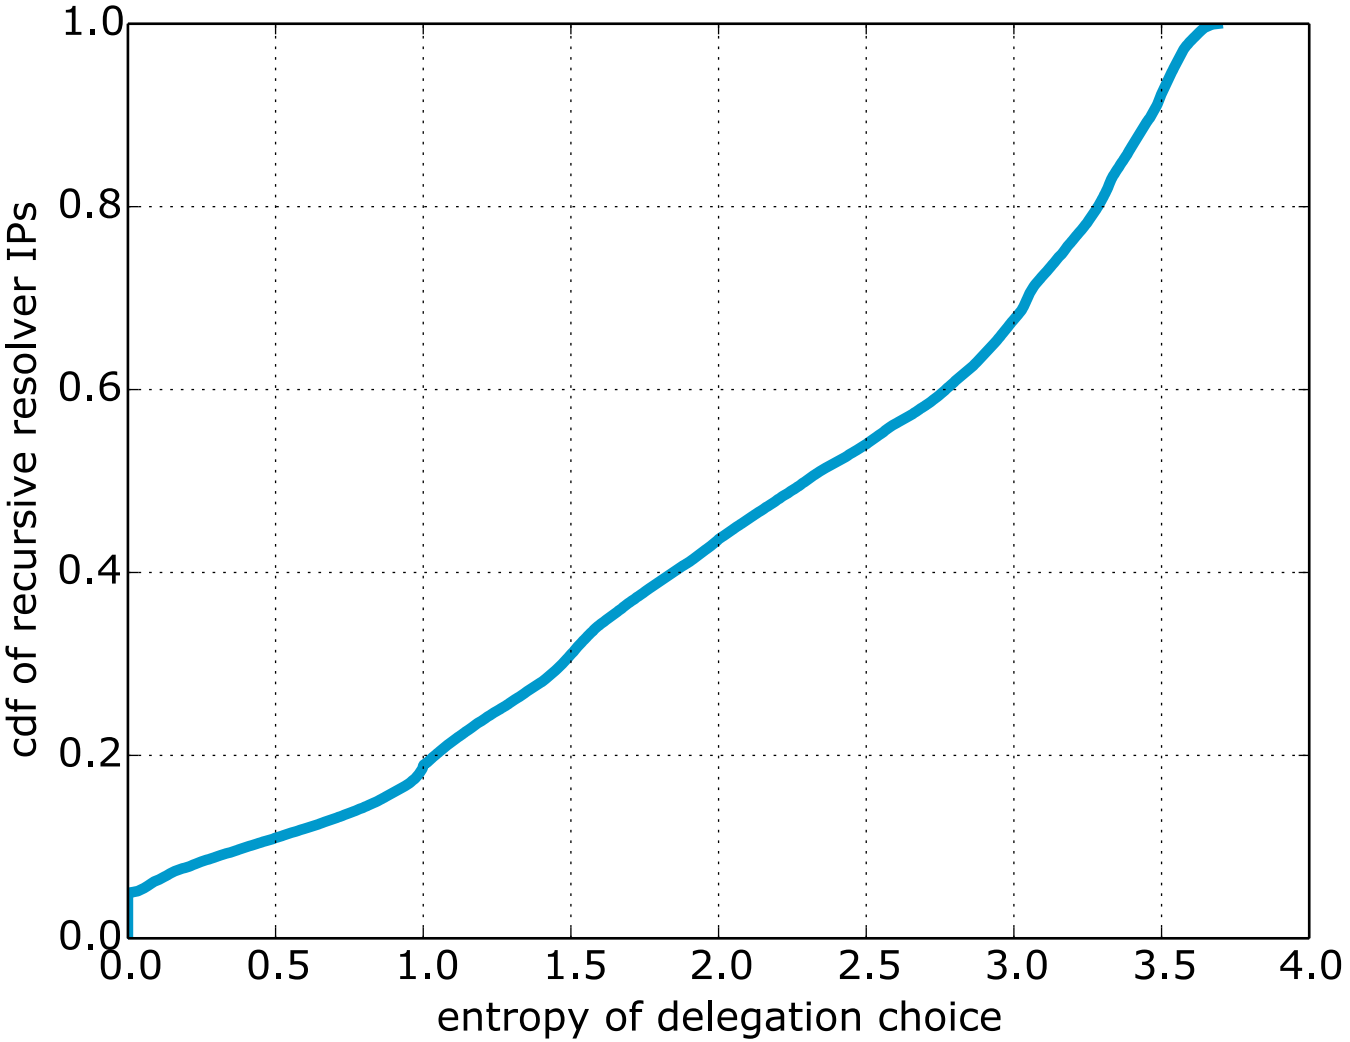


- Spreading queries across delegations adds entropy
- Randomizing source port + transaction ID provides ~31-bits of entropy
- Alternatives for adding entropy do not impact performance

 

# Preventing Cache Poisoning

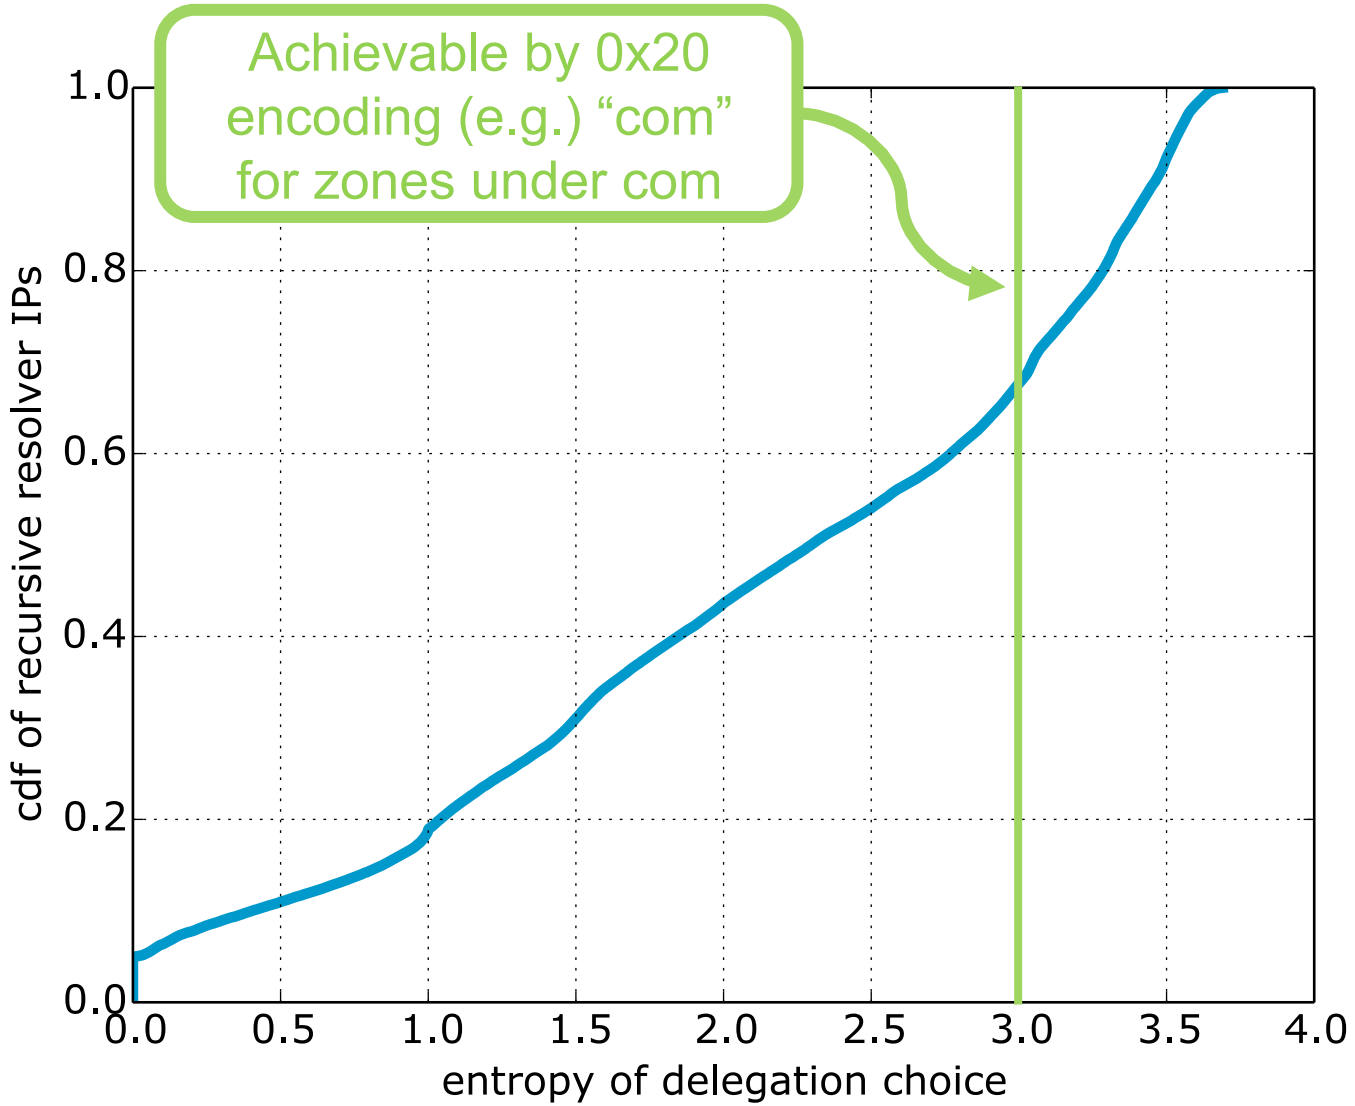


- Spreading queries across delegations adds entropy
- Randomizing source port + transaction ID provides ~31-bits of entropy
- Alternatives for adding entropy do not impact performance

 

# Preventing Cache Poisoning

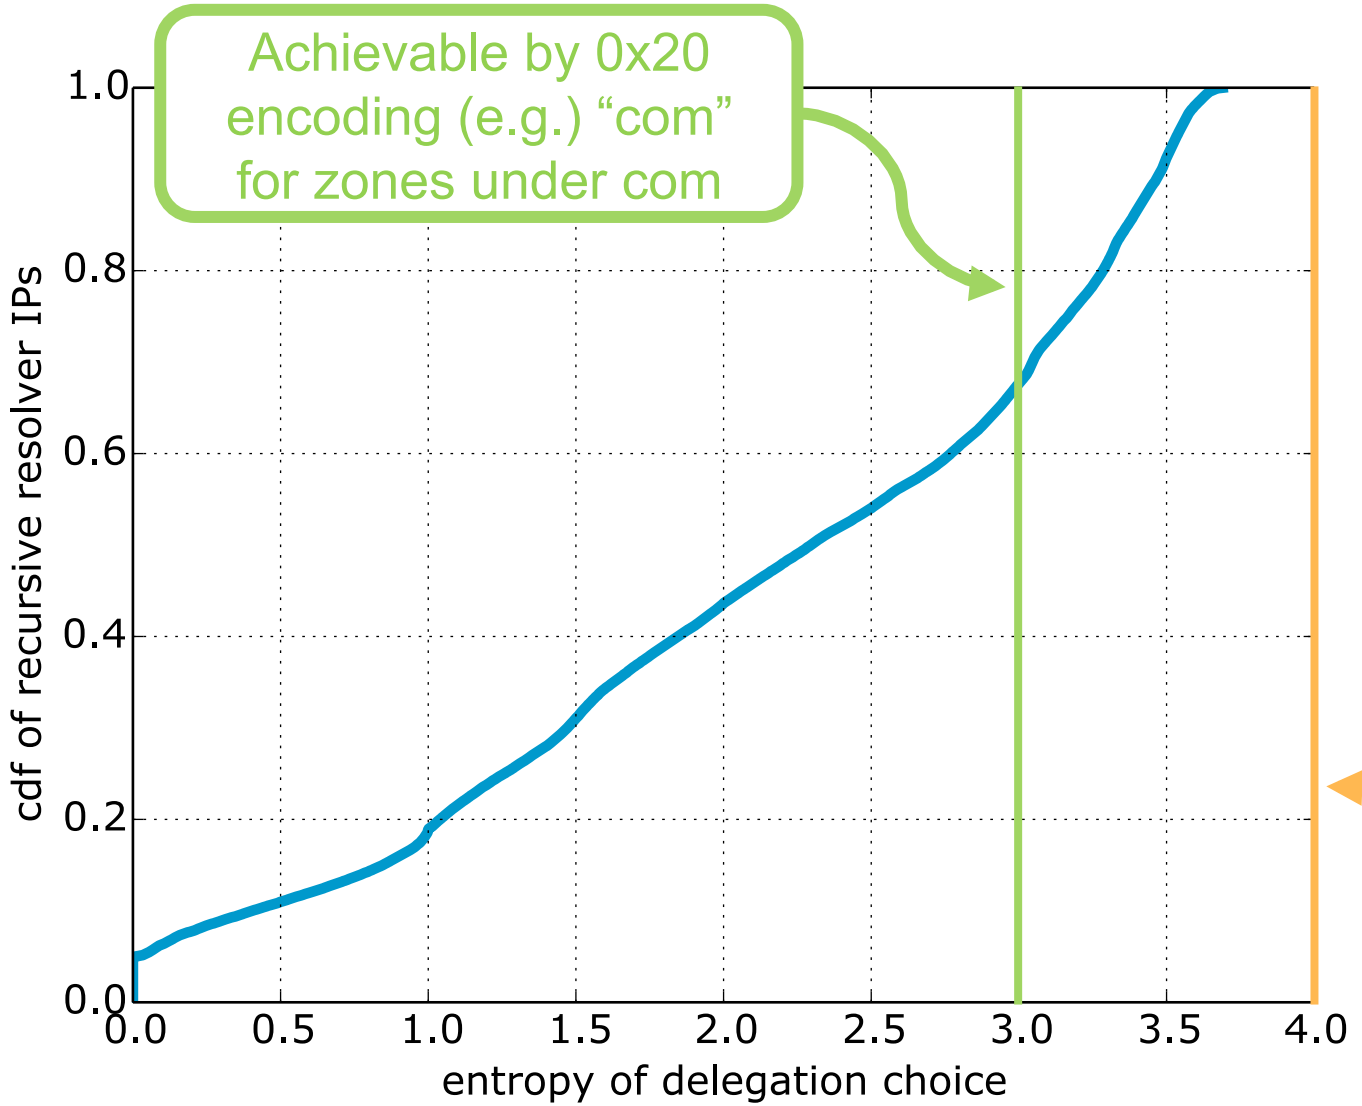


- Spreading queries across delegations adds entropy
- Randomizing source port + transaction ID provides ~31-bits of entropy
- Alternatives for adding entropy do not impact performance

Achievable by randomizing among 16 source IPs (e.g., resolver “pool”)

 

# What about the rest?

- Limited data makes it harder to identify behavior
- Do they behave the same?

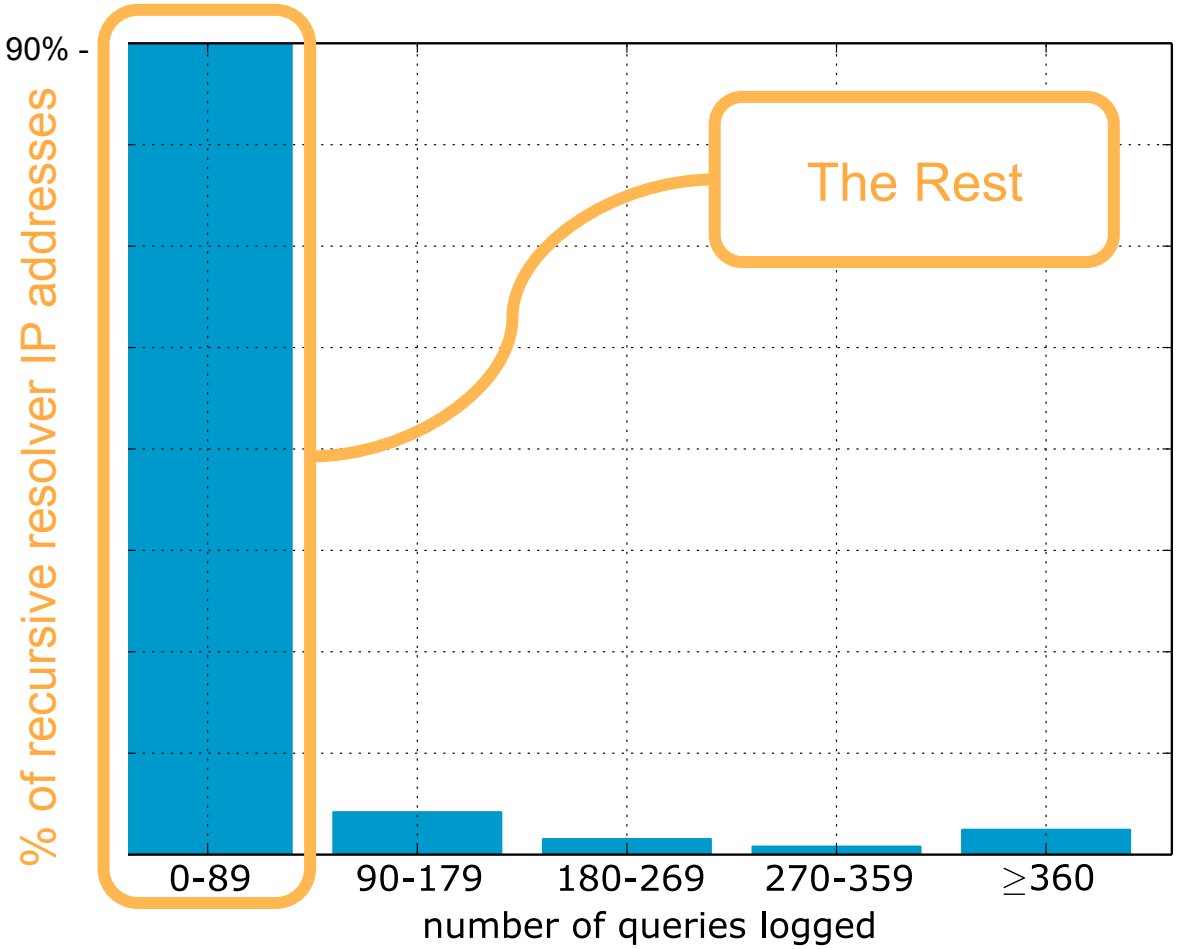


 

# Is there anything special about the 90 queries threshold?

- Threshold of 90 queries is somewhat arbitrary
- If resolvers below the threshold behave similar to those above, then observations can be generalized

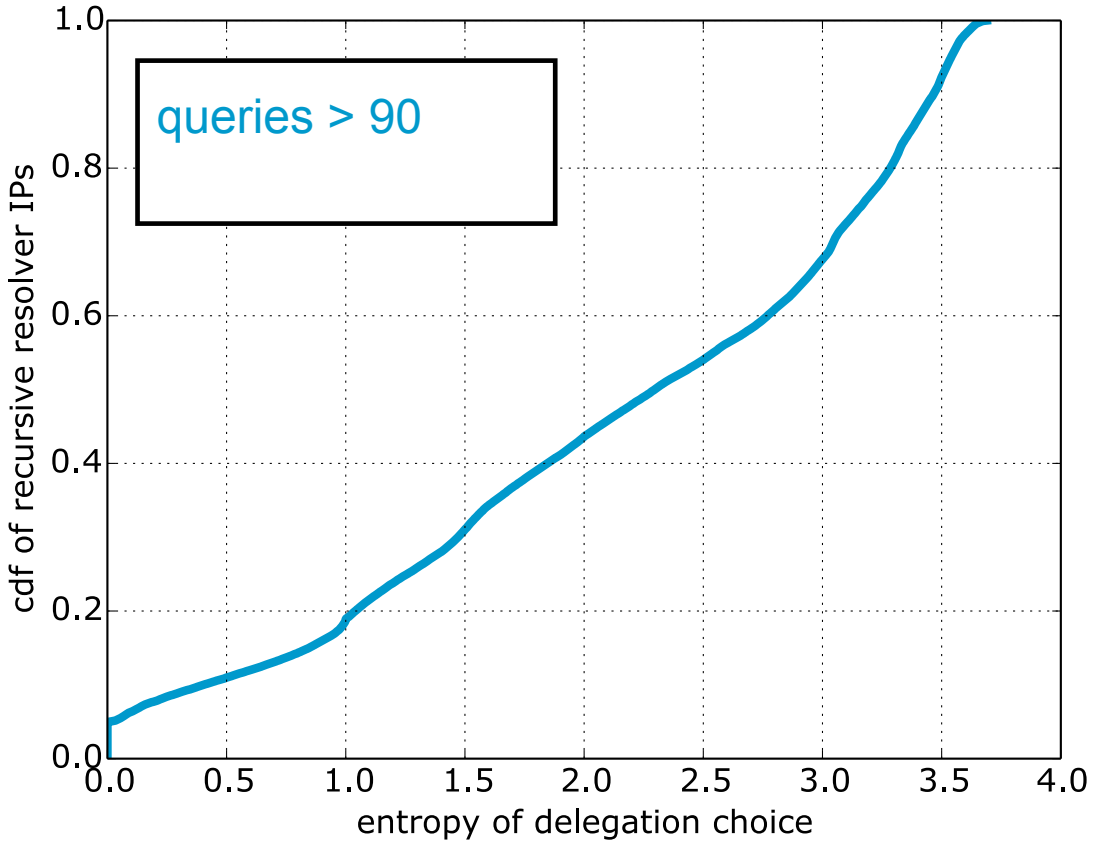


 

# Is there anything special about the 90 queries threshold?

- Threshold of 90 queries is somewhat arbitrary
- If resolvers below the threshold behave similar to those above, then observations can be generalized

A smaller disjoint set of source IP addresses shows similar distribution in delegation choice

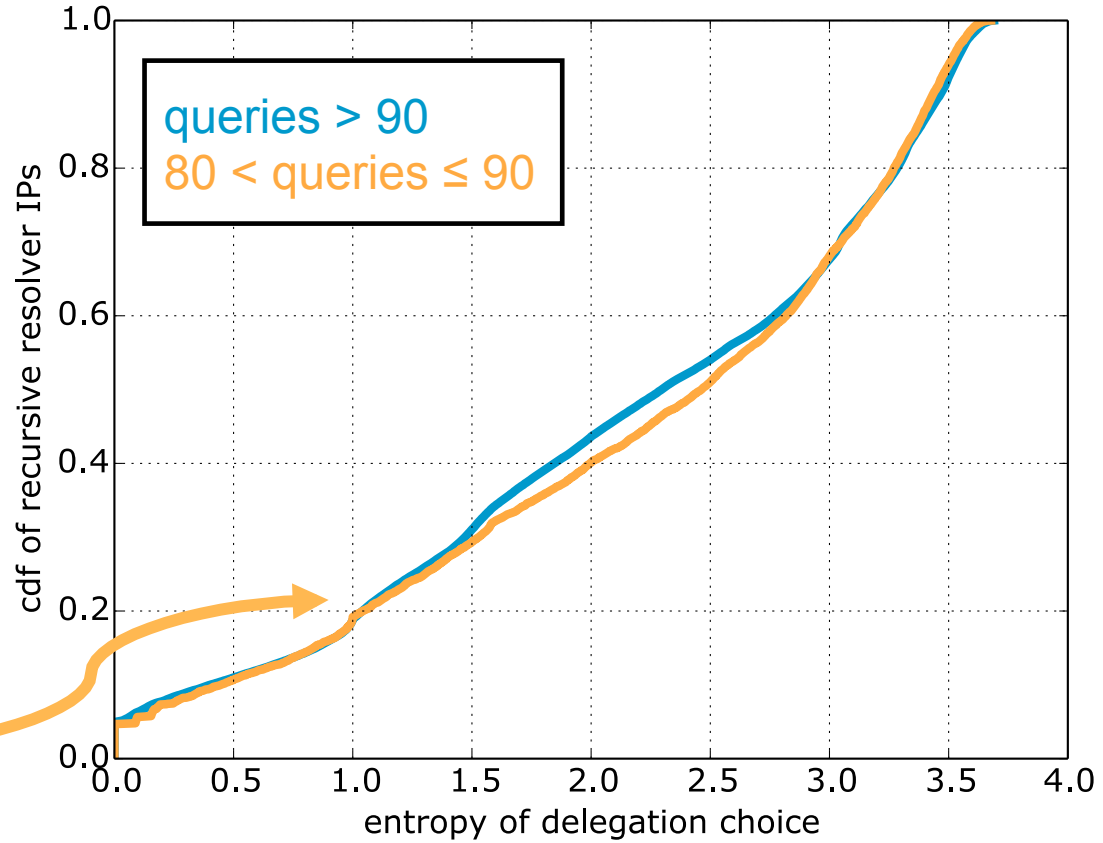


 

# Is there anything special about the 90 queries threshold?

- Preference for lower latency delegations also looks similar

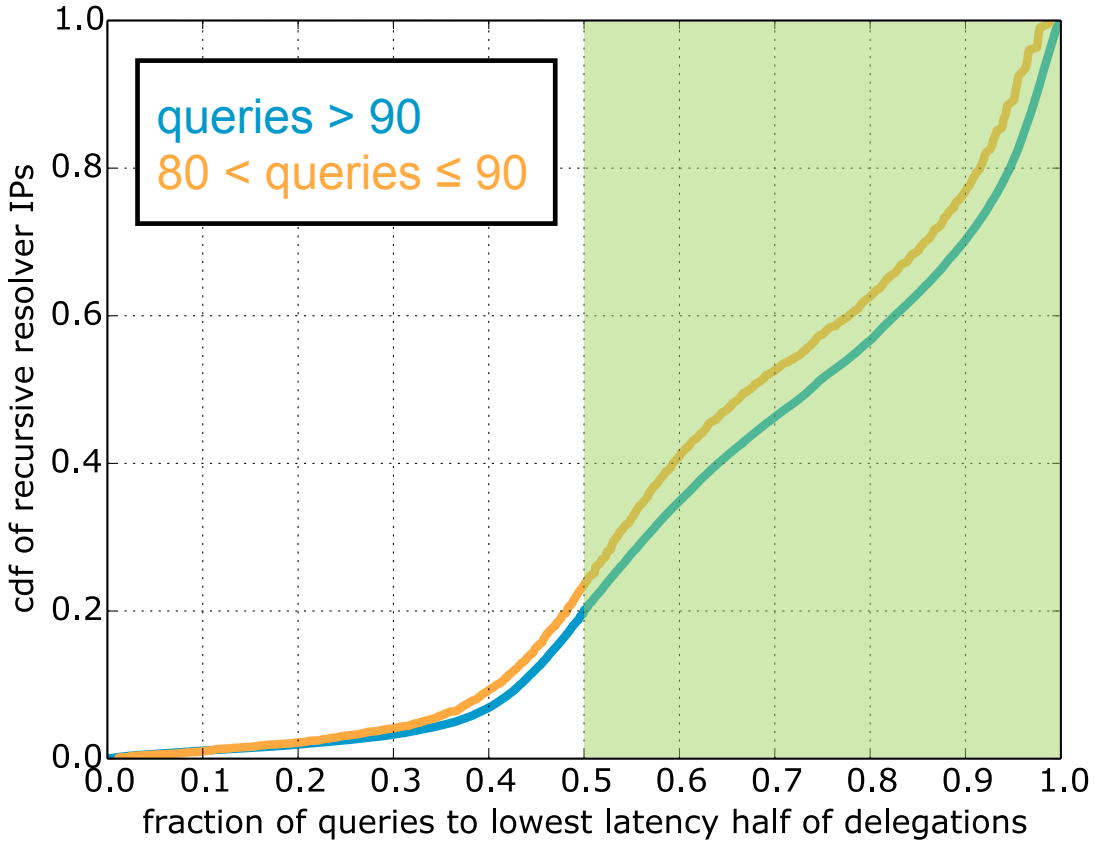


 

# Query rates too low

- Low query rates mean many resolvers will not use the low latency delegations despite algorithms that attempt to identify them

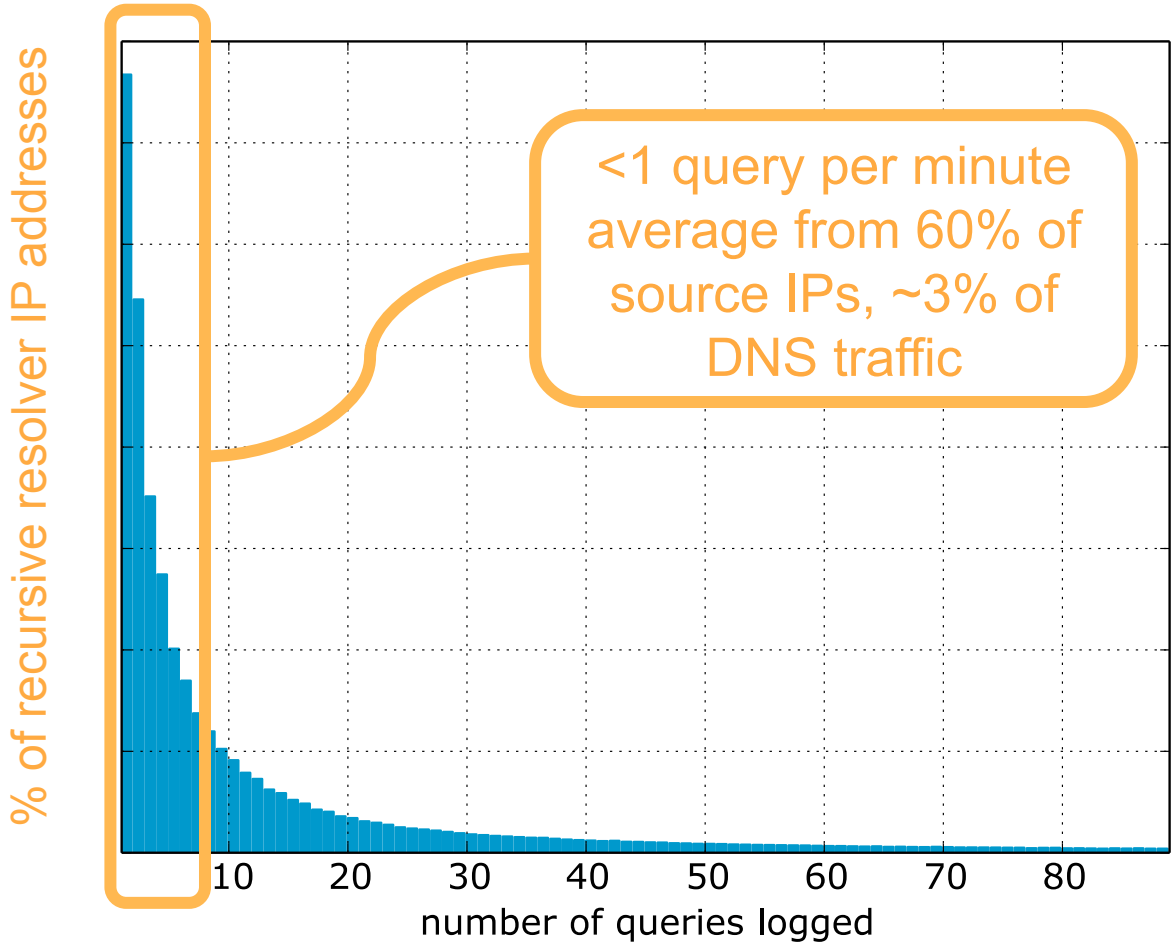


 

# Summary

- <6.7% of resolvers query delegations uniformly
- ~10% of resolvers send nearly all queries to a single delegation
  - Likely the lowest latency delegation
- Remainder attempt to prefer low latency delegations
  - Higher average resolution time over alternatives
  - 60% of resolvers (~3% of DNS traffic) have querying rates low enough that algorithms likely unsuccessful

 

# Suggested improvements

- Authoritative nameserver deployments should strive to offer low latency for all delegations
  - Agrees with findings in other research
- Recursive resolver software can reduce resolution time by using the fastest delegation for the vast majority of DNS queries
  - Probe other delegations rarely
    - Open question: how frequently is good enough?
  - Use other methods for adding entropy to prevent cache poisoning attacks

 

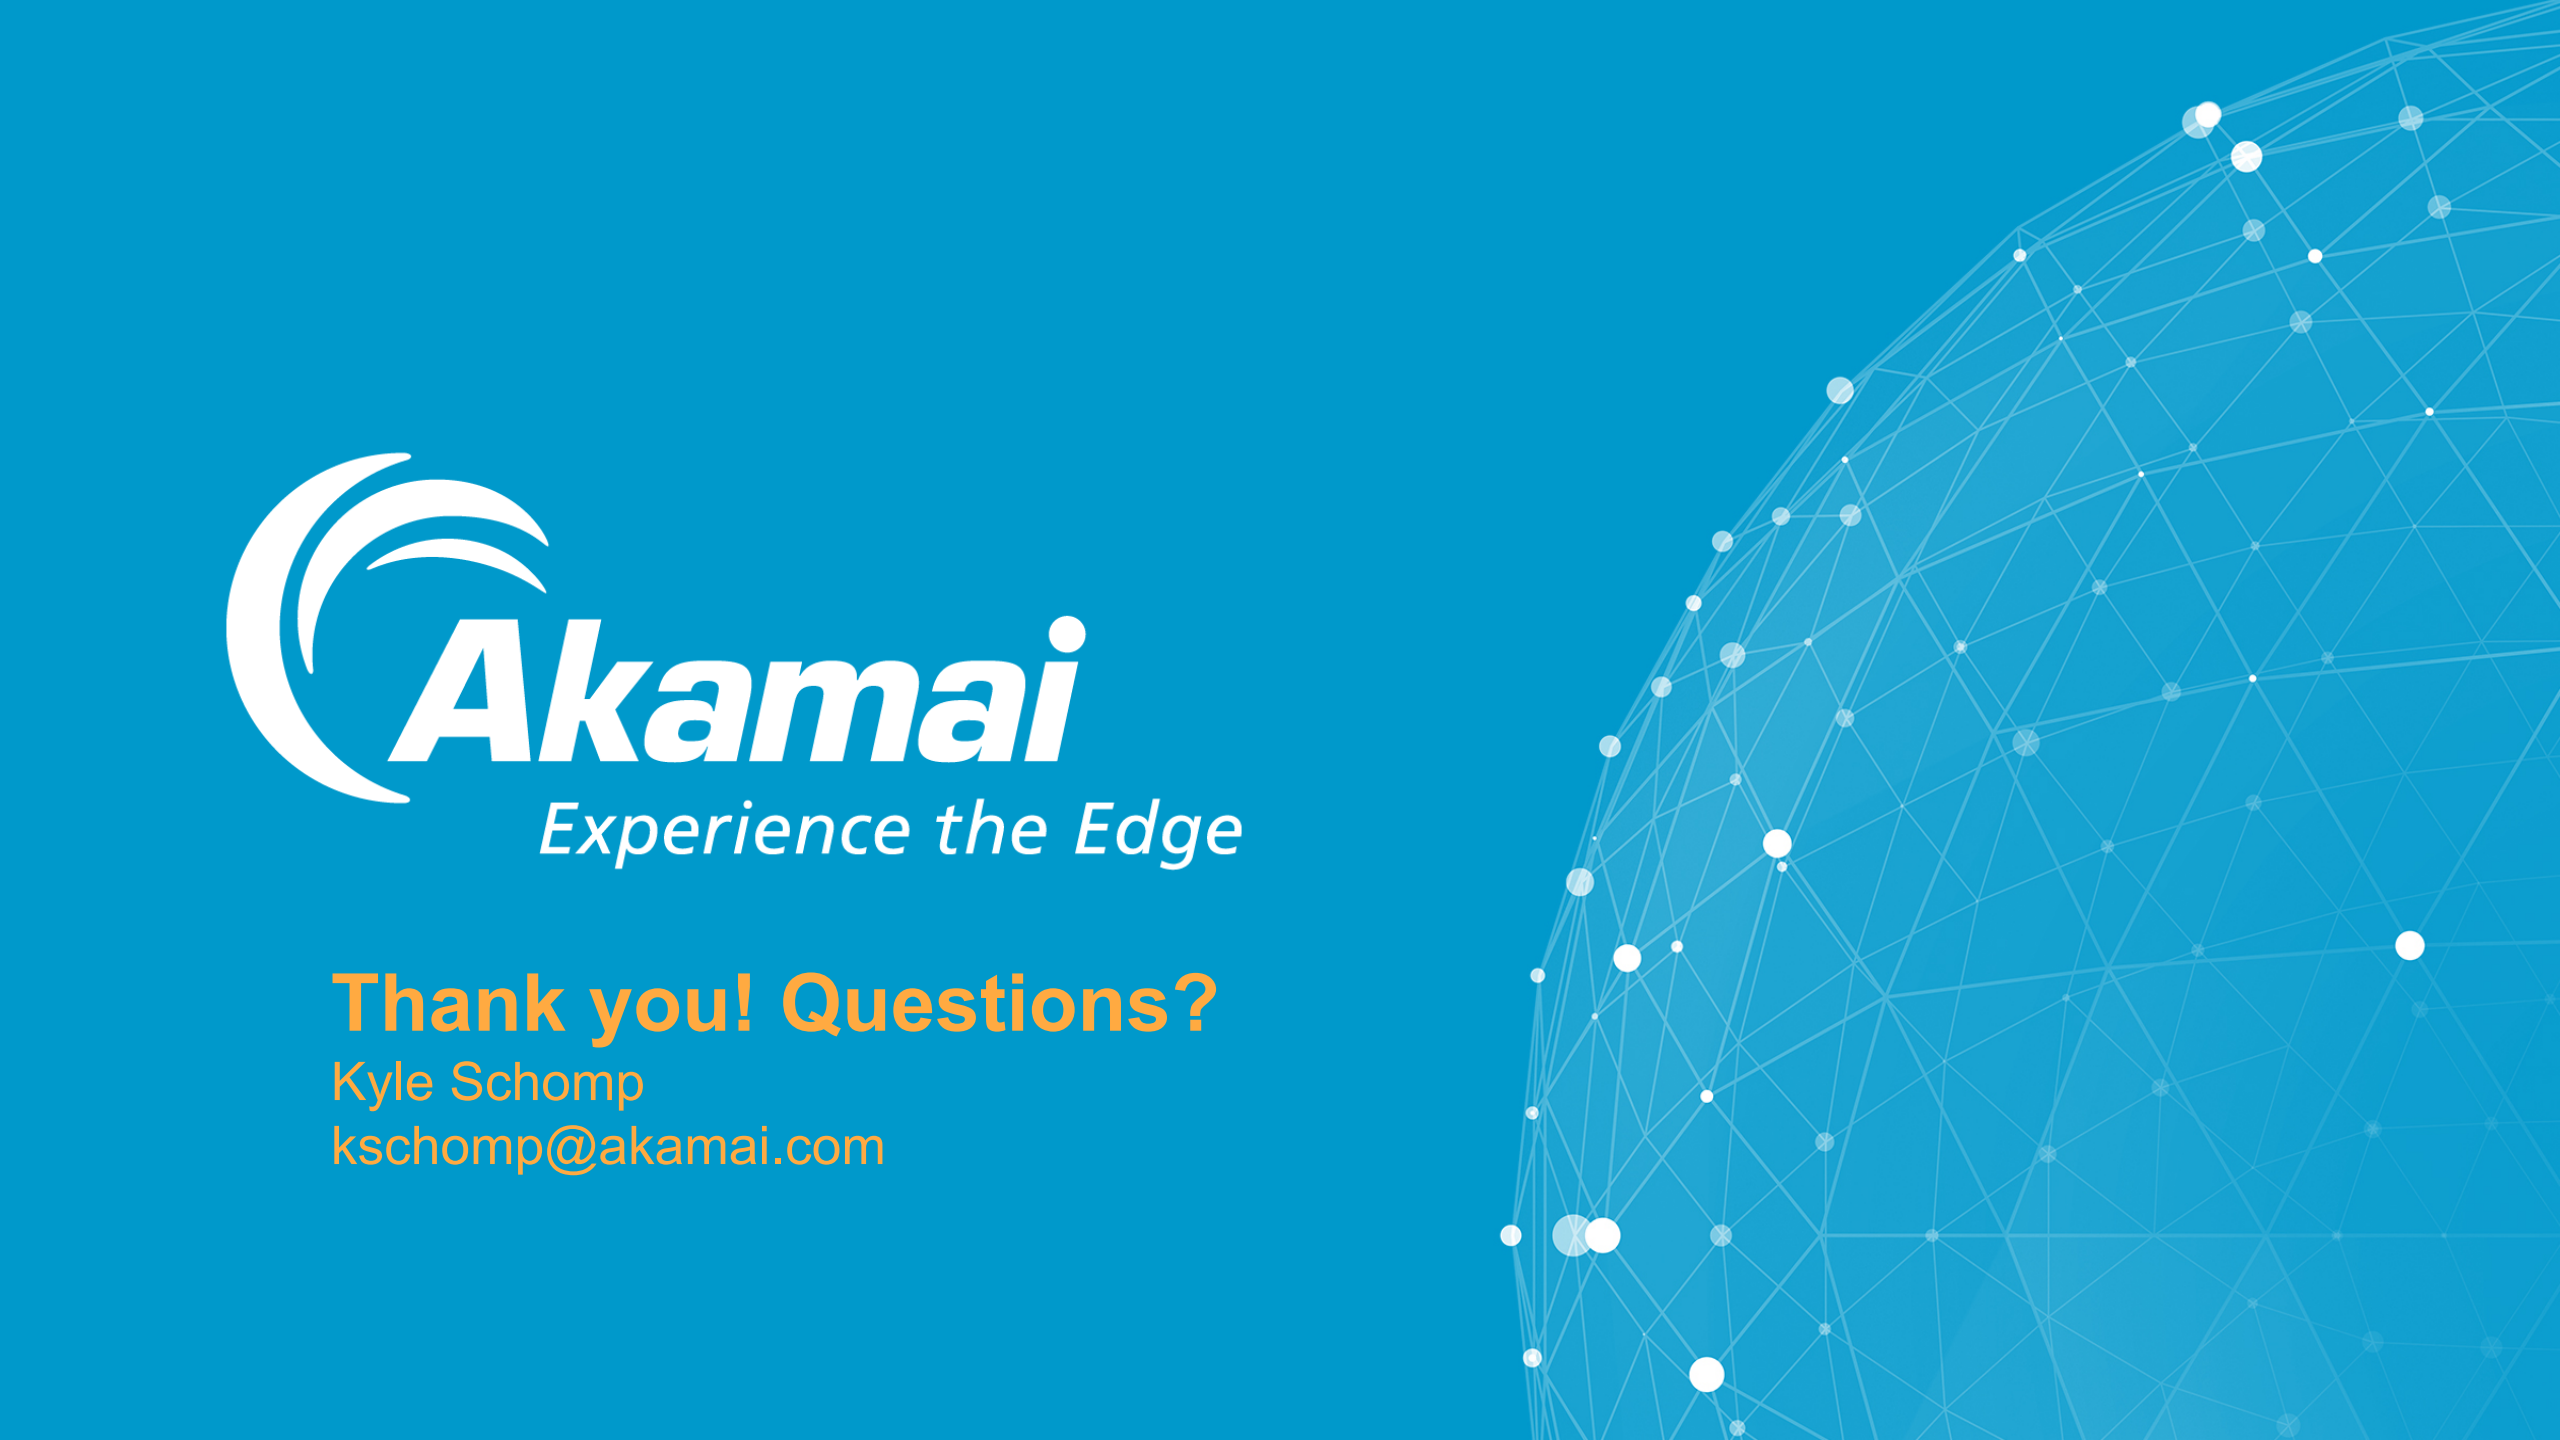


# Thank you! Questions?

Kyle Schomp
kschomp@akamai.com

# EXTRA


Recursive algorithm changes unlikely to help here

<10 queries in 10 min
(using 1 day of logs)

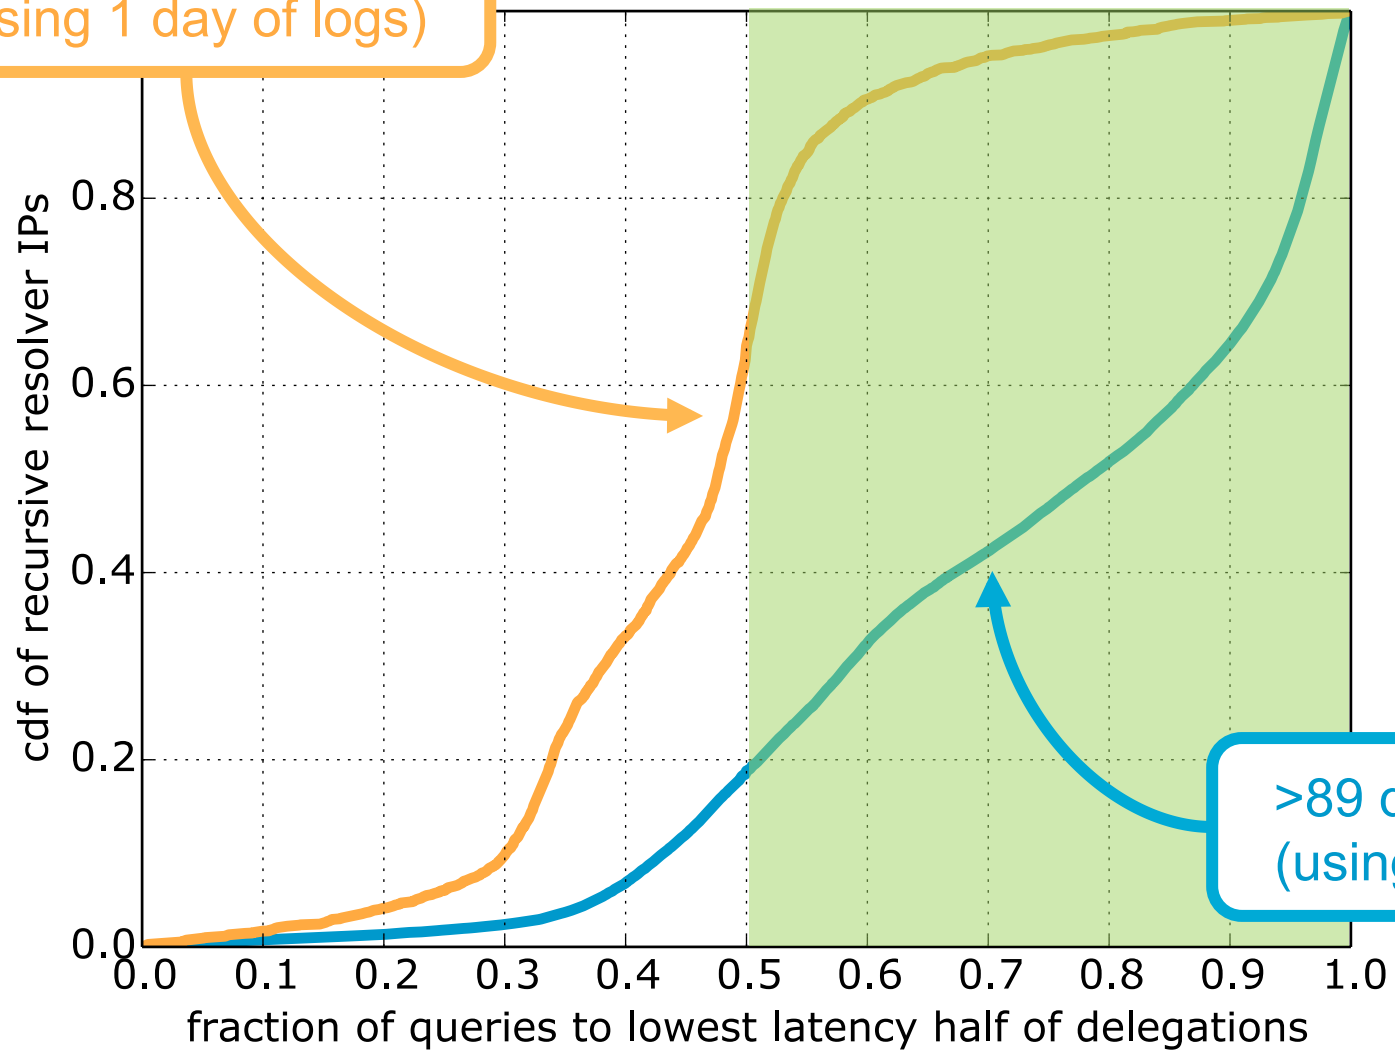


>89 queries in 10 min
(using 10 min of logs)

 