Cache Effect of Shared DNS Resolver

Kazunori Fujiwara^, Akira Sato”, Kenichi Yoshida”

”University of Tsukuba

^Japan Registry Services Co., Ltd (JPRS)

This material comes from K. Fujiwara, A. Sato”, K. Yoshida, "Cache Effect of Shared DNS Resolver", Proceedings of COMPSAC 2017 (IEEE Computer Society Signature Conference on Computers, Software and Applications)
https://doi.org/10.1109/COMPSAC.2017.77

Outline

- Related works
- Issue and evaluation
- DNS Traffic data at Univ. of Tsukuba
- Analysis results
  - Number of DNS queries to authoritative servers
  - Amplification Ratio
  - Frequent Queries and Active Clients
- Conclusion

Related works

- In 2002, Jaeyeon et al. reported on DNS performance and the effectiveness of caching
- In 2012, Fujiwara et al. reported "DNS traffic analysis -- Issues of IPv6 and CDN --"
  - It analyzed a shared resolver in an university.
  - Average cache hit rate was 75.1%.
  - A client query generated 0.00079 Root queries, 0.025 TLD queries and 0.28 other DNS server queries.
- In 2014, Schomp et al. reported "DNS Resolvers Considered Harmful" and proposed
  - removal of shared DNS resolvers
  - use of a full-service resolver at the end clients instead

Issue and evaluation

- Issue
  - A paper proposed removal of shared DNS resolvers
  - Removal of shared cache increases DNS traffic
  - Need to evaluate effect of local/shared cache
- To Estimate Cache Effects, Minimal Resolver Behavior simulation with real queries from clients
  - captured at one of busy resolvers at University of Tsukuba

Data Collection system and captured data

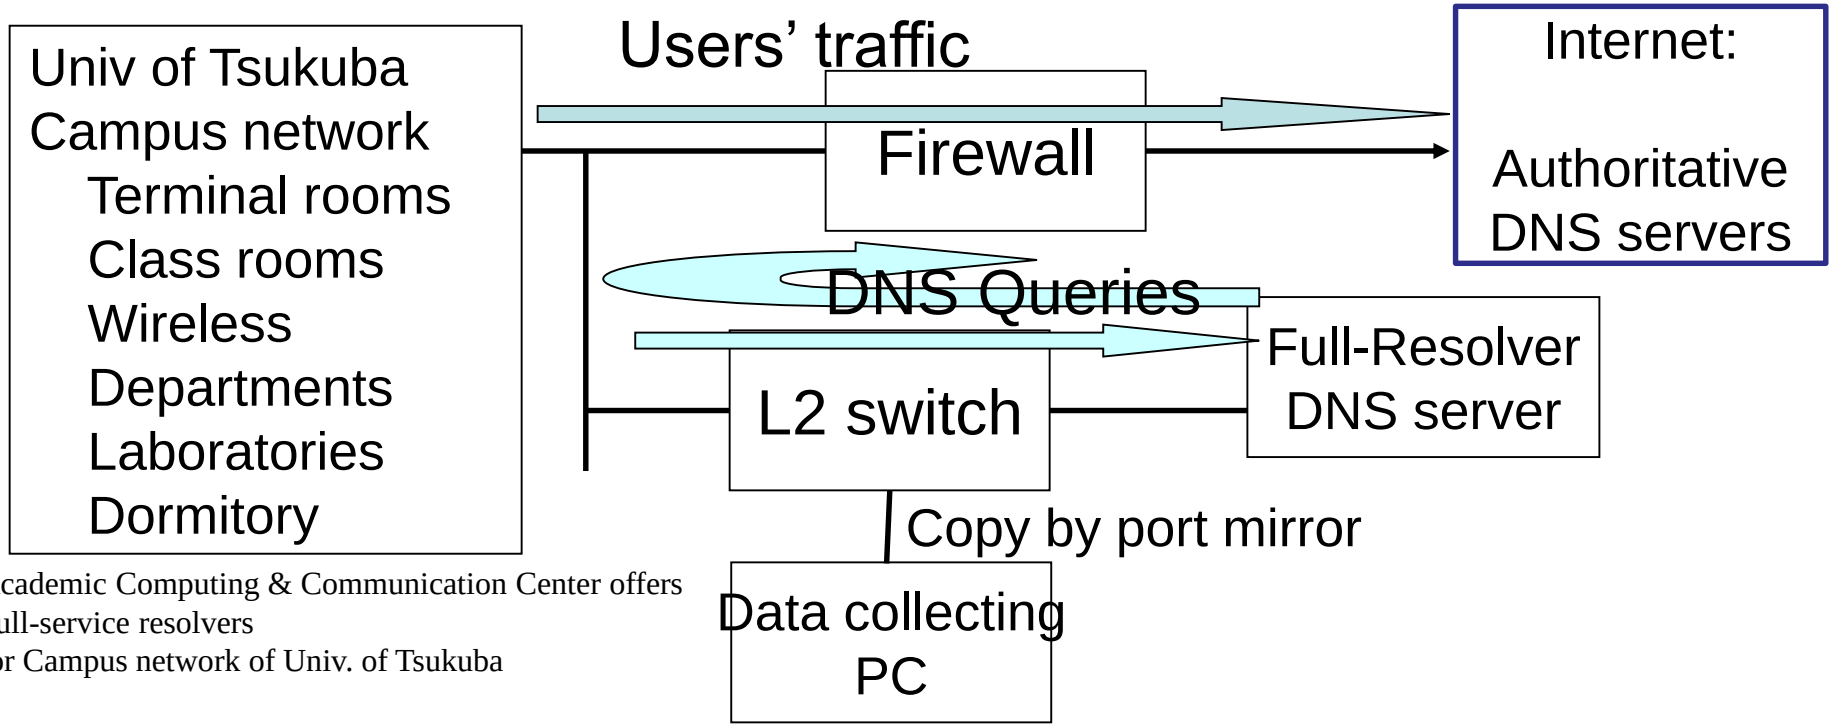

Academic Computing & Communication Center offers
Full-service resolvers
for Campus network of Univ. of Tsukuba

Capturing Date	2016.Dec.01 0:00 ～ 24:00
Capturing Point	Full Resolver of Campus Network
Number of Clients	6,475
Number of QNAMEs	290,862
Average Number of queries / minute	14,650 (244 queries/sec)

Resolver simulation details

- Assumption: Authoritative servers are 3 layered
  - Root, TLD, Organization (second level)
    - Except in jp domain name
  - Each query counts up number of queries counters for Root, TLD, Organization
- Three cache scenarios
  - Without Cache
    - All queries from clients generates queries to Root, TLD, Organization
  - With Local Cache: Cache simulation in each address
  - With Shared Resolver: cache simulation with all queries

Limitations of the simulation

- Ignored: CNAME chain following
- Ignored: Out-of-bailiwick name server name resolution
- Ignored: Domain name existence
- TTL at organization (second, third) level is 300
- Short term evaluation: 1 day

→ This is preliminary result

Result: Number of queries to auth.

Assumption: TTL 300 (5min).

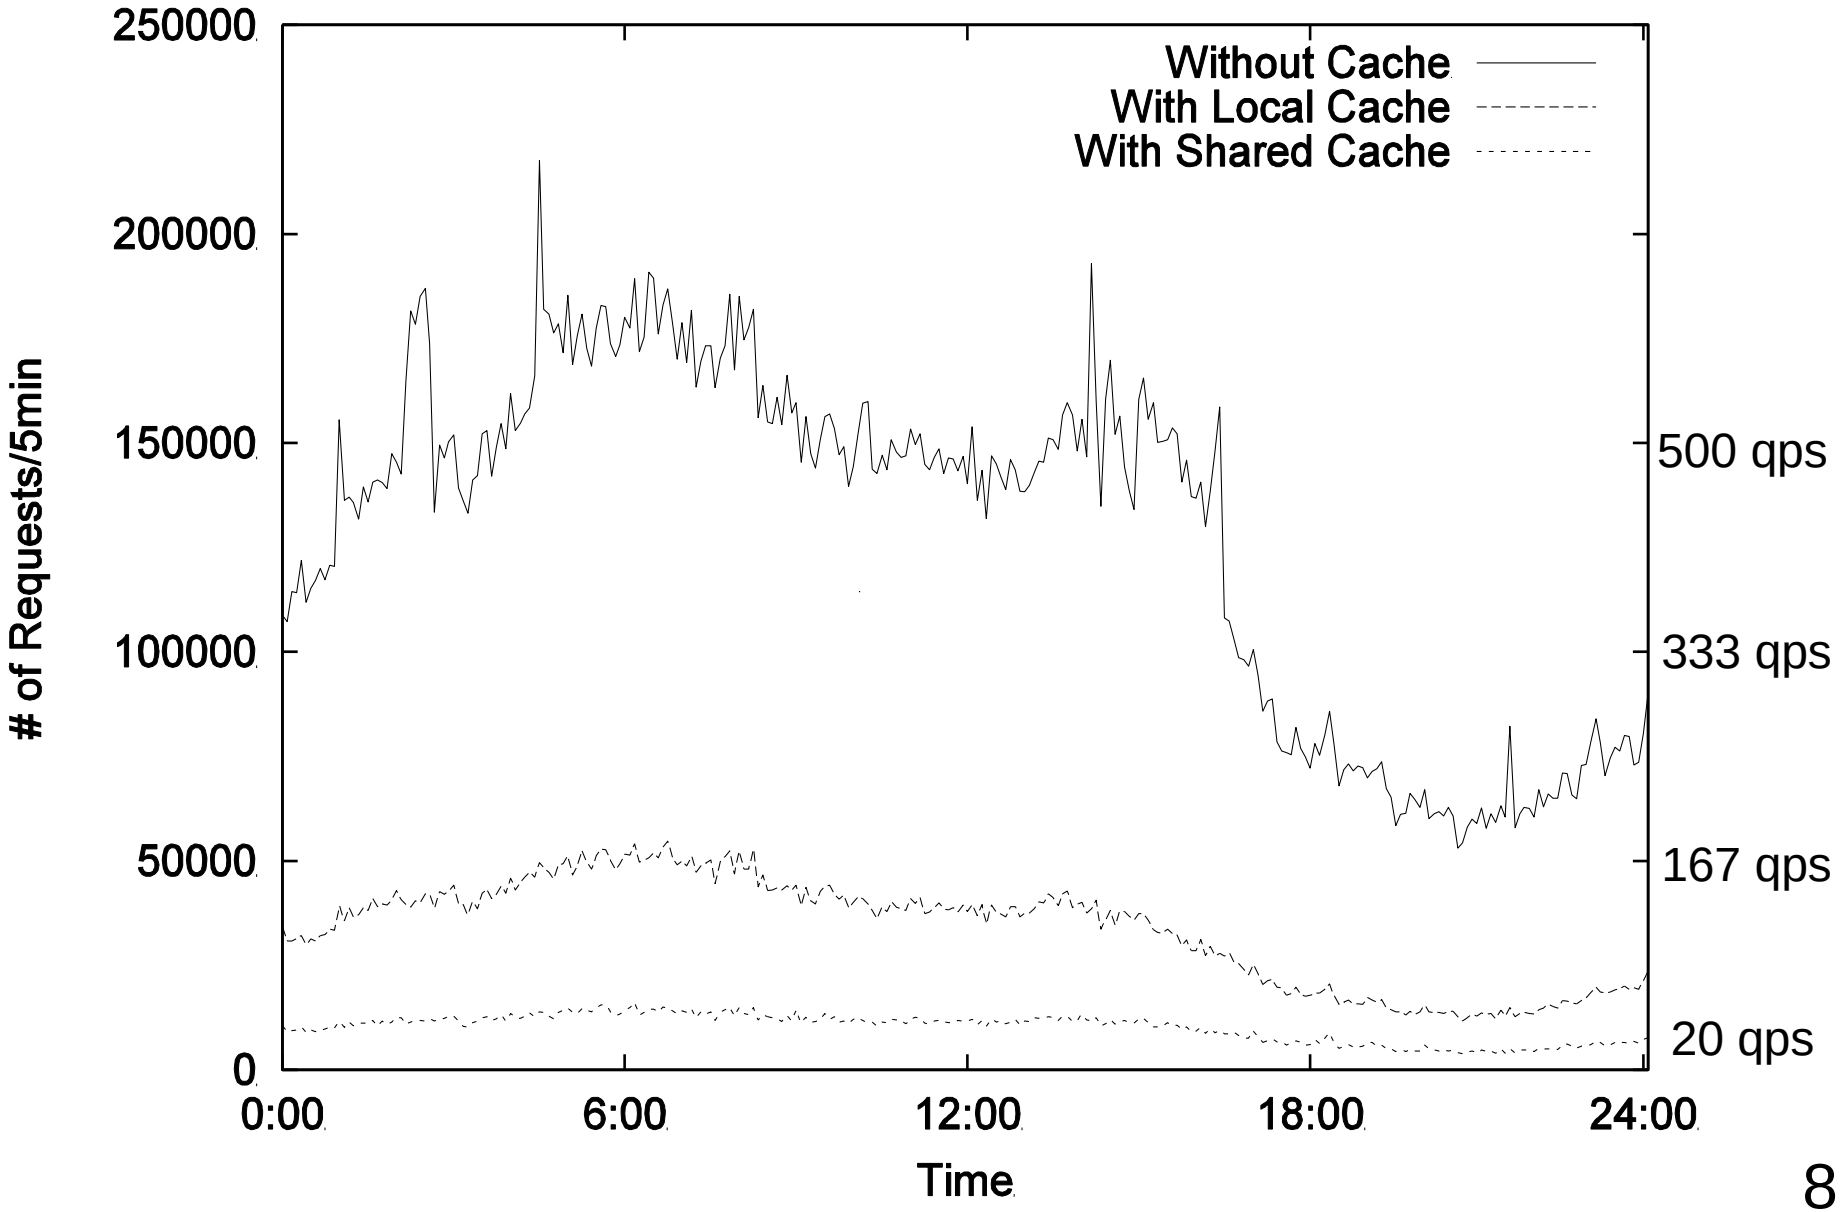

Result: Amplification Ratio

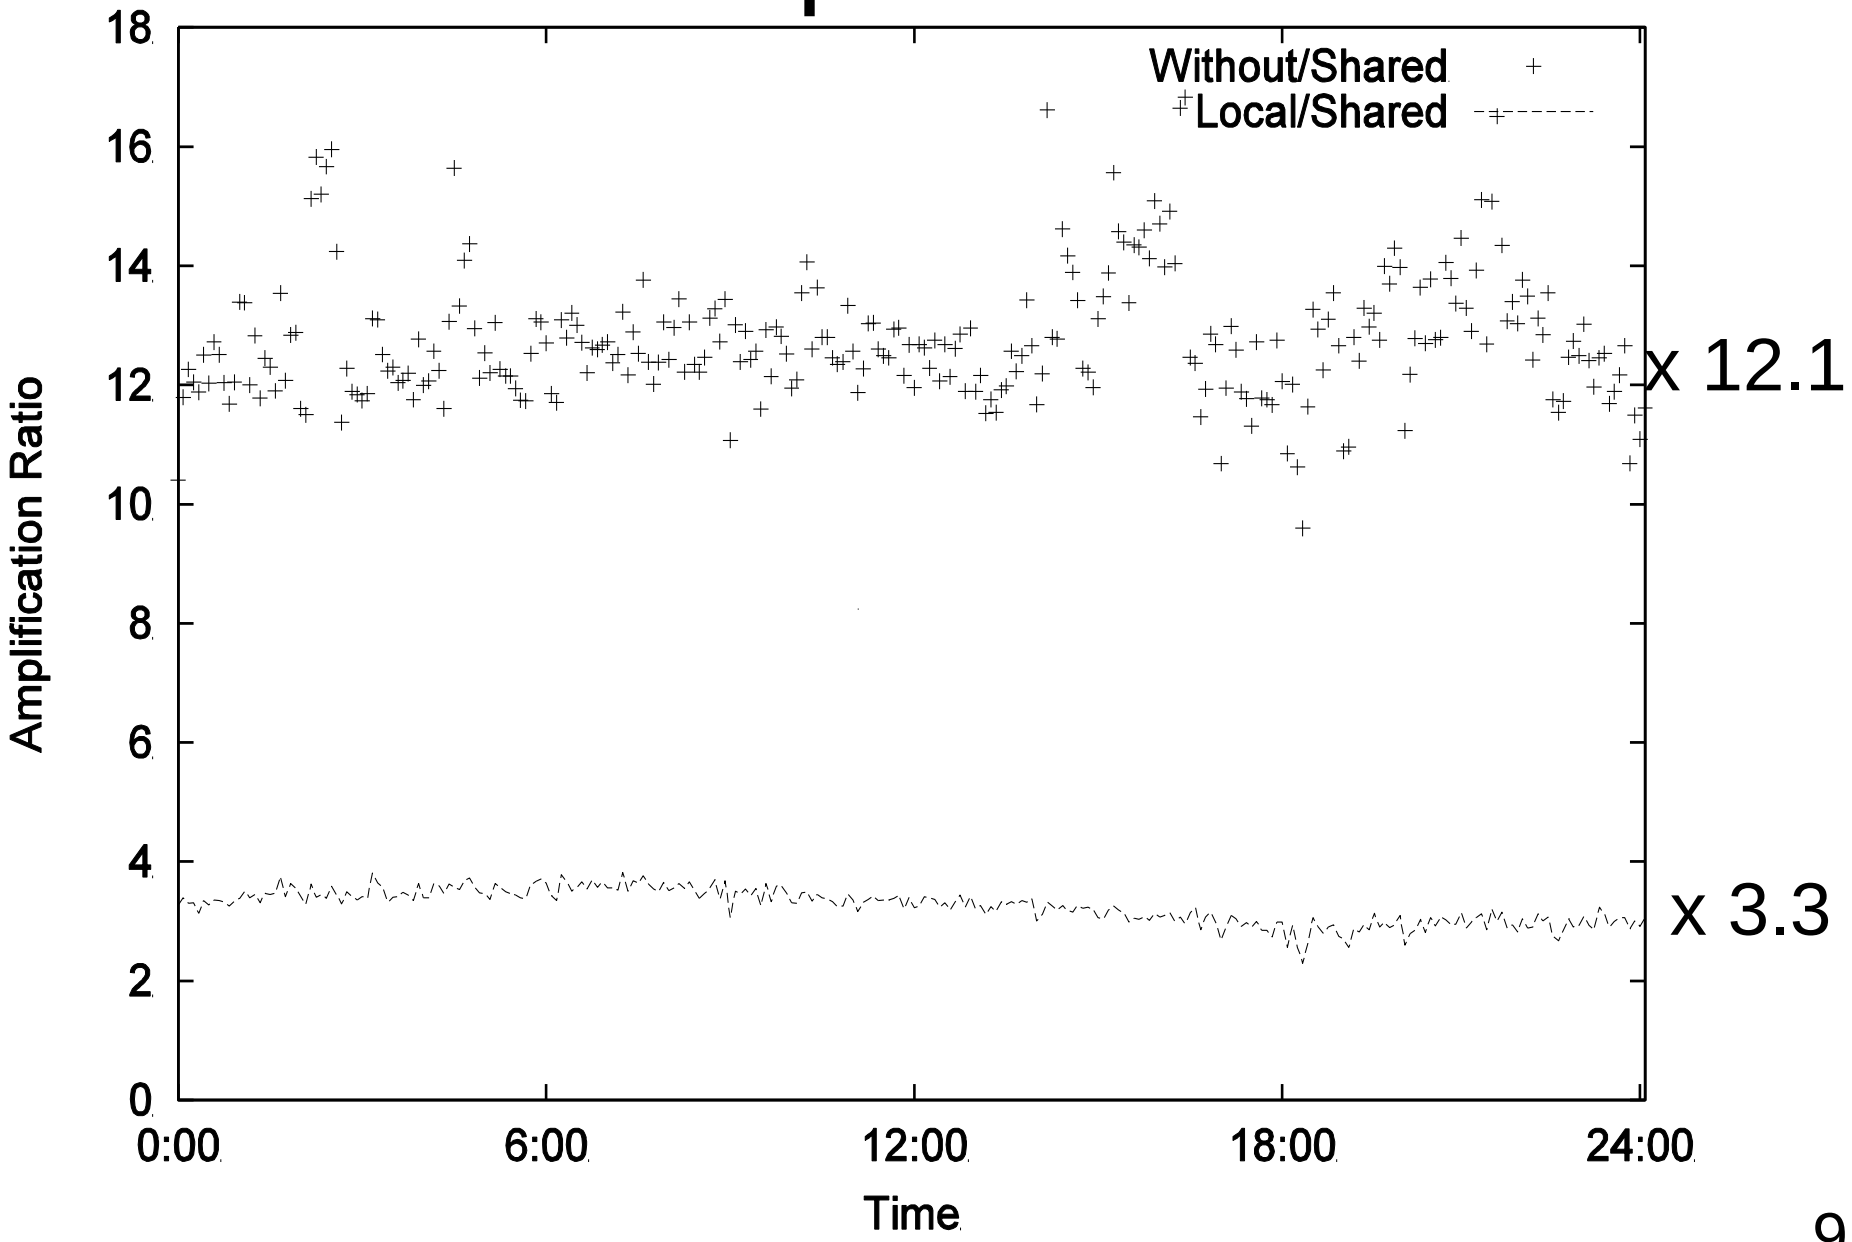

Query Amplification Ratio per TTL

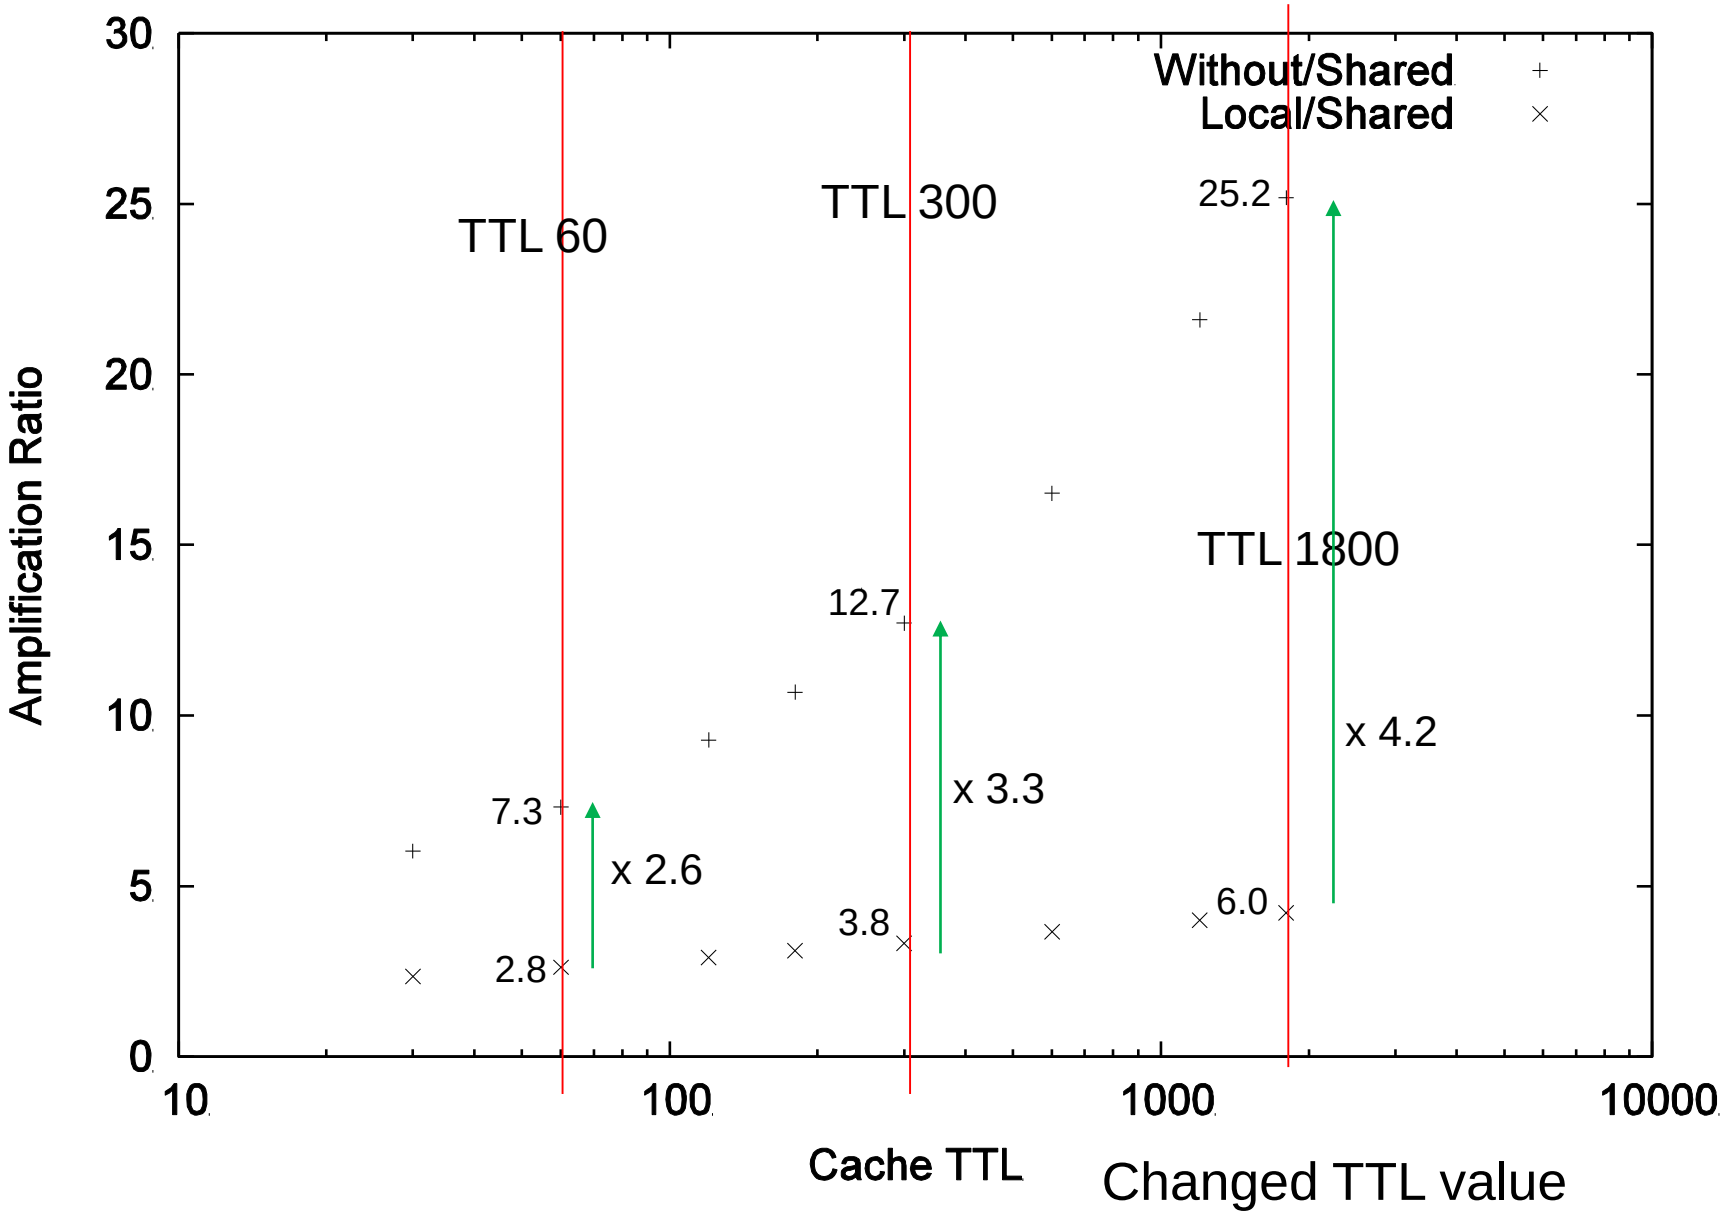

Number of DNS queries

caching mechanism is effective with TLD level

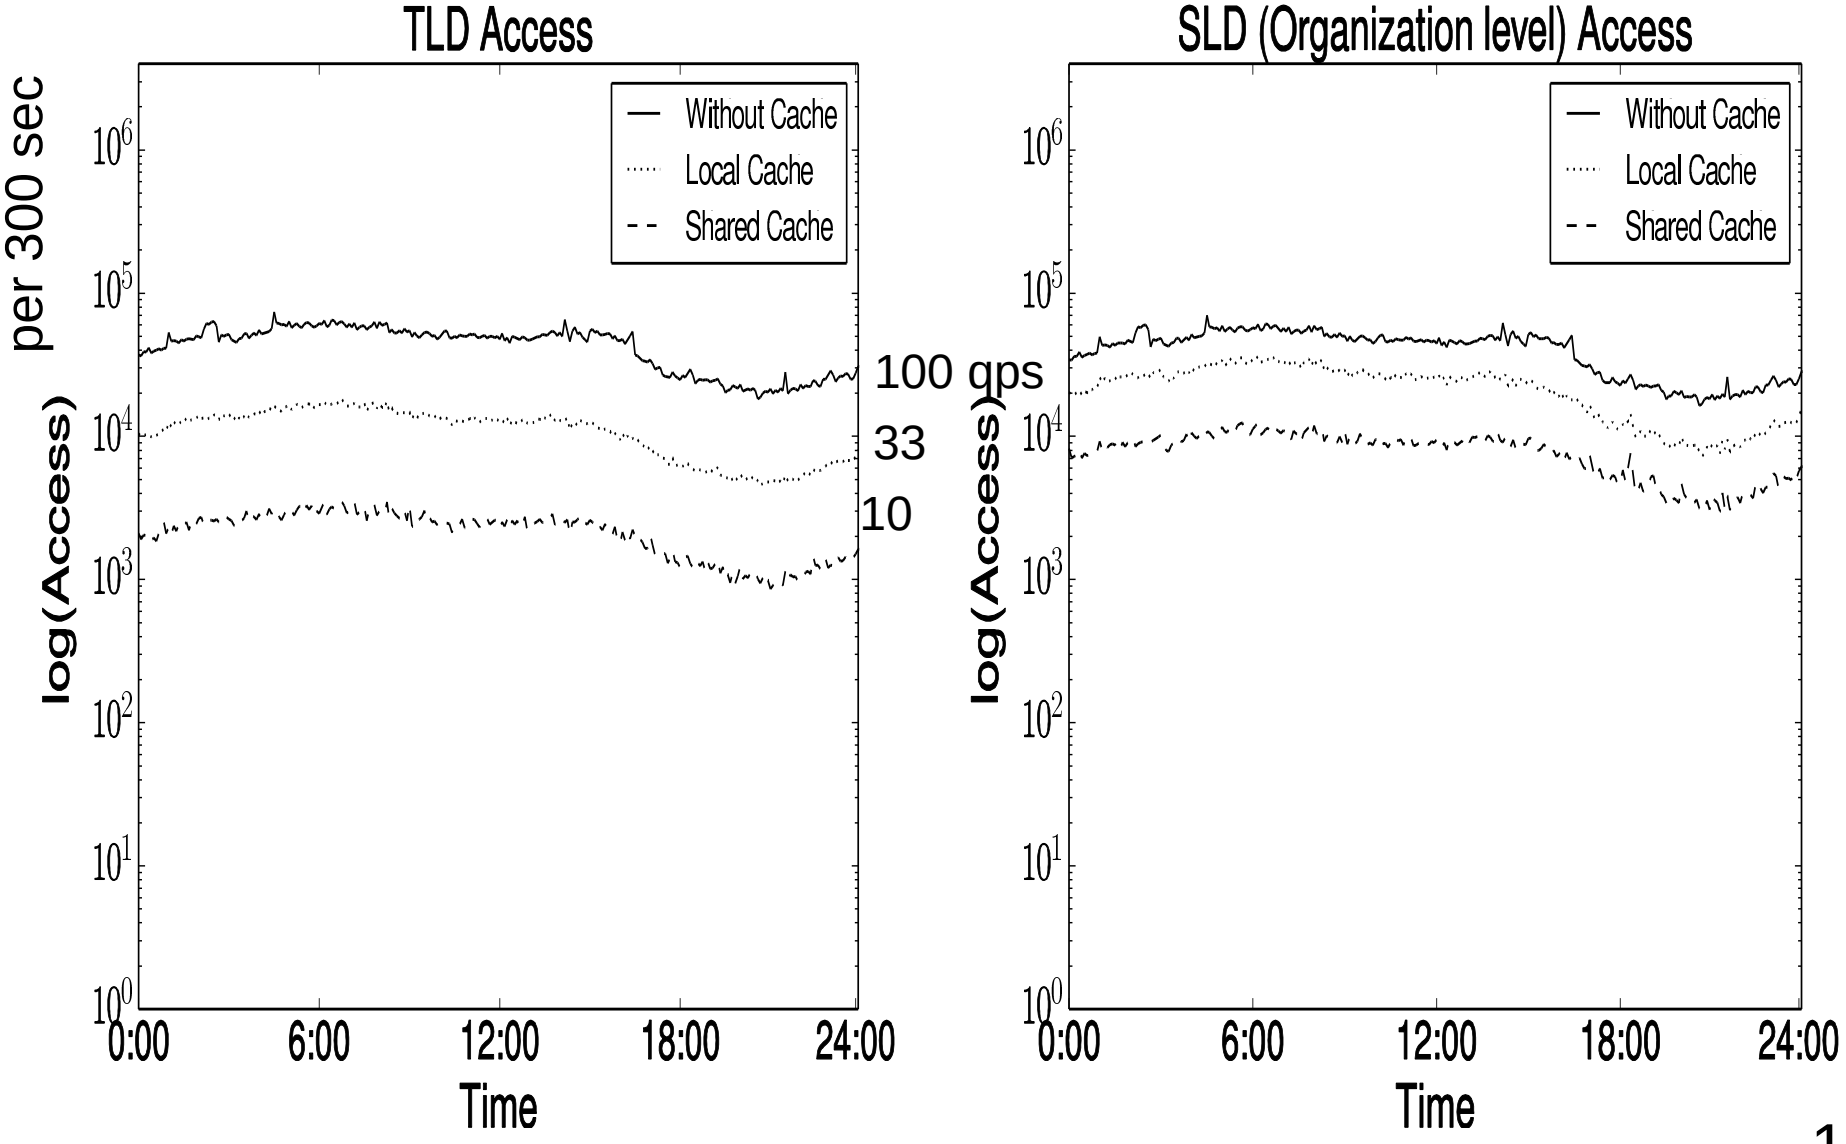

Query Amplification Ratio

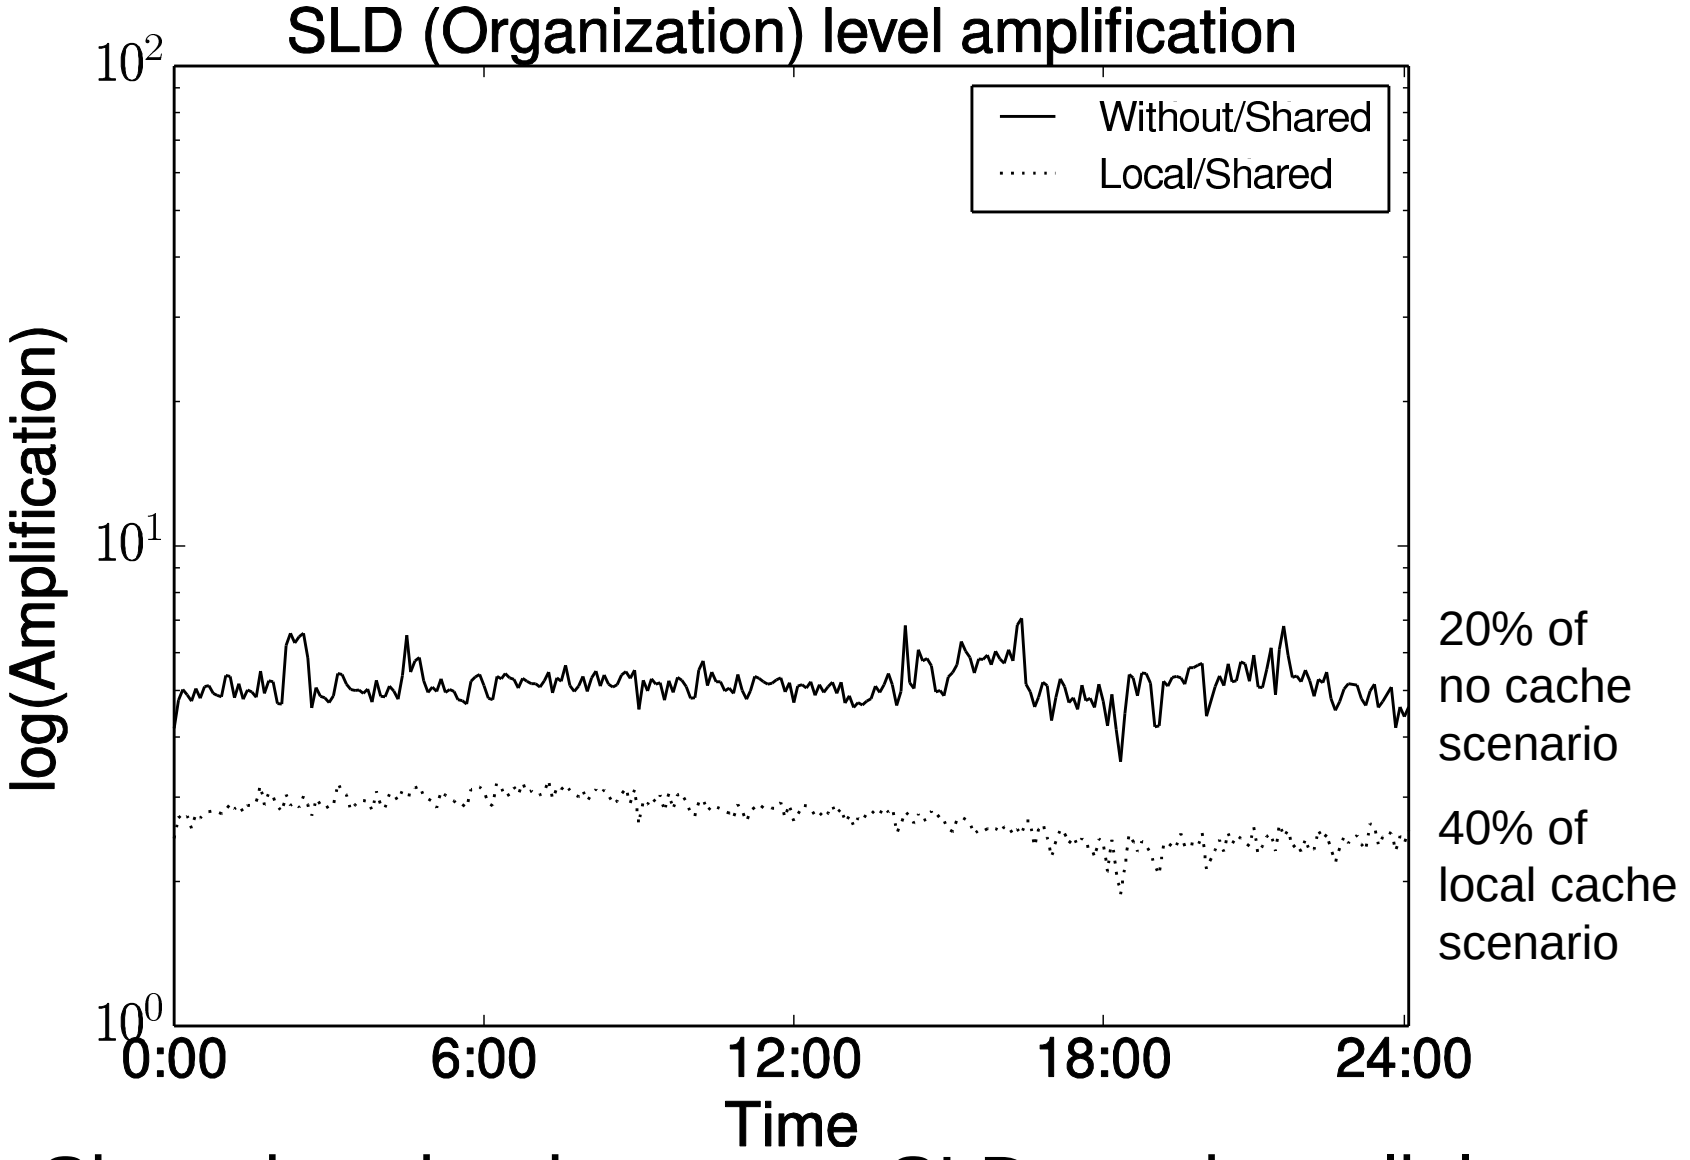

Shared cache decreases SLD queries a little.

Query Amplification Ratio

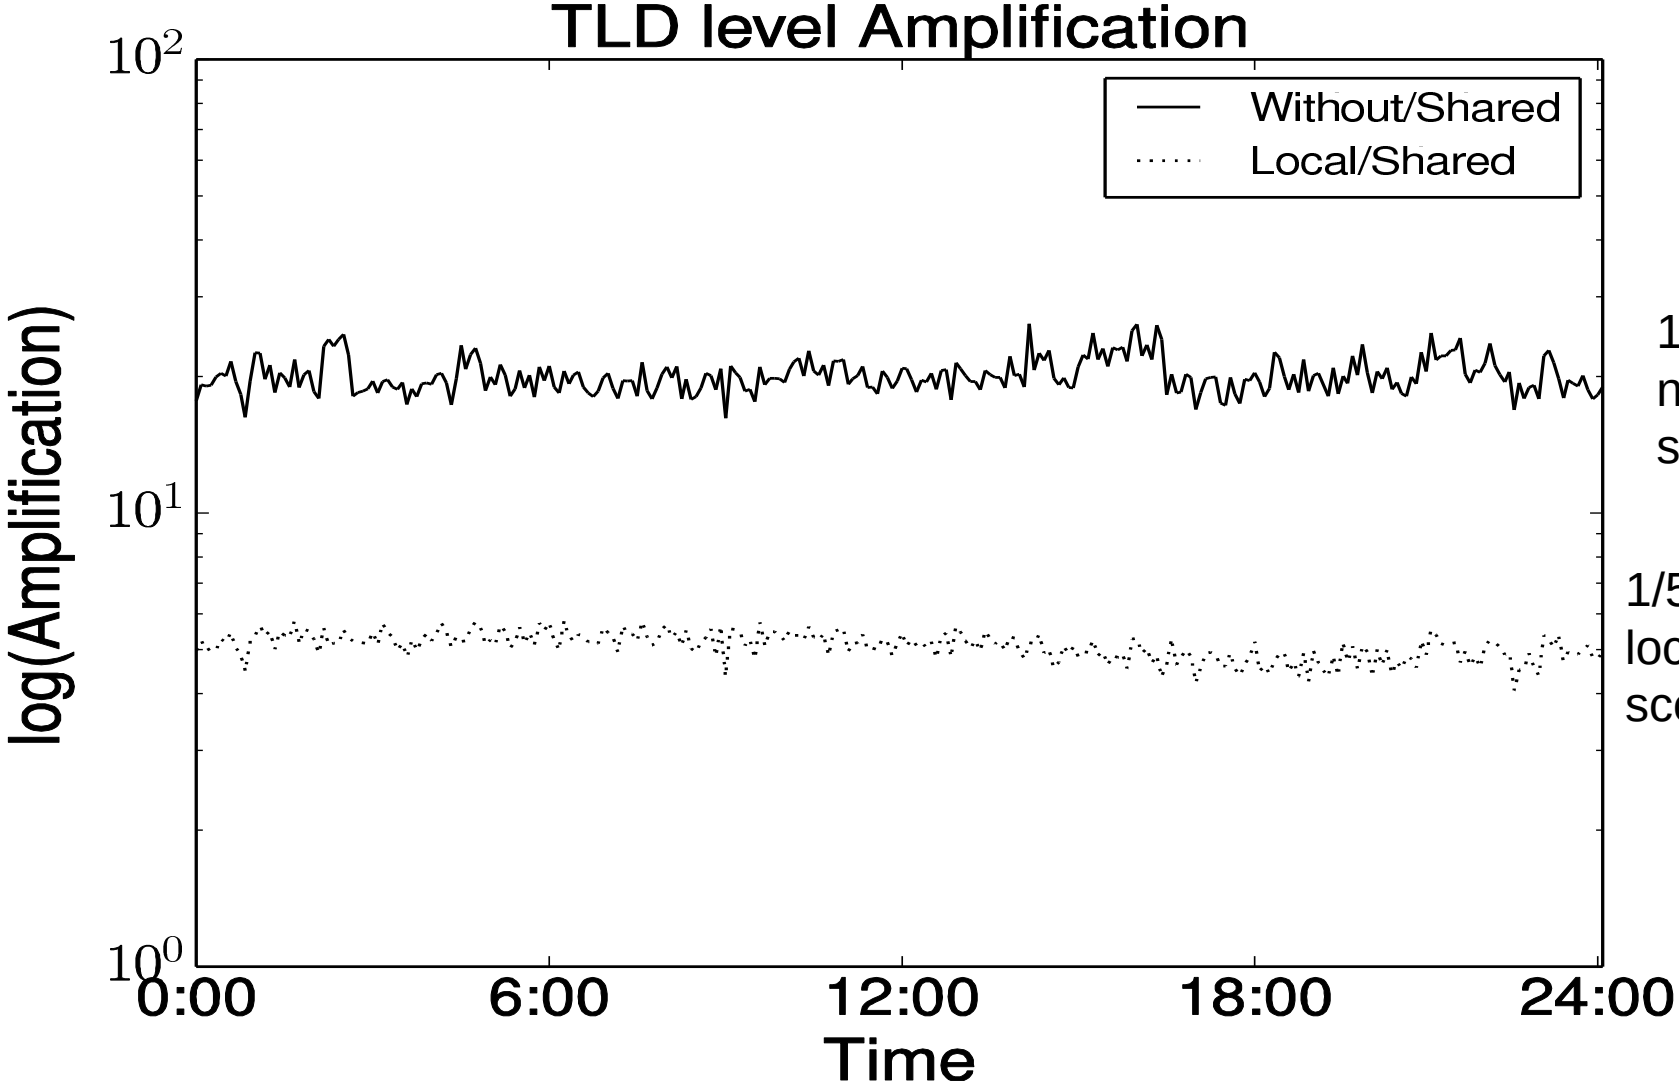

1/20 of no cache scenario

1/5 of local cache scenario

Shared cache decreases TLD queries

Query Amplification Ratio

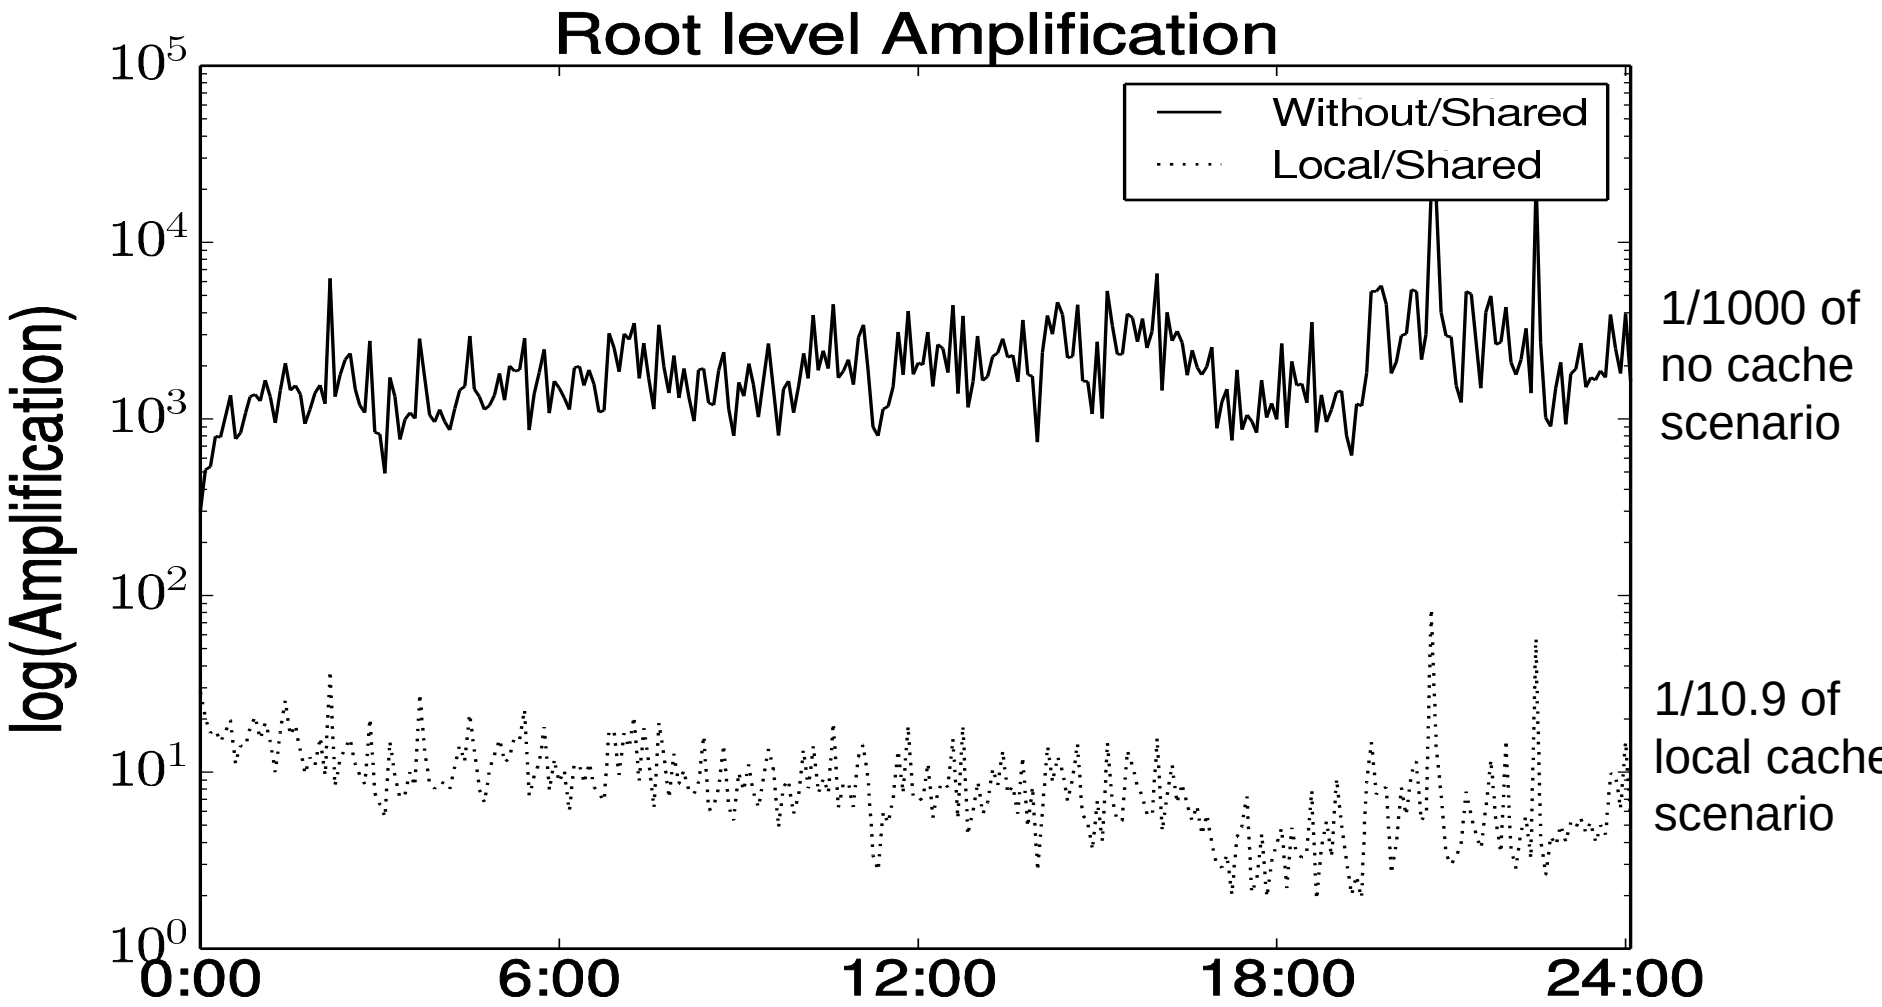

Shared cache decreases root queries

Removing shared cache is harmful

Frequent Queries and Active Clients

- During the experiments, we found:
  - A few clients repeatedly issue many of the same DNS queries in a short interval, and also
  - a few QNAMEs are repeatedly issued in a short interval

Frequent queries (QNAMEs)

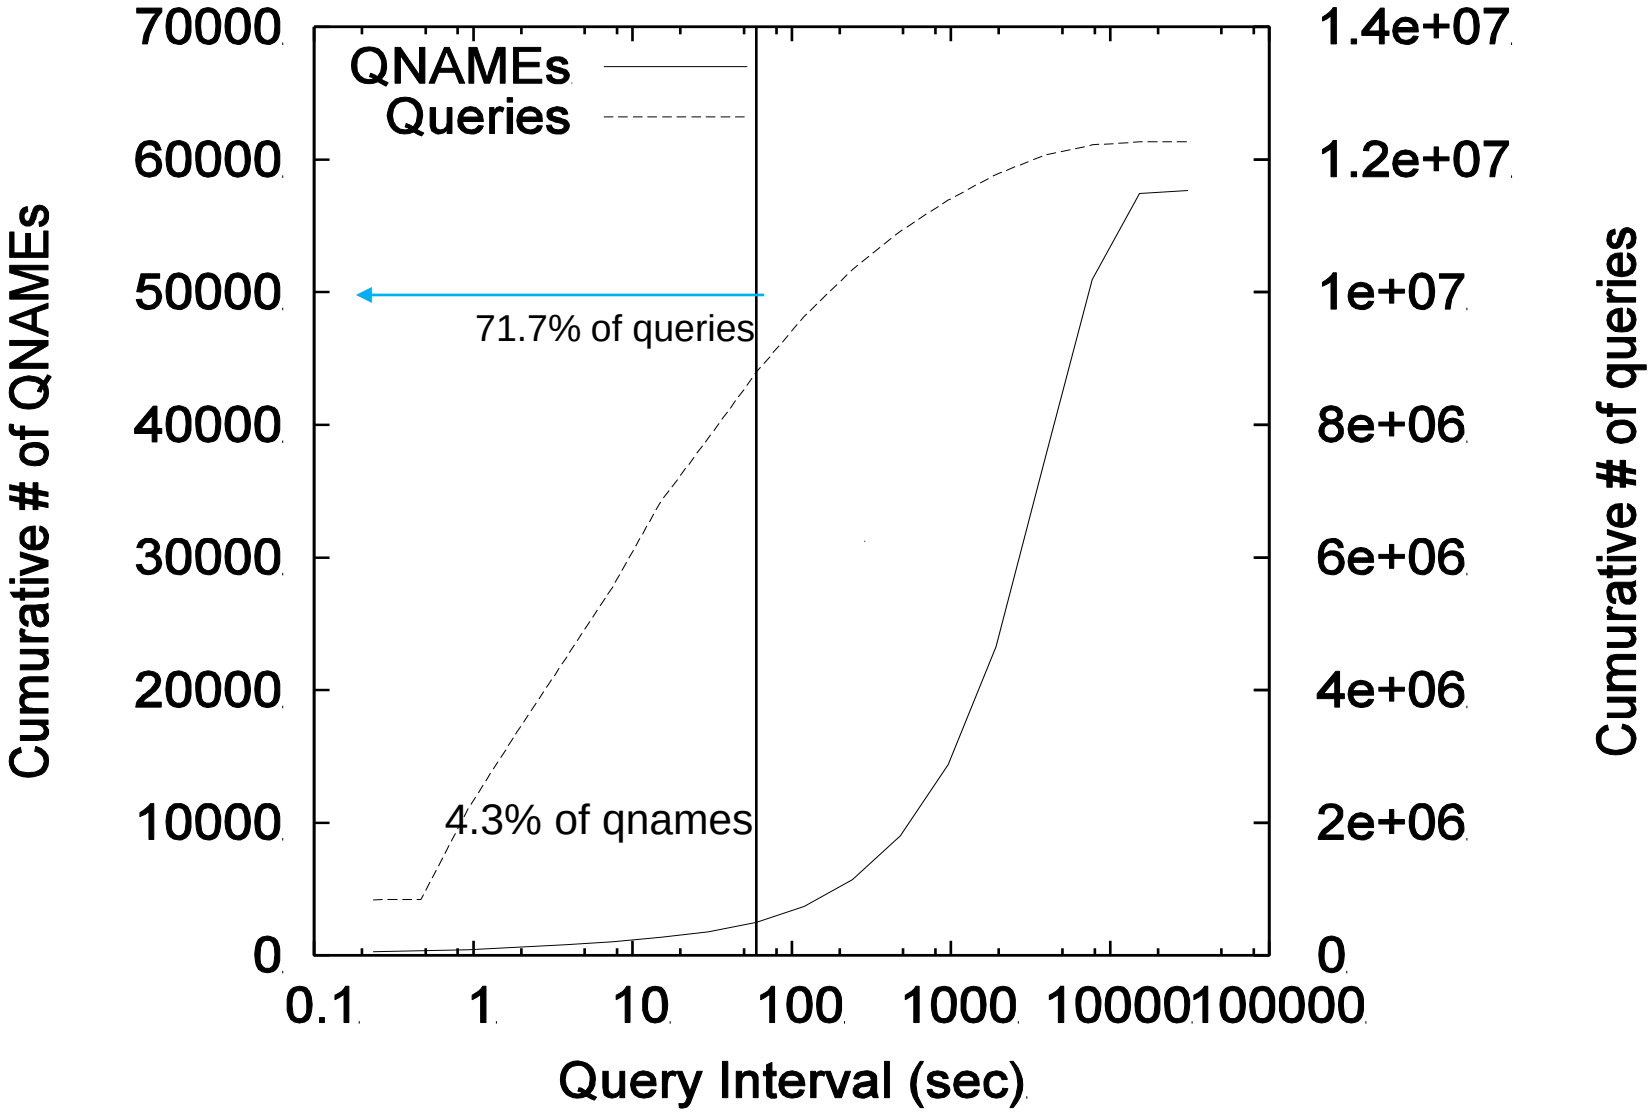

Frequent clients

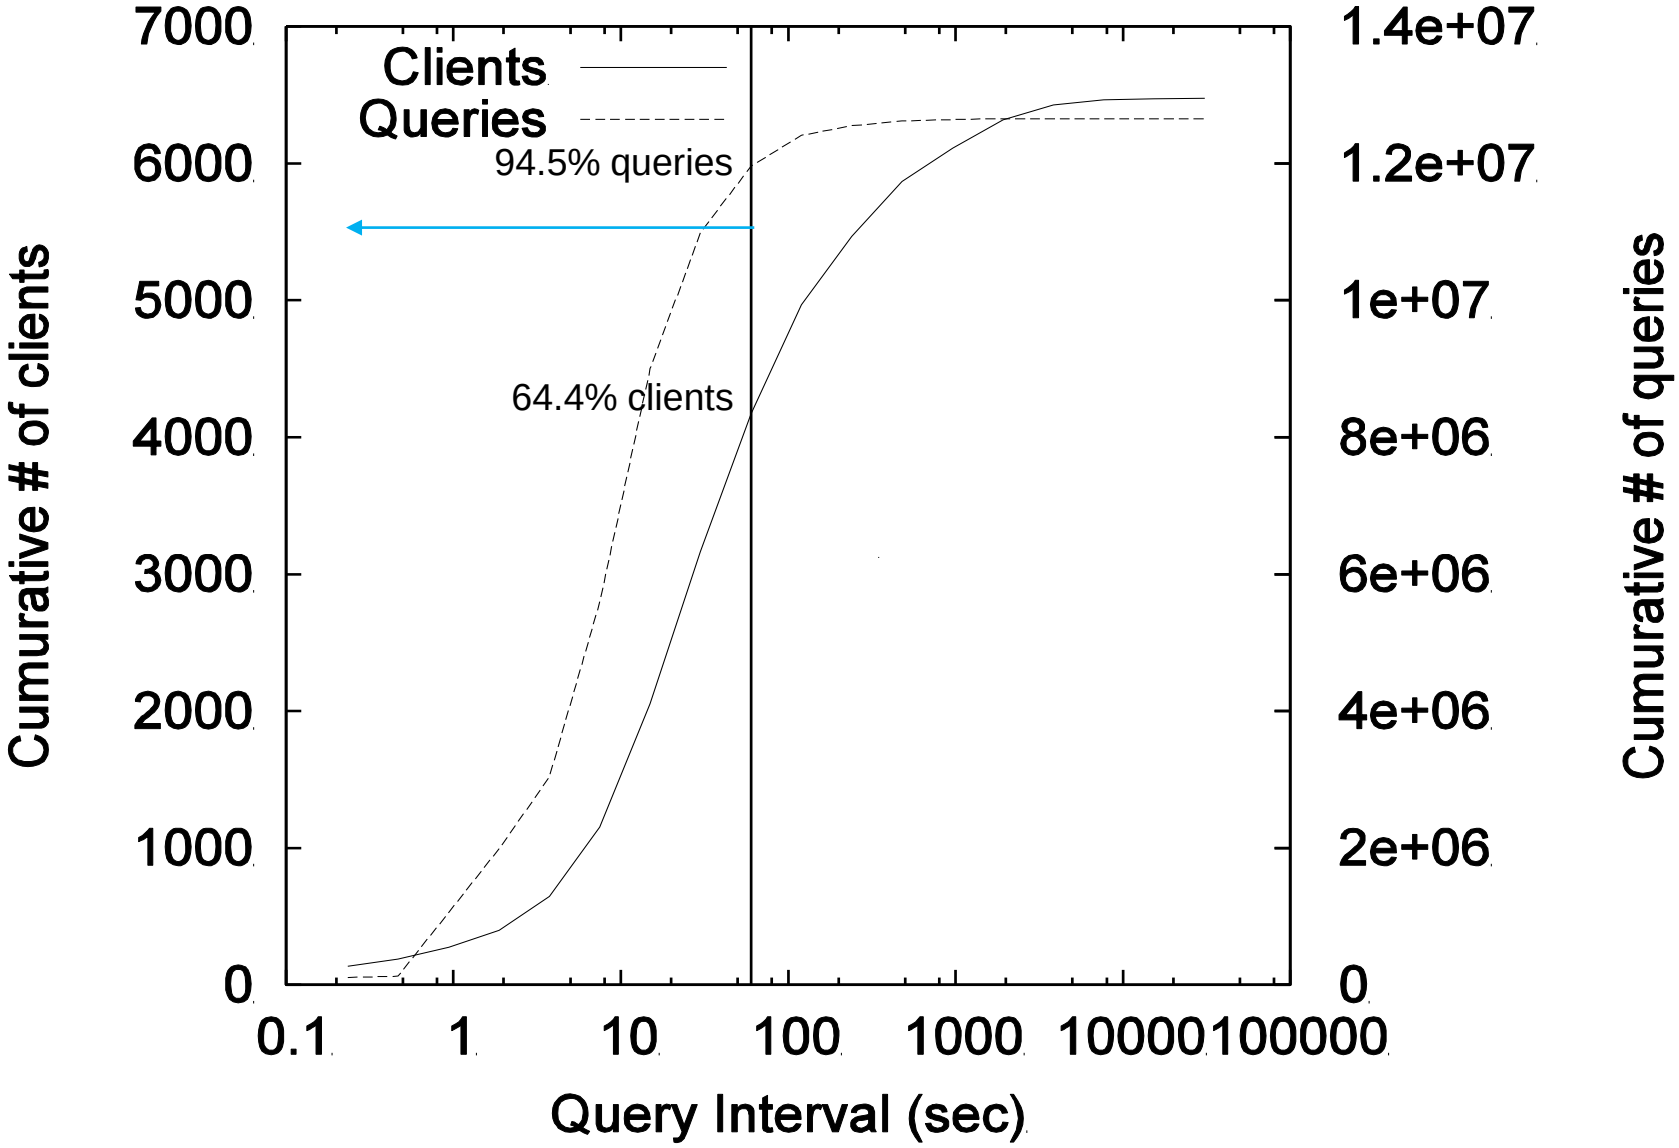

Most active client

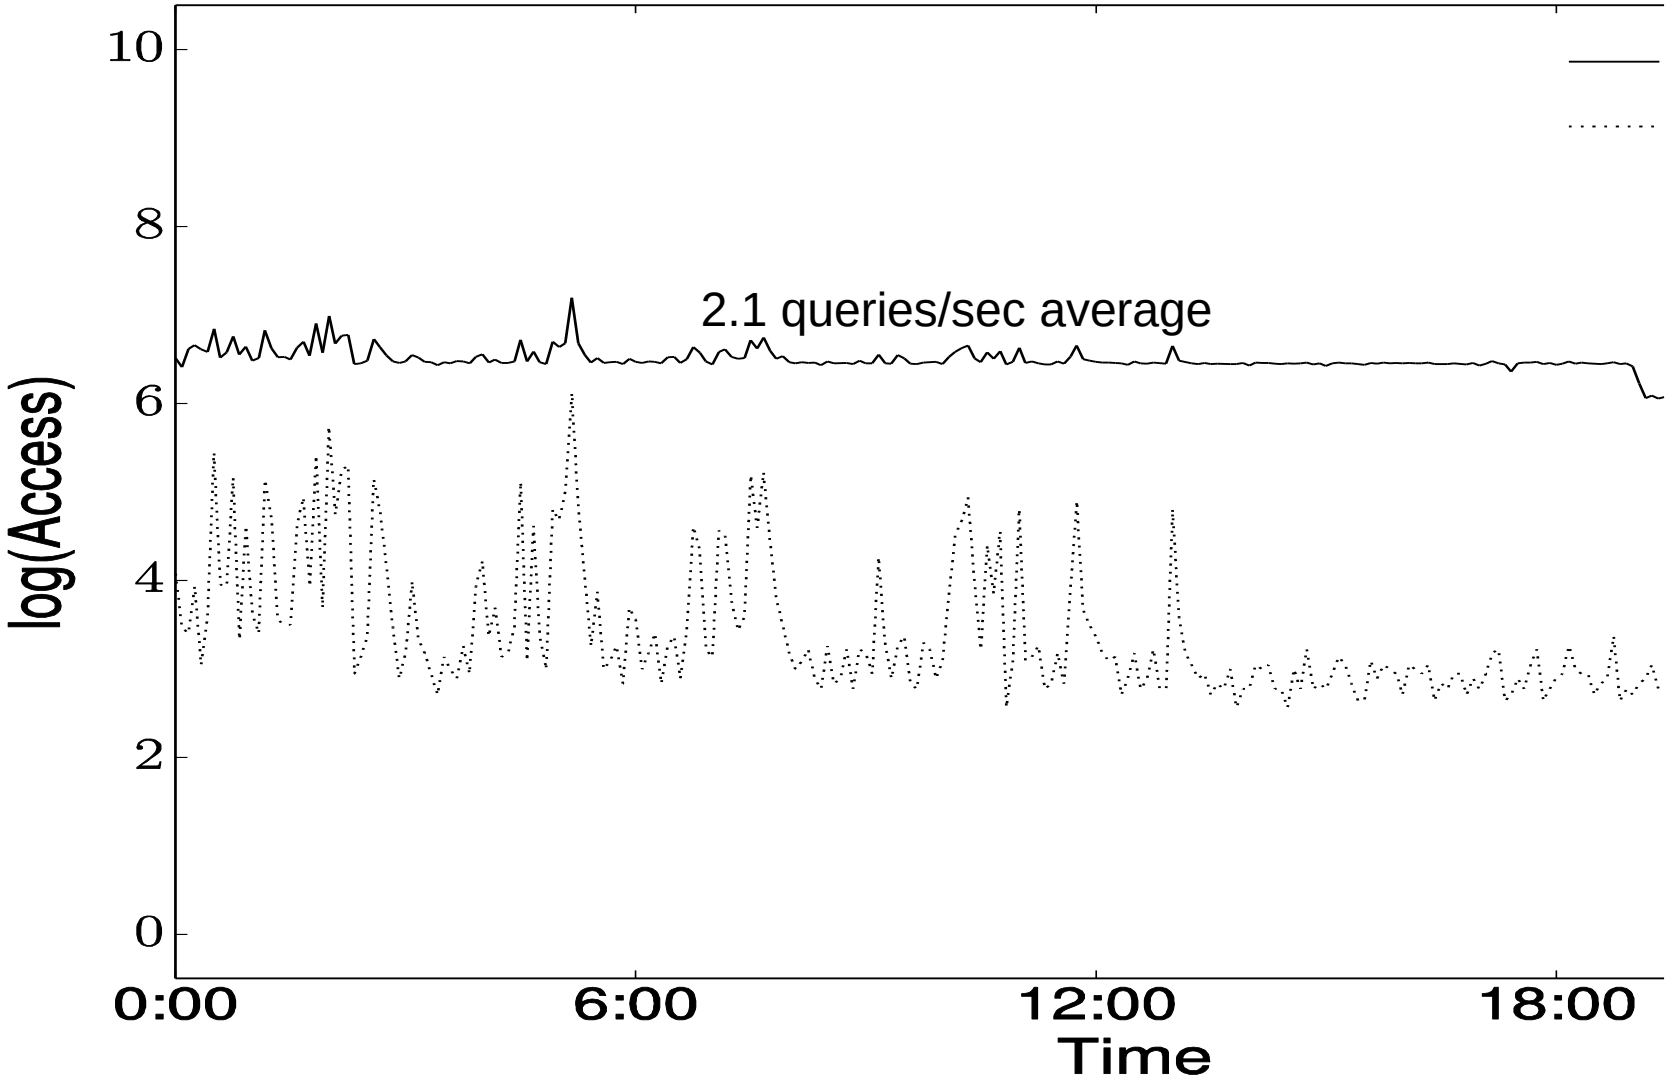

(2.1 qps = 630/300sec, log(630)=6.4

Conclusion

- Removal of shared resolver (replace with local resolver) amplifies the DNS traffic by about 3.3 times.
- The amplification ratio on the root DNS servers is much worse (about 10.9 times).
  - Removal of shared resolver may be harmful
  - (Local cache is useful)
- Some systems (Linux, BSD) lack local cache
  - repeated queries at short intervals (<= 1 min)
  - about 94.5% of client queries
  - the deployment of local cache itself is effective

Future works

- Long term analysis
- Comparison with real traffic
- Detailed full-resolver simulation
- Analysis of new DNS standards that improve name resolution performance
  - RFC 8020: NXDOMAIN: There Really Is Nothing Underneath
  - RFC 8198: Aggressive use of DNSSEC-validated Cache