

Improving end-to-end delay at the application layer

Baptiste Jonglez, Sinan Birbalta, Martin Heusse

► **To cite this version:**

Baptiste Jonglez, Sinan Birbalta, Martin Heusse. Improving end-to-end delay at the application layer. International Summer School on Latency Control for Internet of Services, Jun 2017, Karlstad, Sweden. <hal-01632191>

HAL Id: hal-01632191

<https://hal.inria.fr/hal-01632191>

Submitted on 9 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improving end-to-end delay at the application layer

Baptiste JONGLEZ, Sinan BIRBALTA, Martin HEUSSE
Univ. Grenoble Alpes, CNRS, Grenoble INP, Inria, LIG, F-38000 Grenoble France



UNDERSTANDING DELAY

A reliable, connected-mode transport layer introduces delay [1]:

1. **Connection setup:** can take several Round Trips (RTT), depending on the protocol.
2. **Retransmission:** essentially transforms packet loss into additional delay.
3. **Head-of-line blocking:** harmful when multiplexing several independent communications into a single stream.

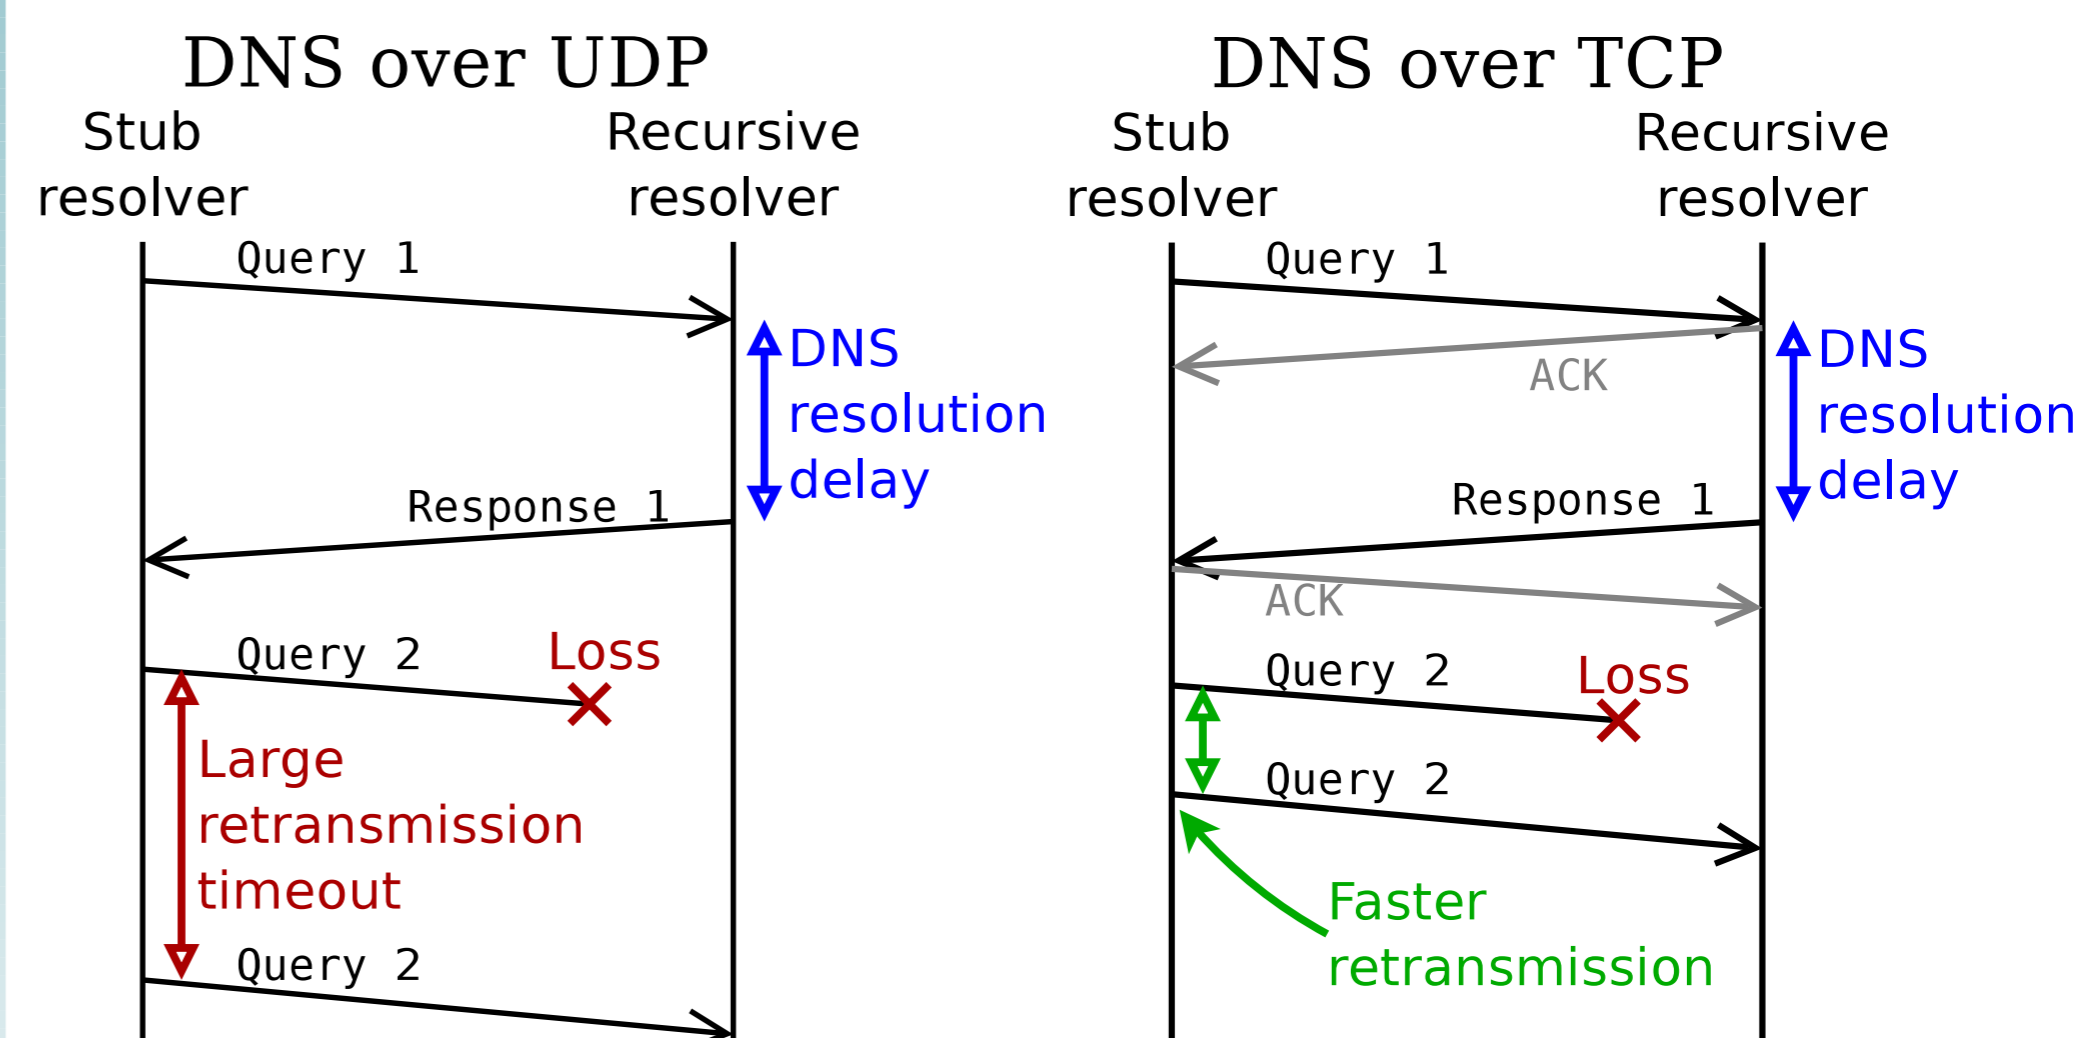
We first focus on improving retransmission delays, especially in the case of *thin streams* [2] or short-lived communication.

KEY IDEA: MEASURE & REUSE

First, we **measure** performance metrics about the network, to improve delay by making better-informed choices (routing, retransmission...). For instance, TCP already estimates the RTT to compute a retransmission timeout.

Then, we propose to **reuse** these performance metrics for several flows. This can be done using proxying or multiplexing, or through application-specific modifications. Reusing performance metrics benefits short-lived connections and *thin streams* [2], for which there is no time to gather performance information.

CONNECTED TRANSPORT FOR DNS



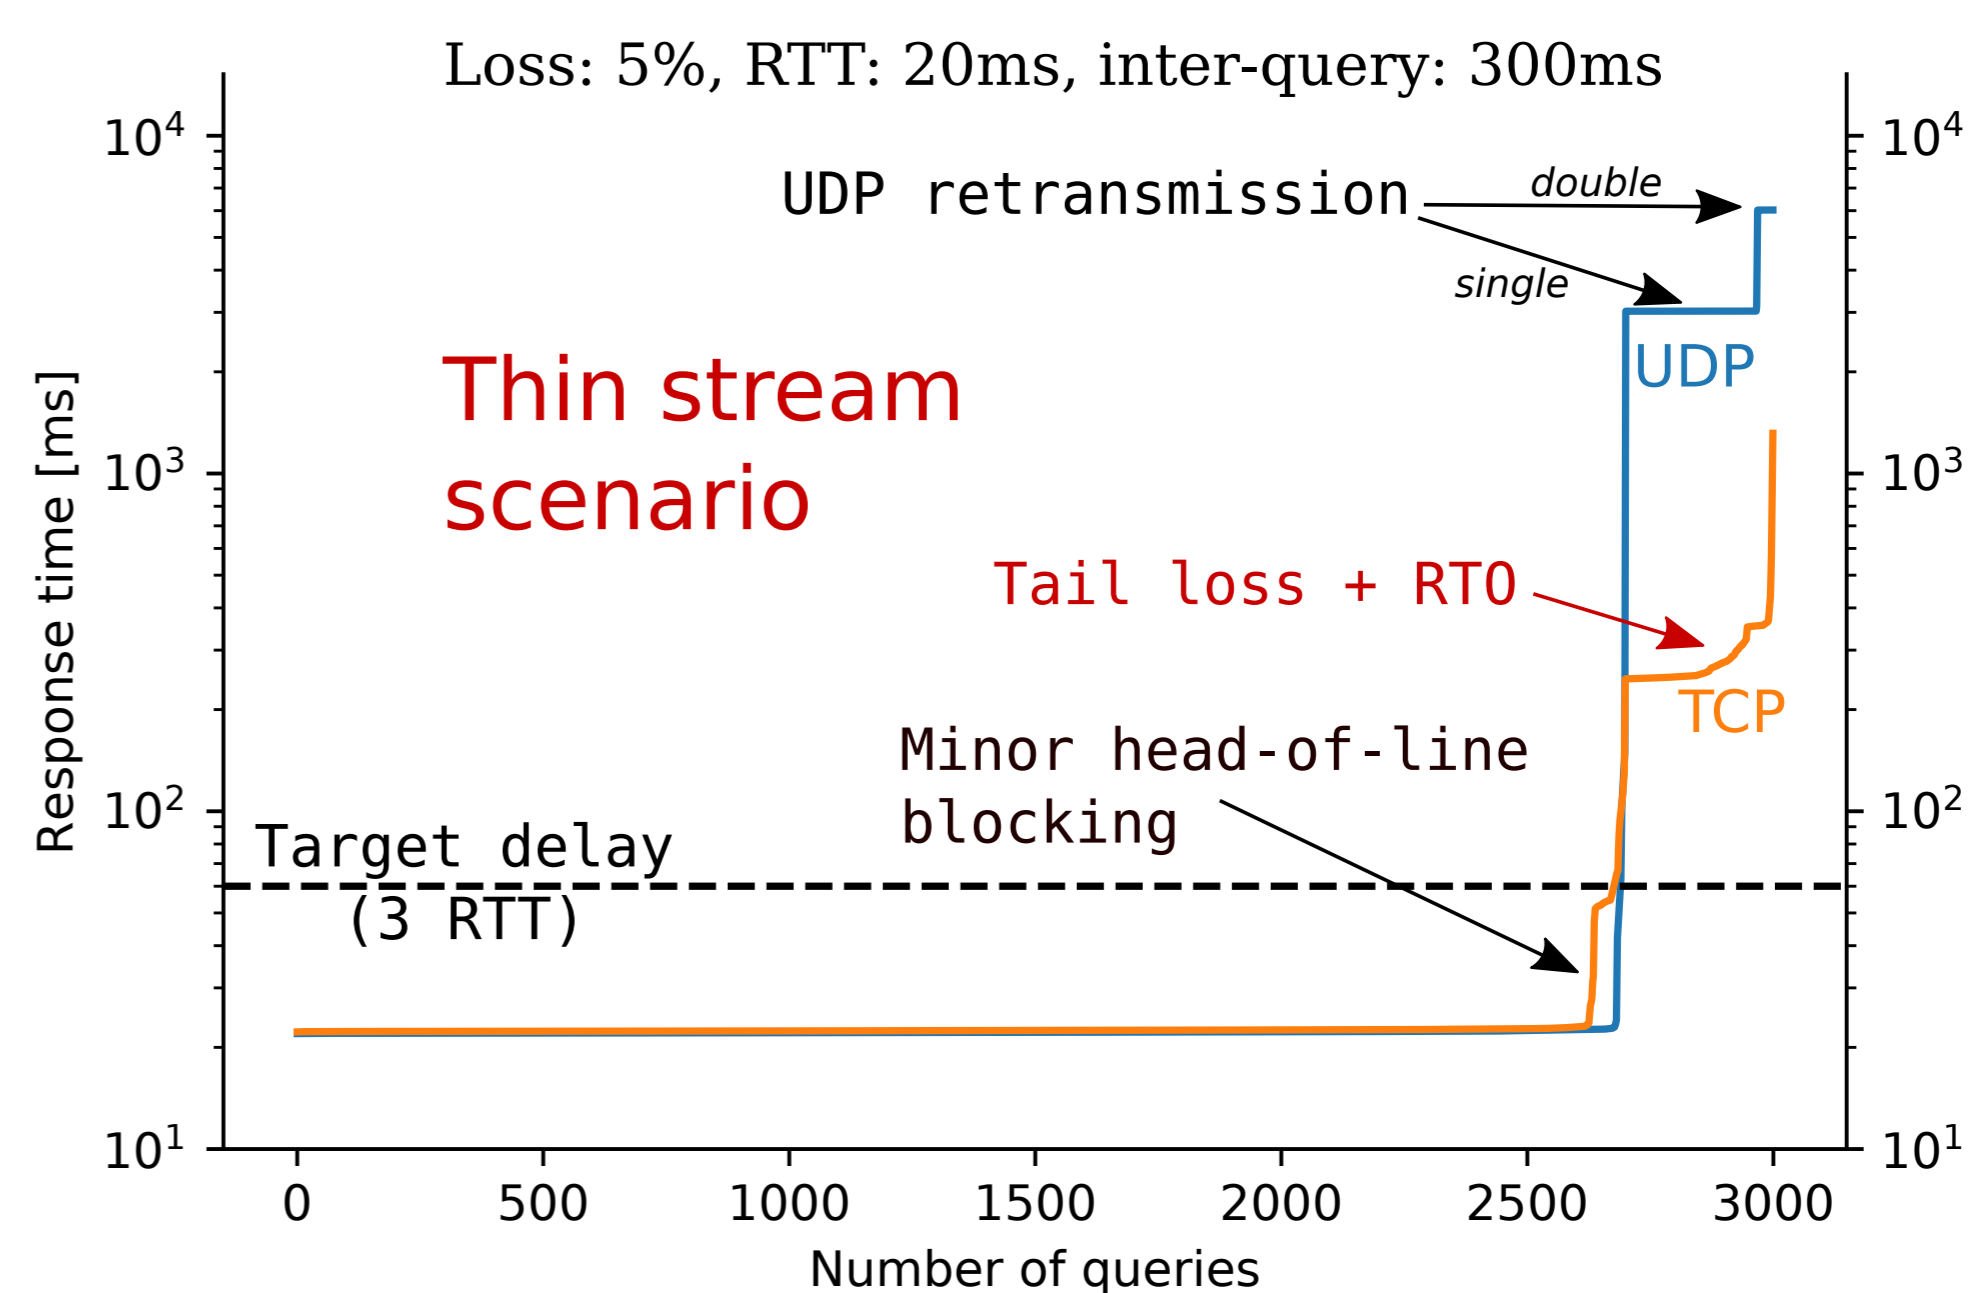
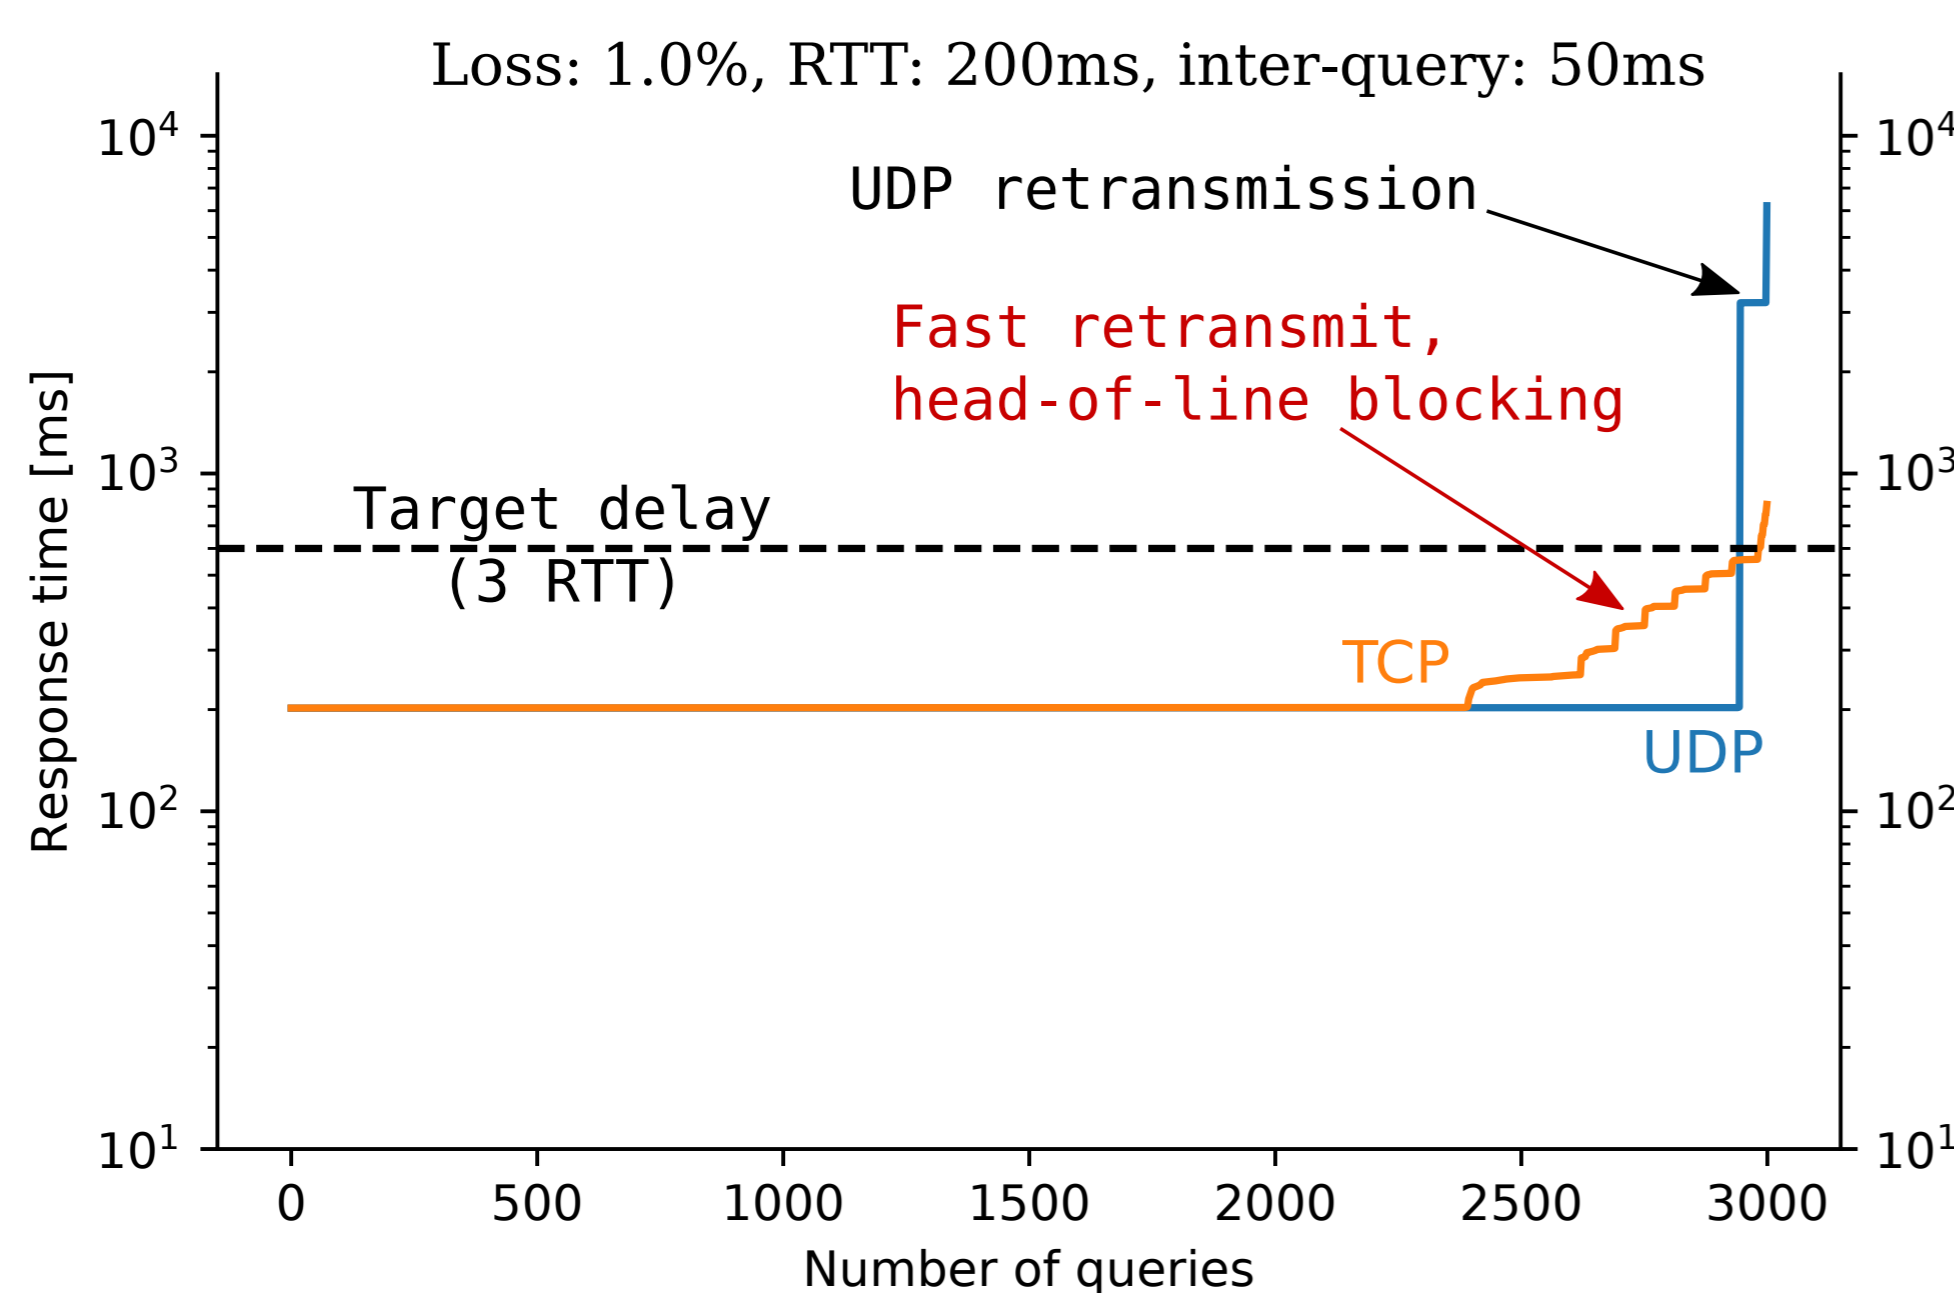
We experiment with **connected transport for DNS** to improve retransmission delay under loss. With UDP, the retransmission timeout is large, to account for the highly variable DNS resolution delay (in blue). With TCP, retransmission at the transport layer can work independently from the application layer. Past RTT measurements are used to adjust the retransmission delay of subsequent requests.

HEAD-OF-LINE BLOCKING

DNS over TCP suffers from head-of-line blocking between DNS requests. We intend to compare end-to-end performance of DNS with TCP, QUIC and SCTP in a future work.

EXPERIMENTAL COMPARISON OF UDP AND TCP FOR DNS

We compare the performance of UDP and TCP on our local testbed. We show the cumulative distribution of response time for DNS queries in two cases: RTT larger than inter-query time, and RTT lower than inter-query time. The UDP retransmission timeout was 3 s.



TCP maintains lower delays than UDP, but may suffer from head-of-line blocking. We also compared several TCP variants, such as Early Retransmit and Tail Loss Probes, with no significant improvement.

FUTURE WORK

Another potential application of the “*measure & reuse*” idea is **multipath transport**. We envision a simpler alternative to Multipath TCP, in which a connection would use a single path at any given time, and migrate between available paths according to long-term performance measurements.

REFERENCES

- [1] B. Briscoe et al. Reducing Internet Latency: A Survey of Techniques and Their Merits. *IEEE Communications Surveys Tutorials*, 18(3):2149–2196, 2016.
- [2] Andreas Petlund. *Improving latency for interactive, thin-stream applications over reliable transport*. PhD thesis, University of Oslo, 2009.