Privacy-Preserving DNS

Analysis of Broadcast, Range Queries and Mixes

Hannes Federrath, Karl-Peter Fuchs,
Dominik Herrmann, Christopher Piosecny
University of Hamburg (Germany)

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Agenda

Missing Privacy in DNS

Characteristics of DNS Traffic

DNS Anonymity Service

Range Queries

Privacy Issue: DNS Resolver learns queries of all users

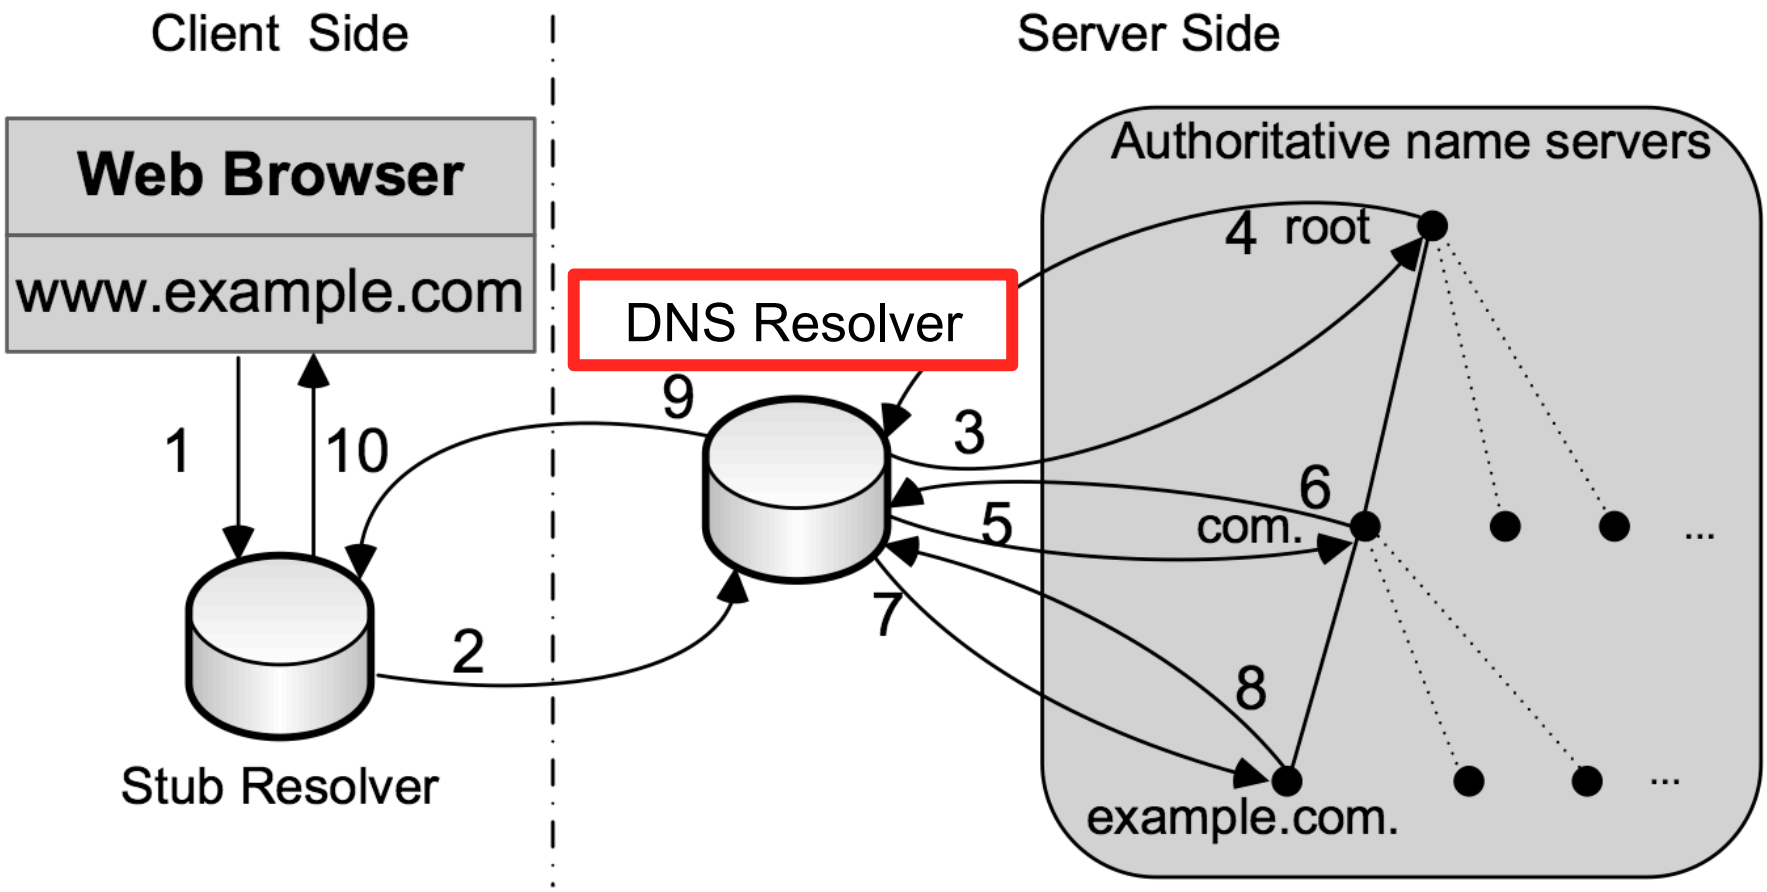

[Lu & Tsudik, 2010]

Third-party DNS Resolvers are increasing in popularity

Google, OpenDNS, Comodo, Norton DNS, ...

Advertised benefits:

Objectives for the DNS Anonymity Service

1. protect privacy of users
   - hide relationship between users and queries from resolver

2. practicable and usable solution
   - very low latency
   - compatibility with existing DNS

Agenda

Missing Privacy in DNS

Characteristics of DNS Traffic

DNS Anonymity Service

Range Queries

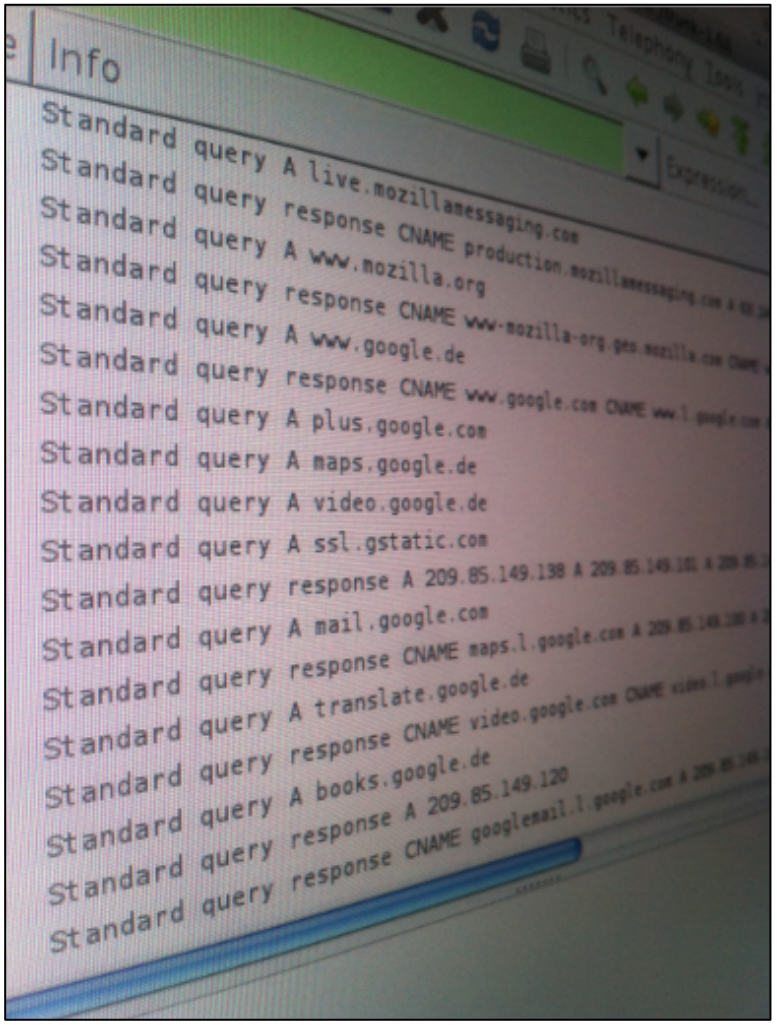

Overview of our DNS dataset

We obtained real-life DNS traces:

- DNS query log of a German university campus network
- >4000 distinct users (on average 2100 active per day)

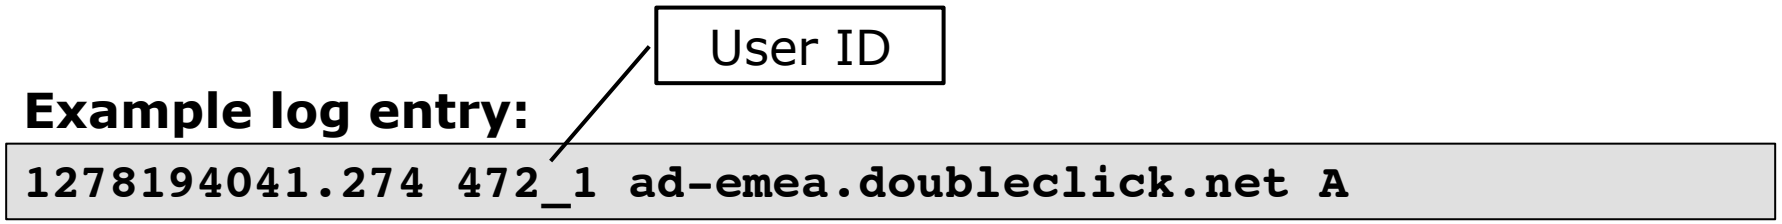

Additionally, for each hostname we have recorded

- TTL value
- query and reply size
- lookup latency (using Google's DNS Resolver)

Characteristics of DNS traffic

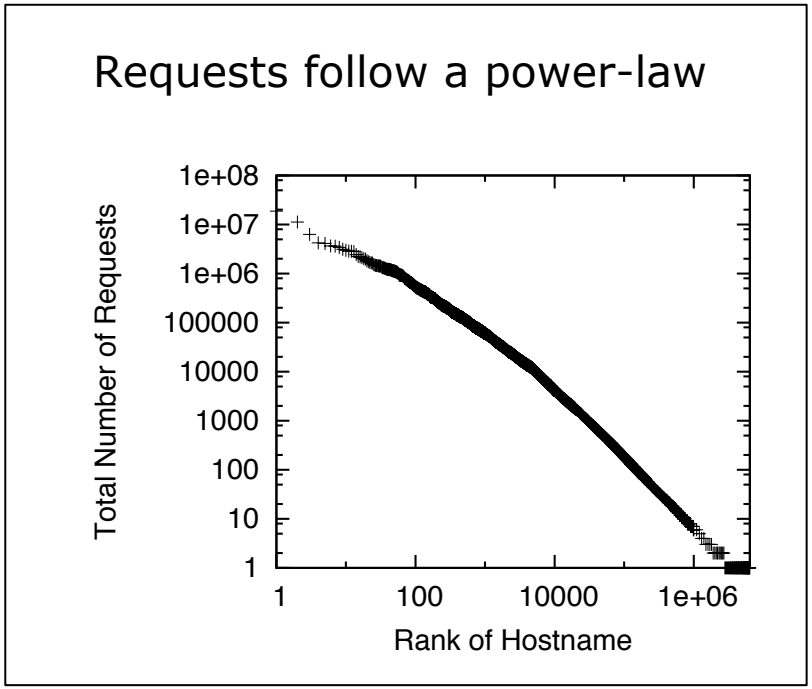

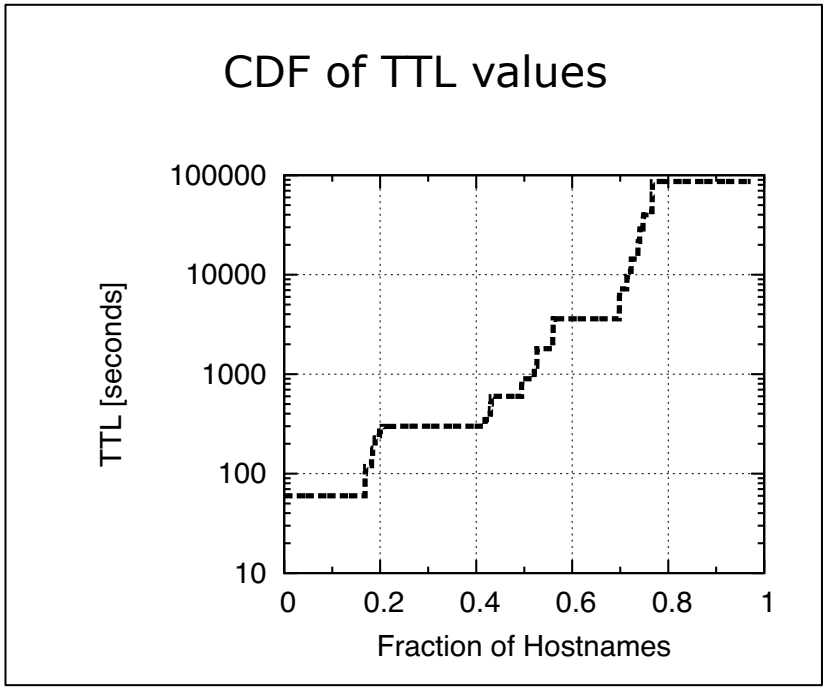

- 80% of queries are for top 10,000 hostnames
- regardless of TTL most RRs remain constant for a long time

Characteristics of DNS traffic

- almost every website visit causes a DNS query burst

Firefox
without prefetching

en.wikipedia.org

upload.wikimedia.org
nn.wikipedia.org
th.wikipedia.org
creativecommons.org
www.wikimediafoundation.org
www.mediawiki.org

Chrome
with prefetching

en.wikipedia.org

geoiplookup.wikimedia.org
commons.wikimedia.org
el.wikipedia.org
en.wikibooks.org
en.wikinews.org
en.wikiquote.org
en.wikisource.org
en.wikiversity.org
en.wiktionary.org
et.wikipedia.org
gl.wikipedia.org
lists.wikimedia.org
simple.wikipedia.org
species.wikimedia.org
wikimediafoundation.org
www.wikilovesmonuments.de

Agenda

Missing Privacy in DNS

Characteristics of DNS Traffic

DNS Anonymity Service

- **Broadcast**
- **Mixes**

Range Queries

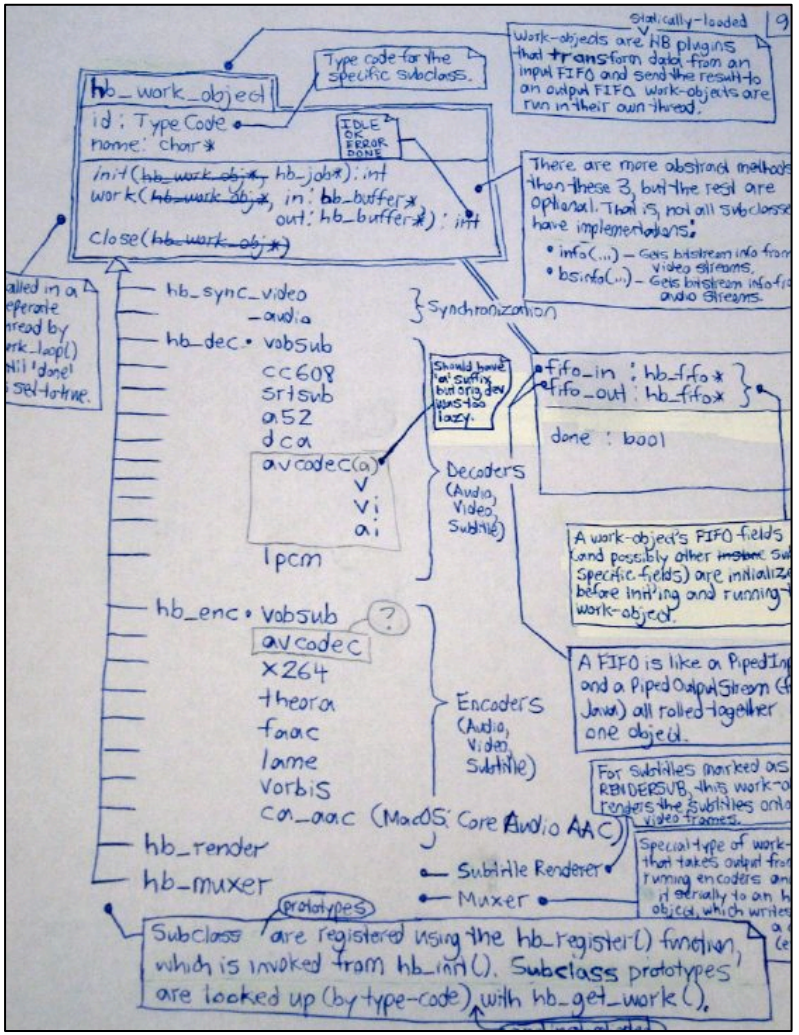

Architecture of the proposed DNS Anonymity Service

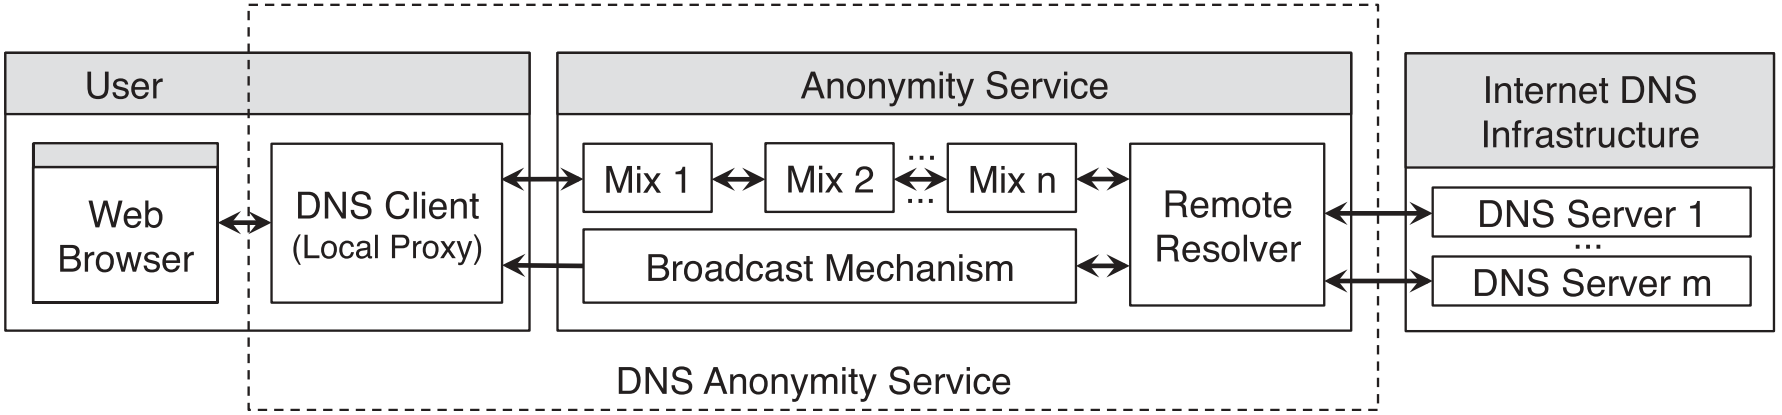

- drop-in replacement for DNS Resolver
- two building blocks
  - broadcast mechanism
  - mixes cascade

Motivation for broadcasting

What if each client had a local copy of the full DNS database?

- clients get **zero lookup latency**
- all DNS queries are **unobservable**

Motivation for broadcasting

What if each client had a local copy of the full DNS database?

- clients get **zero lookup latency**
- all DNS queries are **unobservable**

not practical

We can exploit the power-law distribution of queries!

- compromise: local copy **for most popular hostnames** only

Anonymity Service

- monitors most popular hostnames for updates — Central Update
- provides full copy of database to new clients — Initial Download
- broadcasts changed resource records to clients — Increm. Updates

Evaluate implementation in **trace-driven simulations**

Broadcasting is promising and practicable

Hit Rate	Required Traffic
100 entries 40%	
10,000 entries 83.9%	
100,000 entries 94.5%	

Broadcasting is promising and practicable

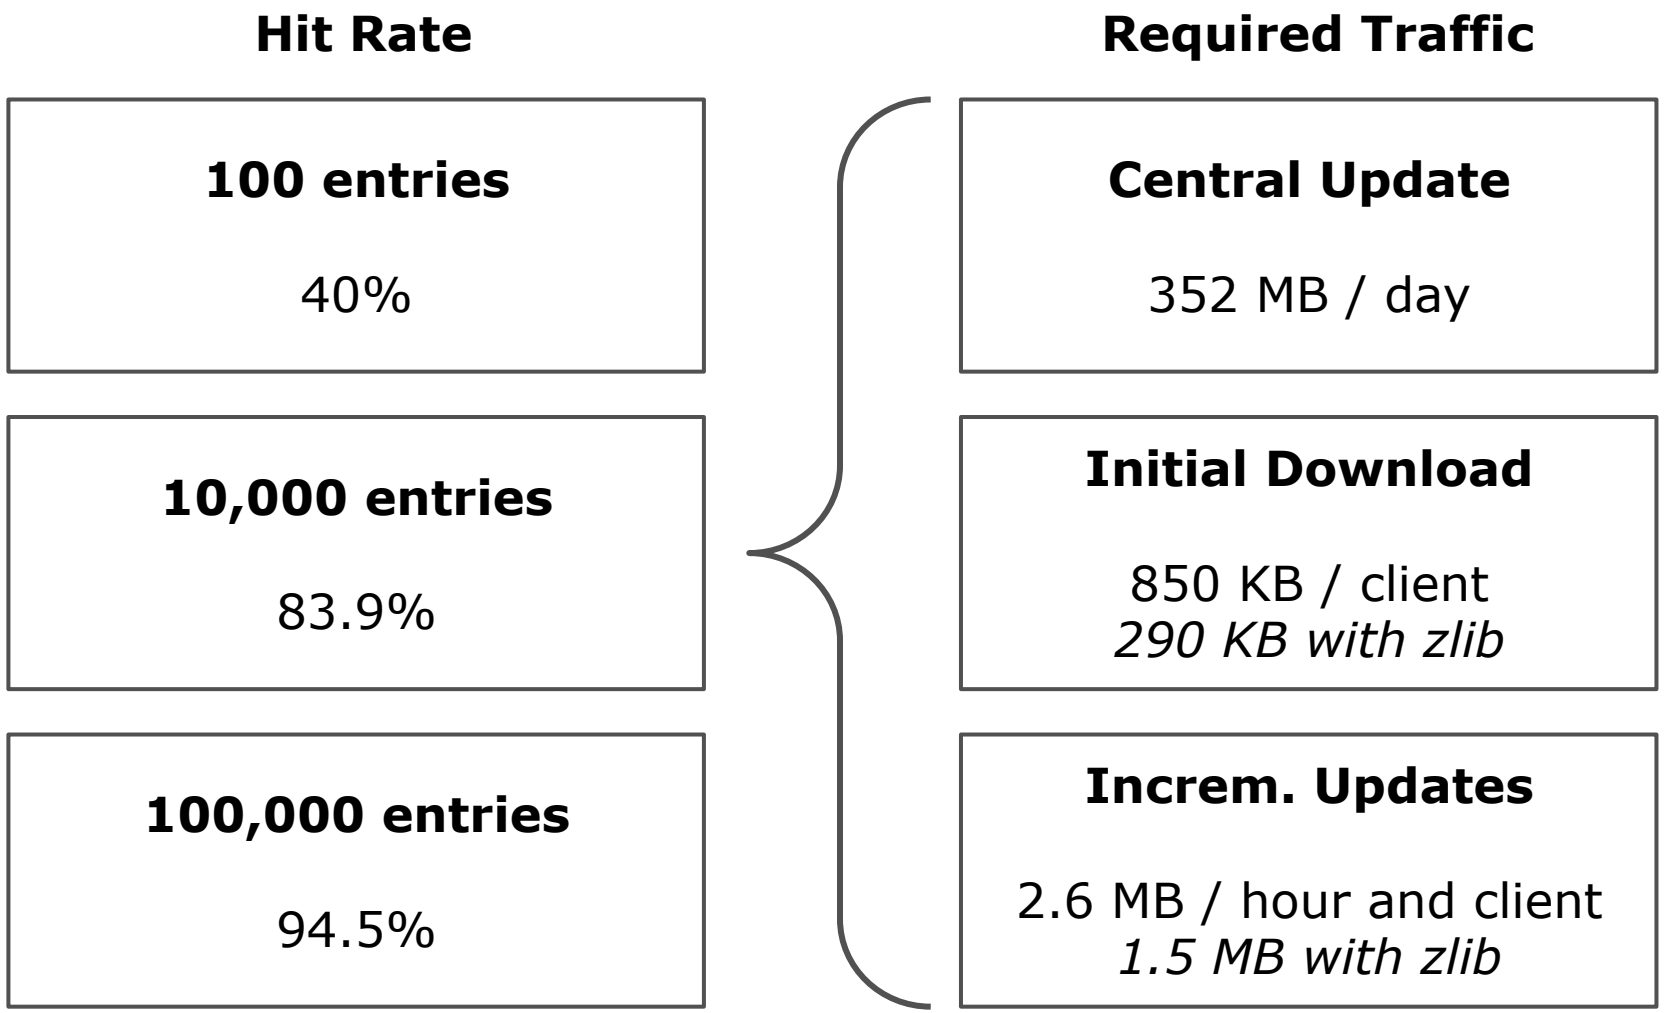

Anonymise remaining queries with mixes

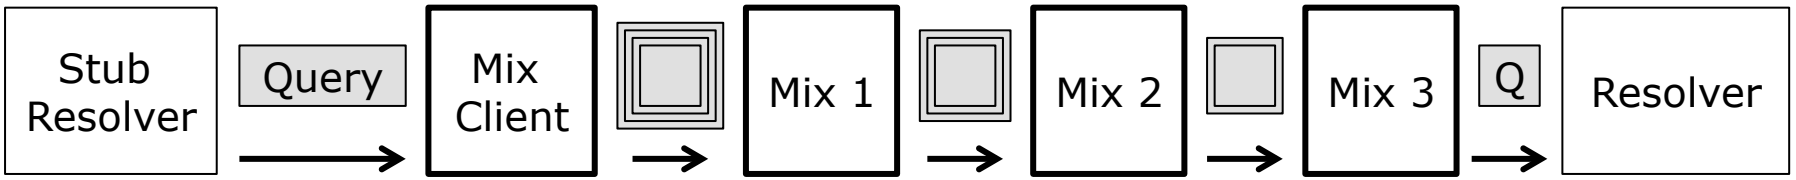

- Motivation:
  - already deployed in practice (Tor, AN.ON)
  - attacker model of practical systems reasonable for DNS
- Performance impact: cryptographic operations, network latency
- Implementation specifics
  - **channels** for low latency (re-established after 60s)
  - **fixed-size messages** (queries: 57 bytes, replies: 89 bytes) to counter traffic analysis
  - Java, BouncyCastle, RSA (2048 bit), AES (128 bit OFB)

Performance evaluation of our implementation

Trace-driven simulation using recorded lookup delays

- 2082 concurrent users
- 107 queries/sec

- DNS traffic increases by 100% (240 KB per day)
- Latency results are also promising

percentile	**50%**	**90%**
without mixes	9.2 ms	46.2 ms
3 mixes (LAN)	10.9 ms	52.0 ms
3 mixes (WAN)	171 ms	274 ms

mix-mix RTT	20ms
client-mix RTT	80ms

- Congestion once >1000 queries/sec issued

→ **Performance of mixes appears to be satisfactory for DNS**

Agenda

Missing Privacy in DNS

Characteristics of DNS Traffic

DNS Anonymity Service

Range Queries

Related Work: Range Queries

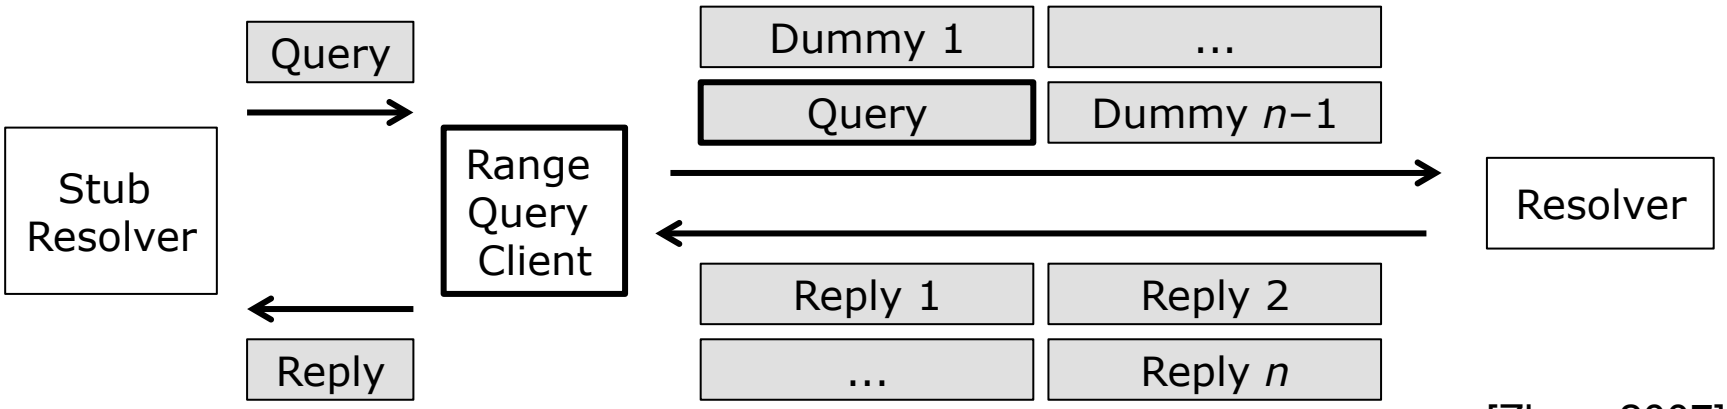

[Zhao, 2007]
[Castillo-Perez, 2008]

- hide actually desired queries using $n-1$ dummy queries
- should offer low latency; but no trace-driven evaluation so far

Also related, but not of interest for us: **PPDNS** [Lu & Tsudik, 2010]

- implements cPIR
- is built on top of CoDoNS

Trace-driven evaluation of range queries

- We implemented a **range query simulator**
  - clients draw *n*–1 dummies randomly from set of all hostnames
  - range queries are compressed using zlib
  - transmitted via TCP to Range Query DNS Resolver
- Trace-driven simulation using recorded lookup delays
- Evaluation using our DNS traces
  → traffic volume increases x4 for n=10, x24 for n=100
- Basic implementation
  - each reply is returned independently to the client
  - latencies do not increase considerably – even for n=1000
- **But:** attacker can exploit dependencies of consecutive queries!

Timing attack based on traffic bursts

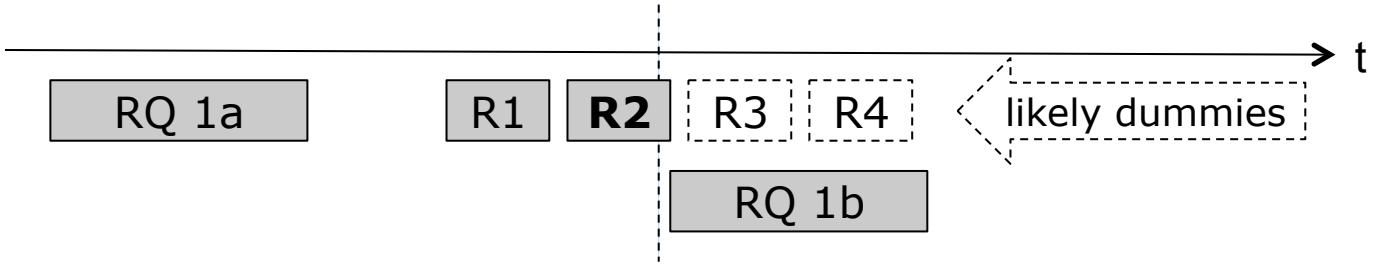

Preventing the timing attack

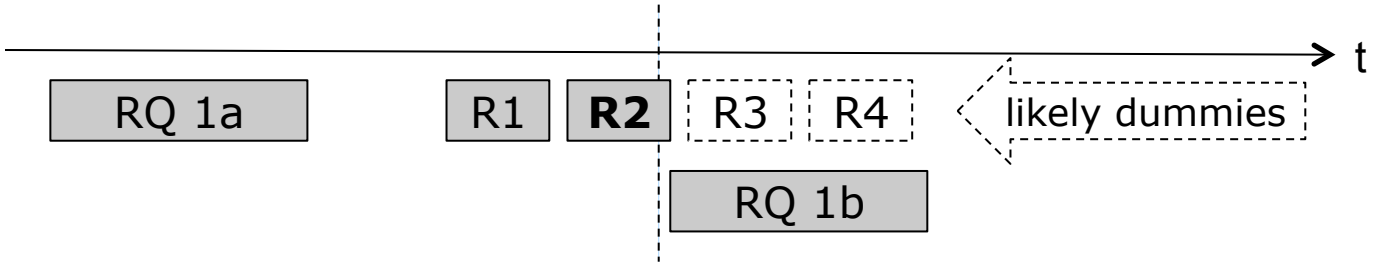

1. stall desired reply

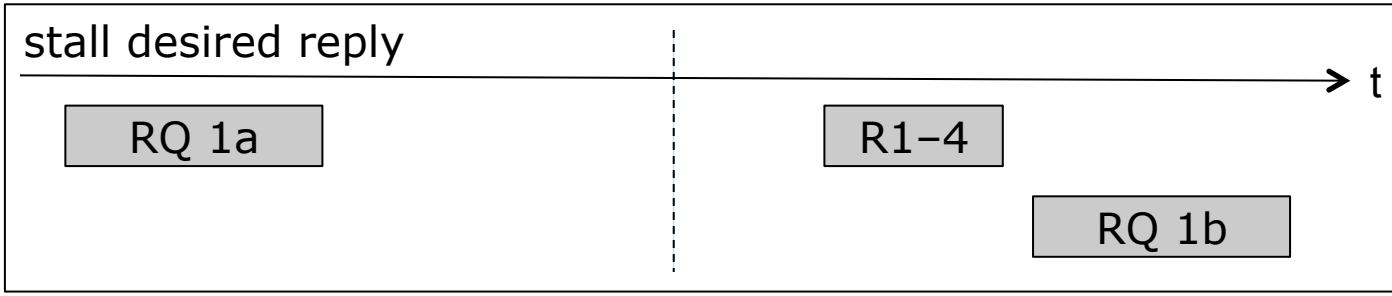

2. delay consecutive query

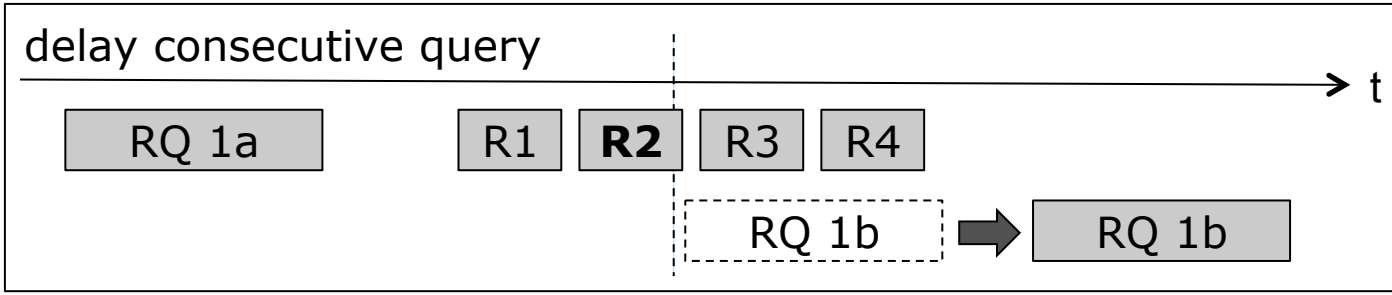

Preventing the timing attack is expensive

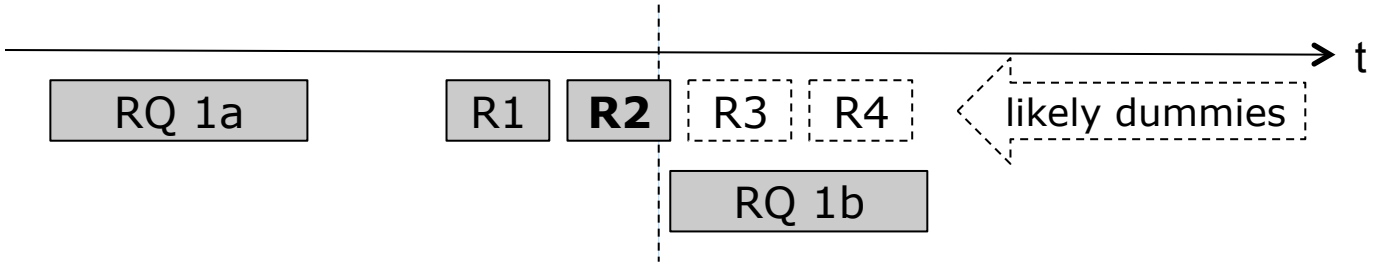

median latency n=10

1. stall desired reply

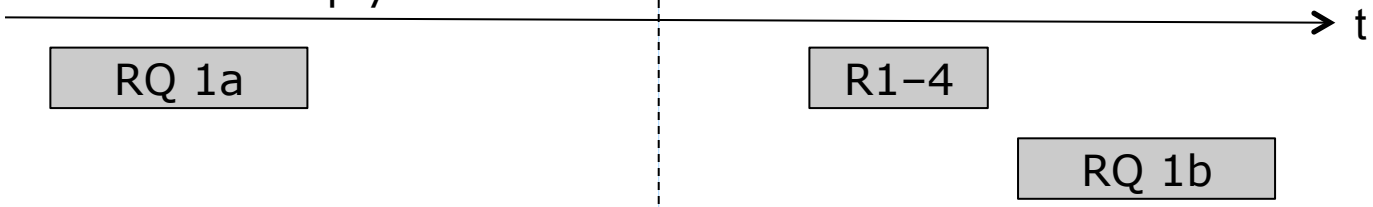

200ms

2. delay consecutive query

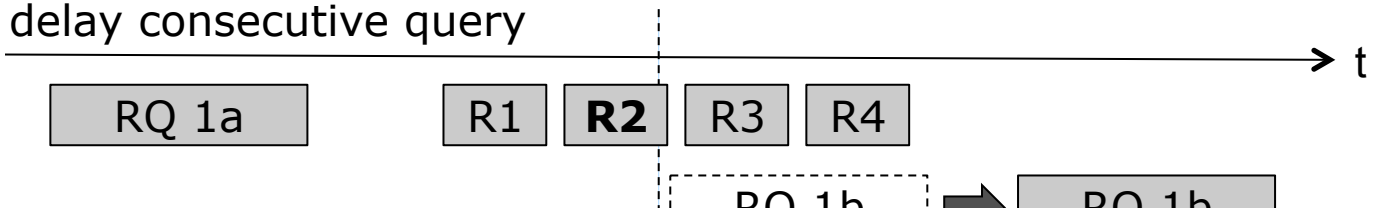

400ms

Open question: how to prevent semantic intersection attack?

Summary

Missing Privacy in DNS

- queries leak to DNS Resolver
- low-latency, practical solution

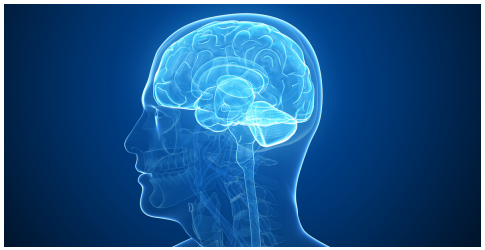

Characteristics of DNS traffic

- power-law distribution
- query bursts

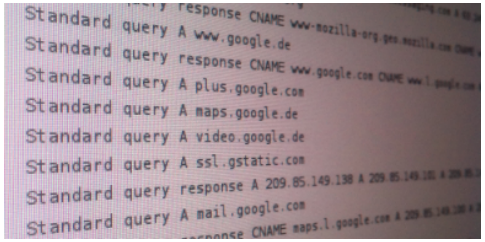

Proposed DNS Anonymity Service

- broadcast: zero latency + unobservability
- mixes: satisfactory performance

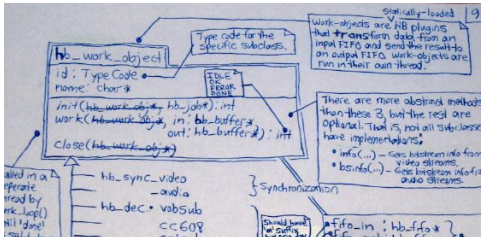

Evaluation of Range Queries

- fast for isolated queries
- preventing timing attack is expensive