

Need for Low Latency: Media over QUIC

ACM Mile-High Video, May 2023


Zafer Gurel, Tugce Erkilic Civelek and Ali C. Begen

A Quick Look into the “Quick” QUIC Protocol


- Main new features over TCP+TLS+H2
 - Connection establishment latency: 0-RTT (or 1-RTT)
 - Customizable congestion control, improved retransmission machinery
 - Multiplexing w/o HoL blocking
 - Connection migration: Moving between network interfaces without renegotiating the session
- Reliable and prioritized delivery
 - Decoupled retransmissions, congestion control and flow control
 - Stream prioritization management by the sender
- Separate logical streams (for data that can be processed independently) within a physical connection
 - In-order delivery within a stream but no in-order delivery guarantee among streams
 - Partial reliability thru resetting the stream to expire an unack'ed message (if the message is mapped to a single stream)
- Encrypted delivery
 - Multiple connections needed if QoS is needed (network cannot identify the streams)

H2
Inspired by
Google's SPDY

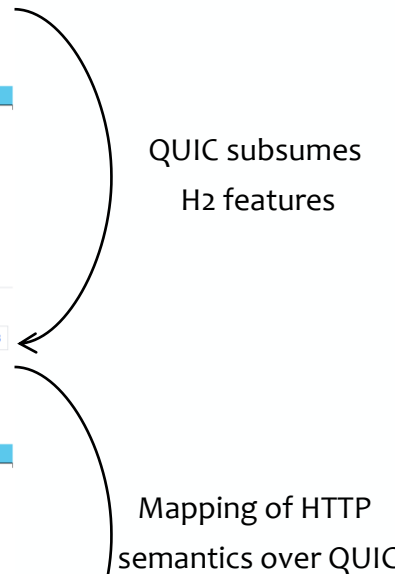

Hypertext Transfer Protocol Version 2 (HTTP/2) RFC 7540



QUIC: A UDP-Based Multiplexed and Secure Transport RFC 9000



HTTP/3 RFC 9114




QUIC


H3

How Quickly QUIC can Replace the Not-So-Quick TCP

TCP



QUIC

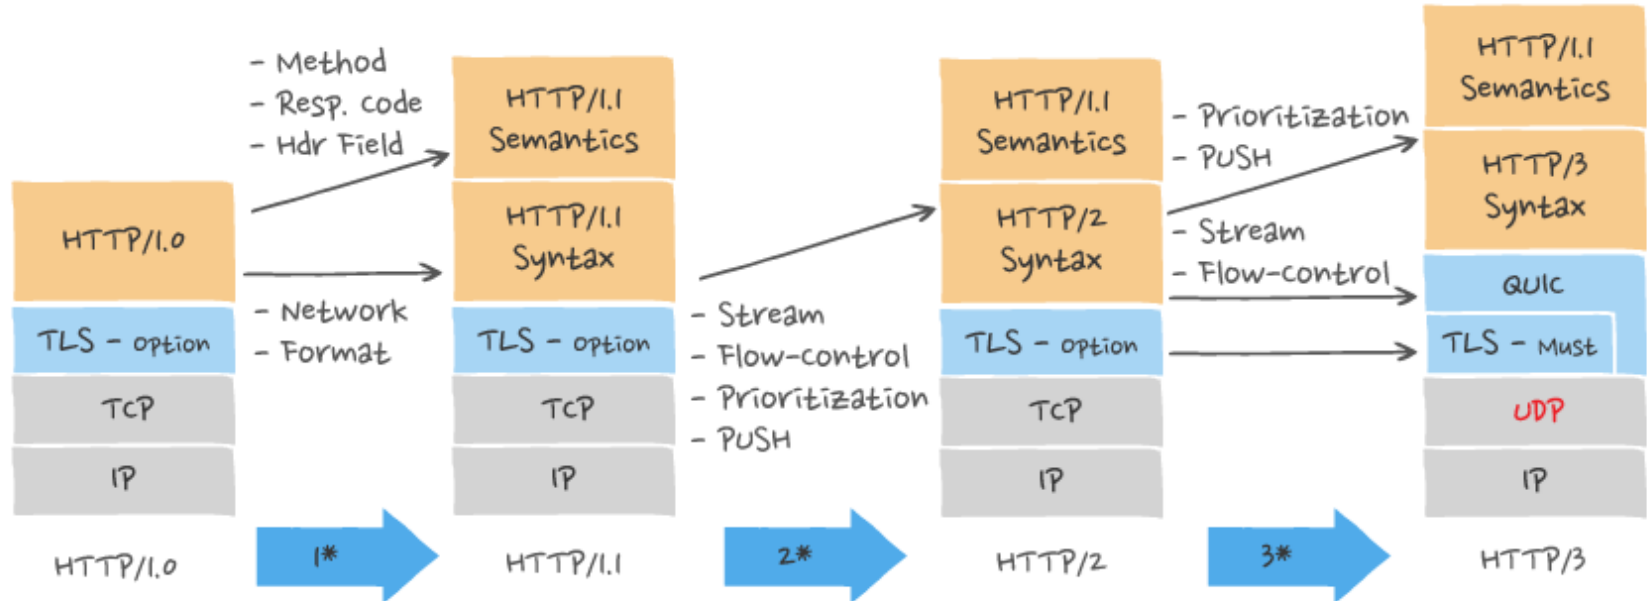


Source: <https://blog.cloudflare.com/http-3-vs-http-2/>

ACM MHV - May 2023

HTTP/1.1 vs. H2 vs. H3

HTTP protocol stack transition and comparison



Source: <https://ably.com/topic/http-2-vs-http-3>

Earlier Research Showed that

- For timely delivery, QUIC may perform better than TCP in congested environments
 - We still need a custom application-layer protocol to reap all the benefits QUIC provides at the transport layer
- Existing adaptive streaming methods
 - have been highly tuned for HTTP/1.1 and 2 running on top of TCP
 - do not give remarkably better results with H3 running over QUIC

Unless the streaming application is aware of QUIC's unique features,
the improvements will be limited

Sample reading:

Quickly starting media streams using QUIC – ACM Packet Video'18

Evaluating QUIC performance over web, cloud storage and video workloads – IEEE Trans. NSM 2021

Take the red pill for H3 and see how deep the rabbit hole goes – ACM MHV'22

High Latency

Typical Latency

Low Latency

Very Low Latency

near-Real Time Latency

DASH/HLS
10s segments

DASH/HLS
6s segments

DASH/HLS
2s segments

Cable, IPTV, satellite,
over-the-air broadcast

DASH/HLS
1s segments

DASH/HLS
chunked segments

Live sports, game
streaming, eSports

Social media, messaging

Online gambling,
betting, auctioning

Videoconferencing,
VoIP

Cloud gaming

Surgical robots

45
seconds

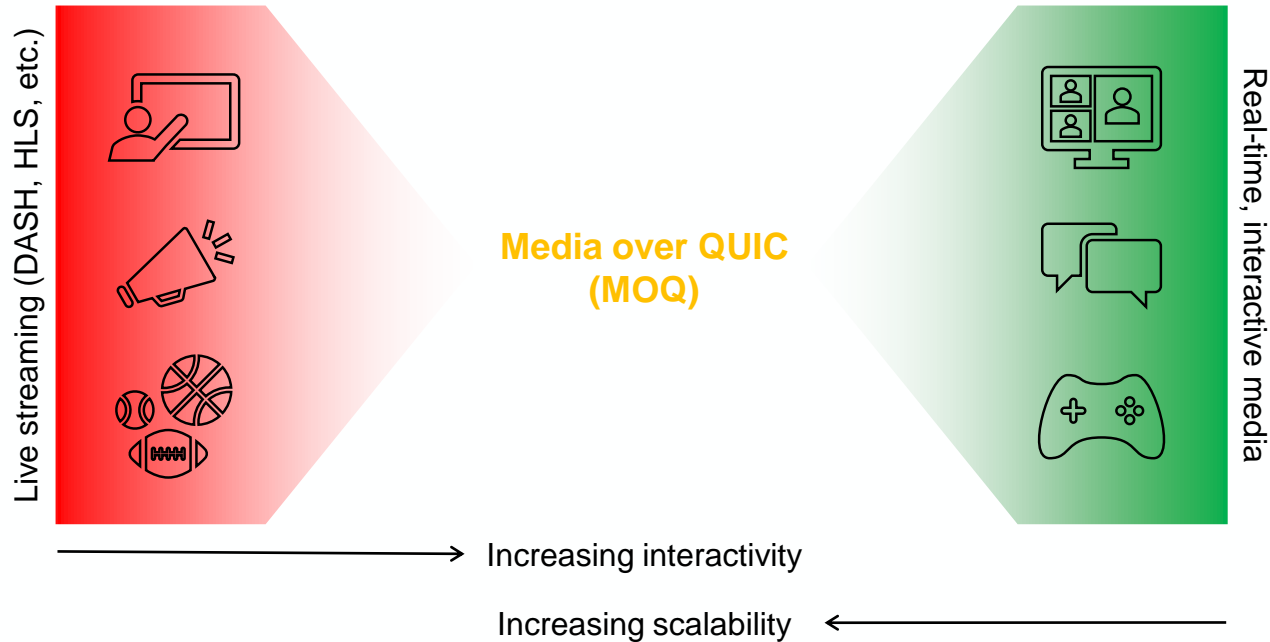
10
seconds

1
second

1
millisecond

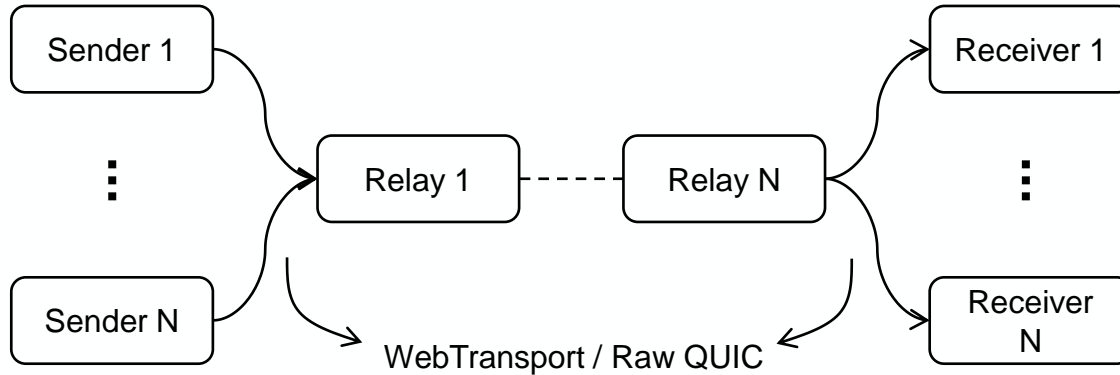
The New Media over QUIC (MOQ) Working Group

Using QUIC for media transmission in one-to-one, one-to-many and many-to-one applications requiring interactivity




Simplified Diagram for E2E Deployment of MOQ

Pub/Sub model: push as opposed to pull



Motivation and Goals

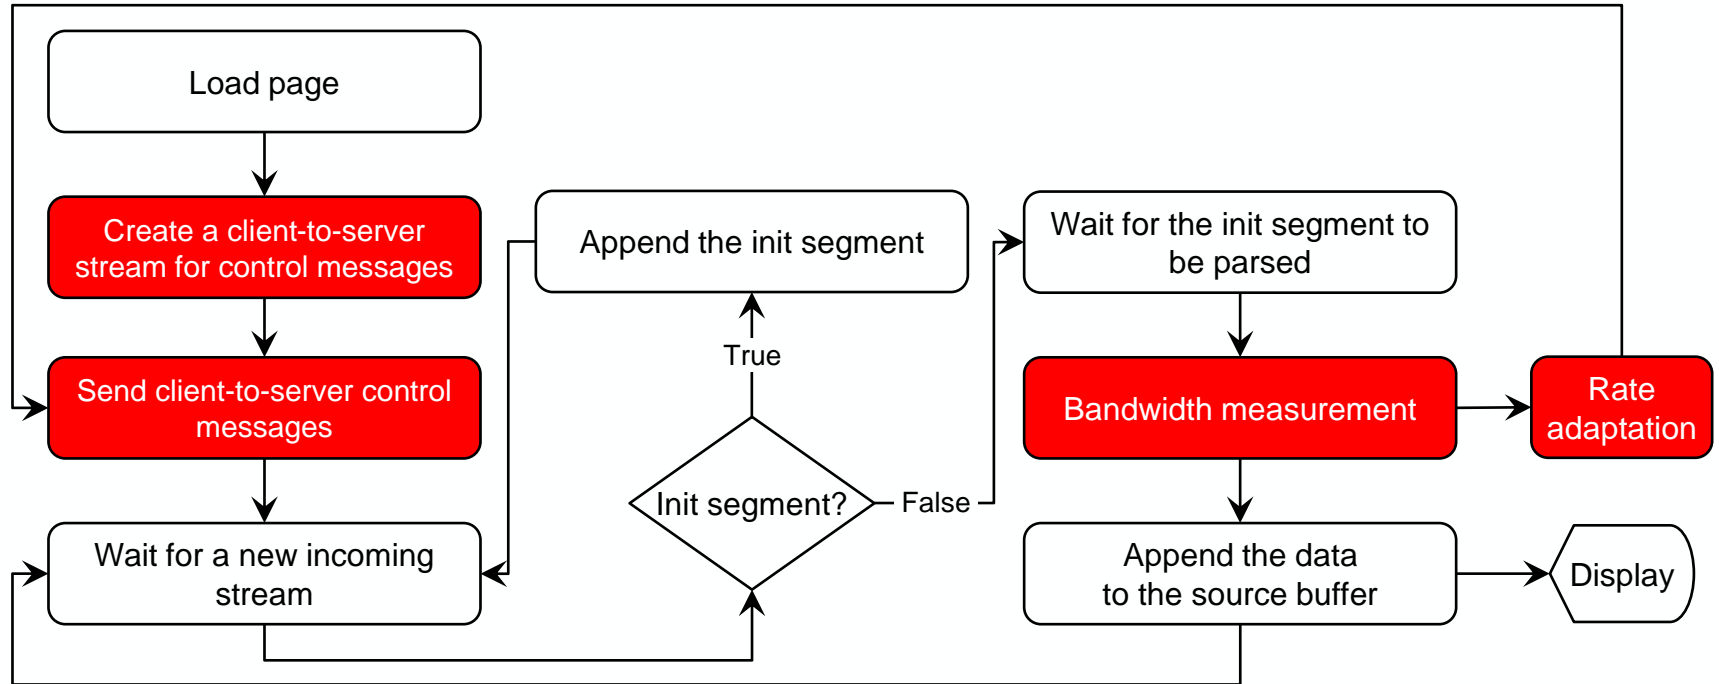
Do a gap analysis of Warp and create an open-source testbed



- 1 Server-to-client informational messages
Wallclock time synchronization (in progress)
Client-to-server control messages
- 2 Passive and active bandwidth measurements
- 3 Enhanced user interface and logging


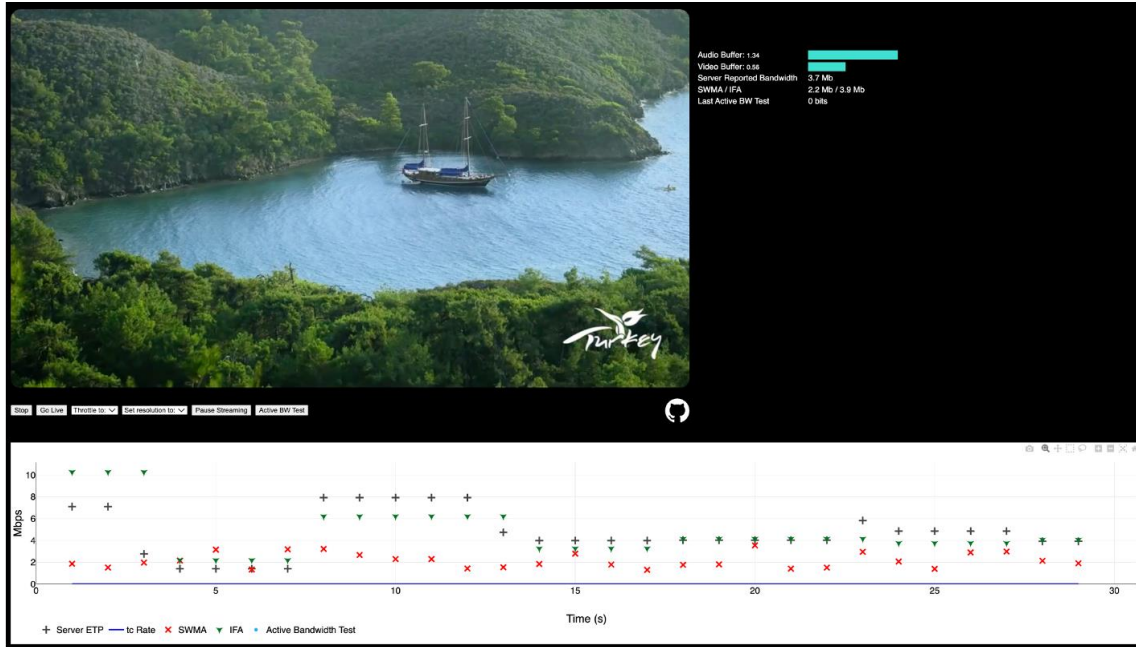
Warp Client-Side Flowchart

Our additions and enhancements are shown in **red** boxes



Public Demo

<https://moq.streaming.university>



Ongoing Work

- Testing with servers at distinct geographies
- Replacing quic-go with quiche for BBR support
- Using WebCodecs API for low-level access to media frames
- Trying different strategies for stream prioritization
 - Varying the number of streams per segment and their priorities
- Labeling and storing log data in the cloud
- Comparative analysis of the related implementations (Meta, Cisco, others?)



ALWAYS HOPE, THERE IS

Grad students wanted!

One hour
here is seven
years on
earth

Great, I will
do my PhD
here



AMAAAS*

- Reach out to any of us for questions
 - zafer.gurel@ozu.edu.tr
 - tugce.civelek@ozu.edu.tr
 - ali.begen@ozyegin.edu.tr
- Acknowledgments
 - Luke Curley for his help/discussions
 - AWS for the EC2 support to host our demo application

This deck (and many others) are posted at
<https://ali.begen.net>

* AMAAAS: Ask me almost anything about streaming