T-DNS: Connection-Oriented DNS to Improve Privacy and Security

IETF 89 DNSE BOF

USC/ISI

John Heideman

Zi Hu

Liang Zhu

Verisign Labs

Allison Mankin

Duane Wessels

Why Consider T-DNS

- Privacy – Lacking encryption, vanilla DNS is susceptible to eavesdropping; especially so given widespread use of WiFi and third-party recursive DNS services.
- Spoofing – UDP's connectionless nature makes it ideal for use in reflection/amplification attacks.
- Fragmentation – Large DNS responses are increasingly common, leading to IP fragmentation and a new set of security concerns.

Downsides

- TCP setup adds 1 RTT.
  - Amortized with reuse and pipelining.
  - Reduced with TCP Fast Open.
- Have to worry about middleboxes that interfere with TCP/53.
- TLS setup adds another 2 RTTs.
  - Reduced with Session Resumption.
- Have to think about certificate validation.
- Have to think about failover.

Proposed: New EDNS0 bit “TO”

a.k.a. STARTTLS for DNS

1. Establish TCP connection.
2. Client sends (dummy) query with TO bit set.

   “Hey, let’s upgrade this connection to TLS!”
3. Server responds with TO bit set.

   “Yeah, I’m down with that!”
4. TLS session negotiation commences.

Performance Improvements

Problem	Solution	Status
TCP setup	Connection reuse	stub—recursive: good recursive—auth: poor
TCP setup	TCP fast open [draft-ietf-tcpm-fastopen-07]	In Linux Requires application change
TLS setup	Session Resumption [RFC 5077]	In GnuTLS
Stop-and-Wait	Pipelining	poor
Head-of-line blocking	Out-of-Order Processing [RFC 5966]	clients: good servers: poor

Implementation Status

- TO bit server-side in Unbound.
  - partial support for OOOP
- TO bit client-side in custom client.
  - supports pipelining and OOOP
- A DNS proxy
  - Accepts UDP/TCP/TLS on server side
  - Outbound UDP to upstream server

Further Information

- draft-hzhwm-start-tls-for-dns-00
- T-DNS: Connection-Oriented DNS to Improve Privacy and Security
  - ftp://ftp.isi.edu/isi-pubs/tr-688abs.htm
- http://www.isi.edu/ant/tdns/index.html

Appendix

Simulated Connection Reuse stub-to-recursive

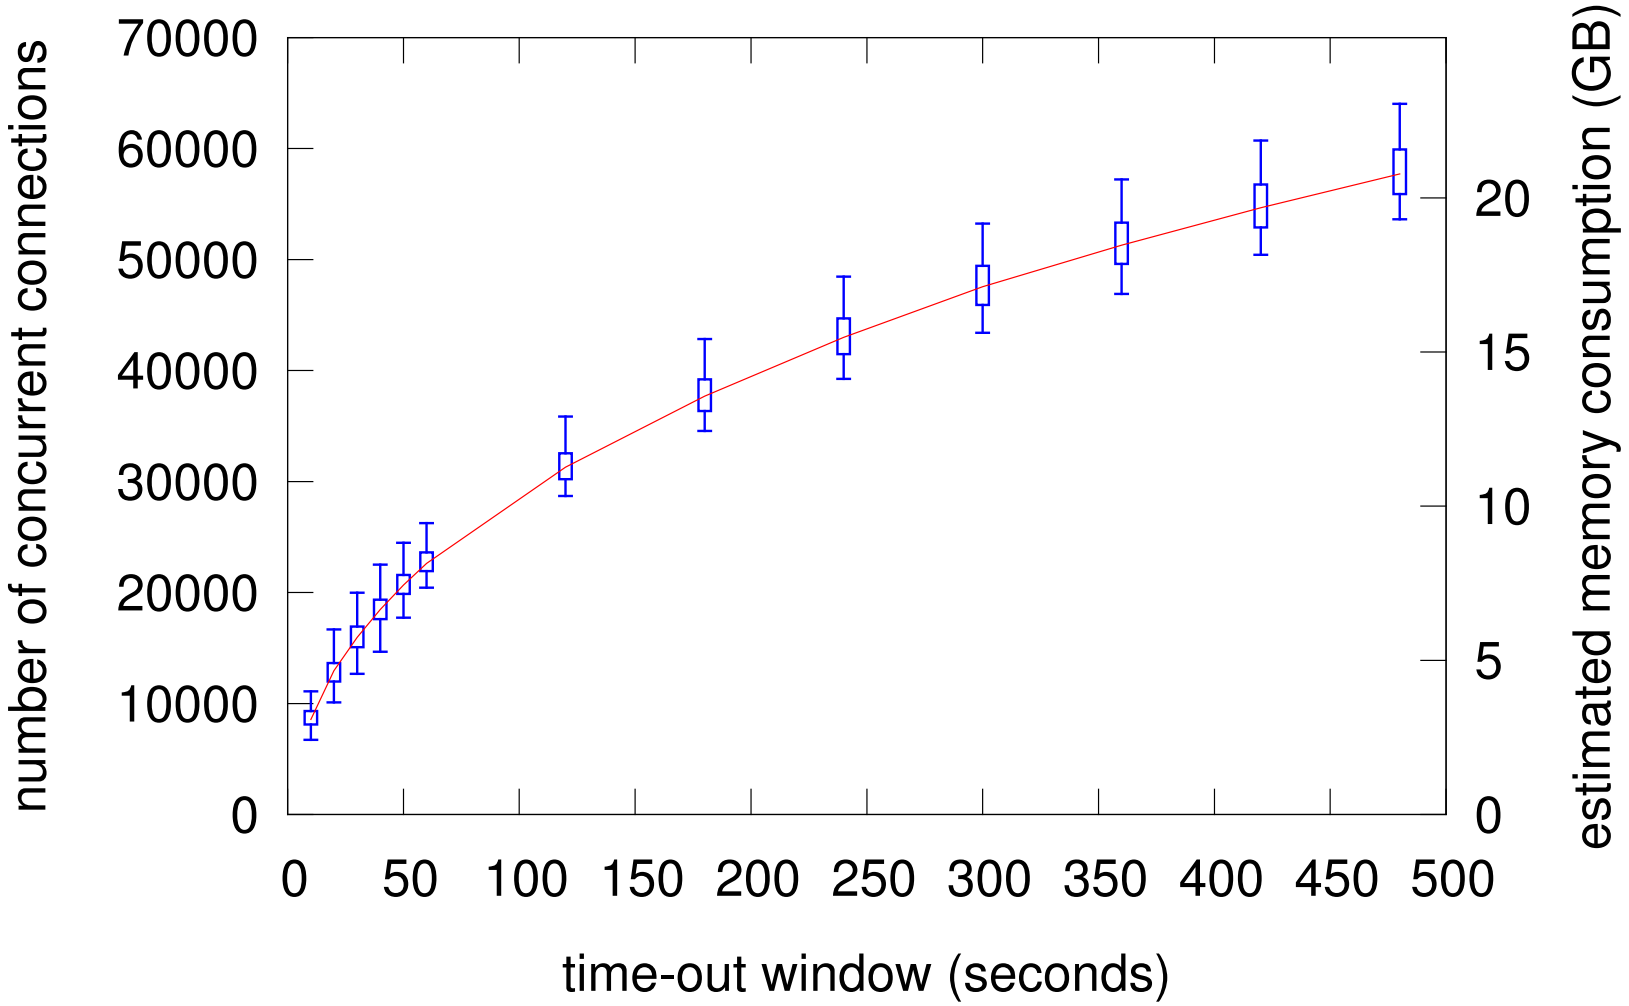

Simulated Connection Reuse recursive-to-authoritative

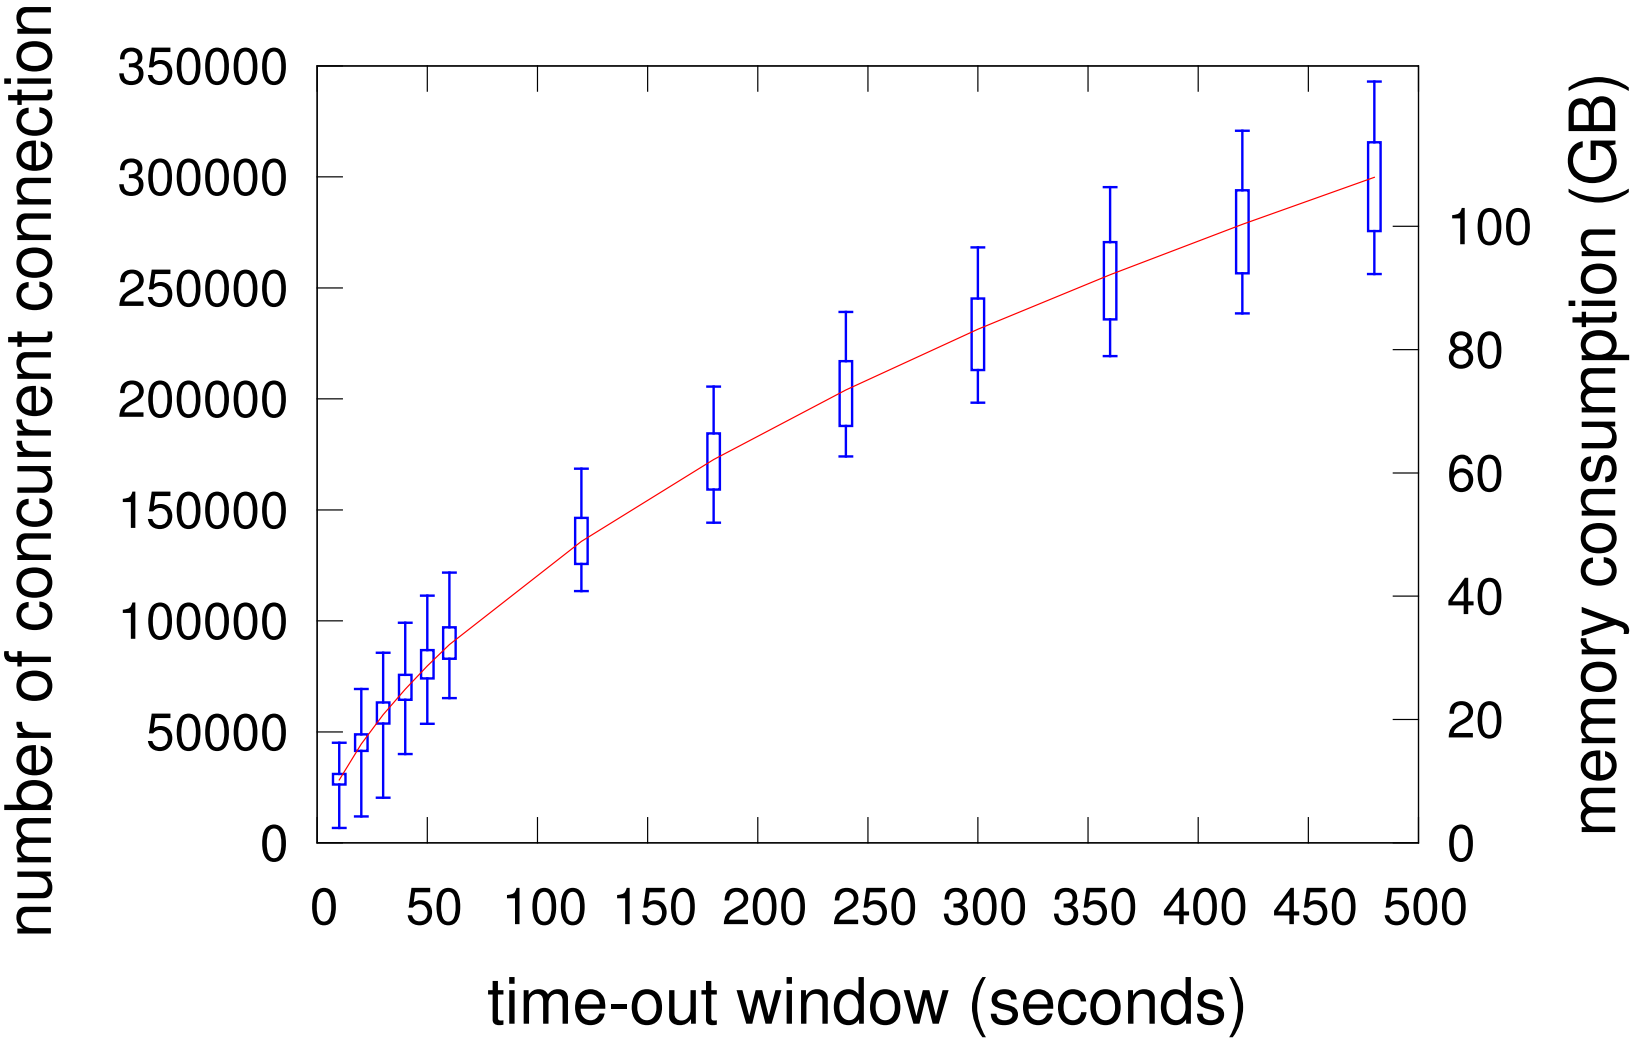

Latency Measurements