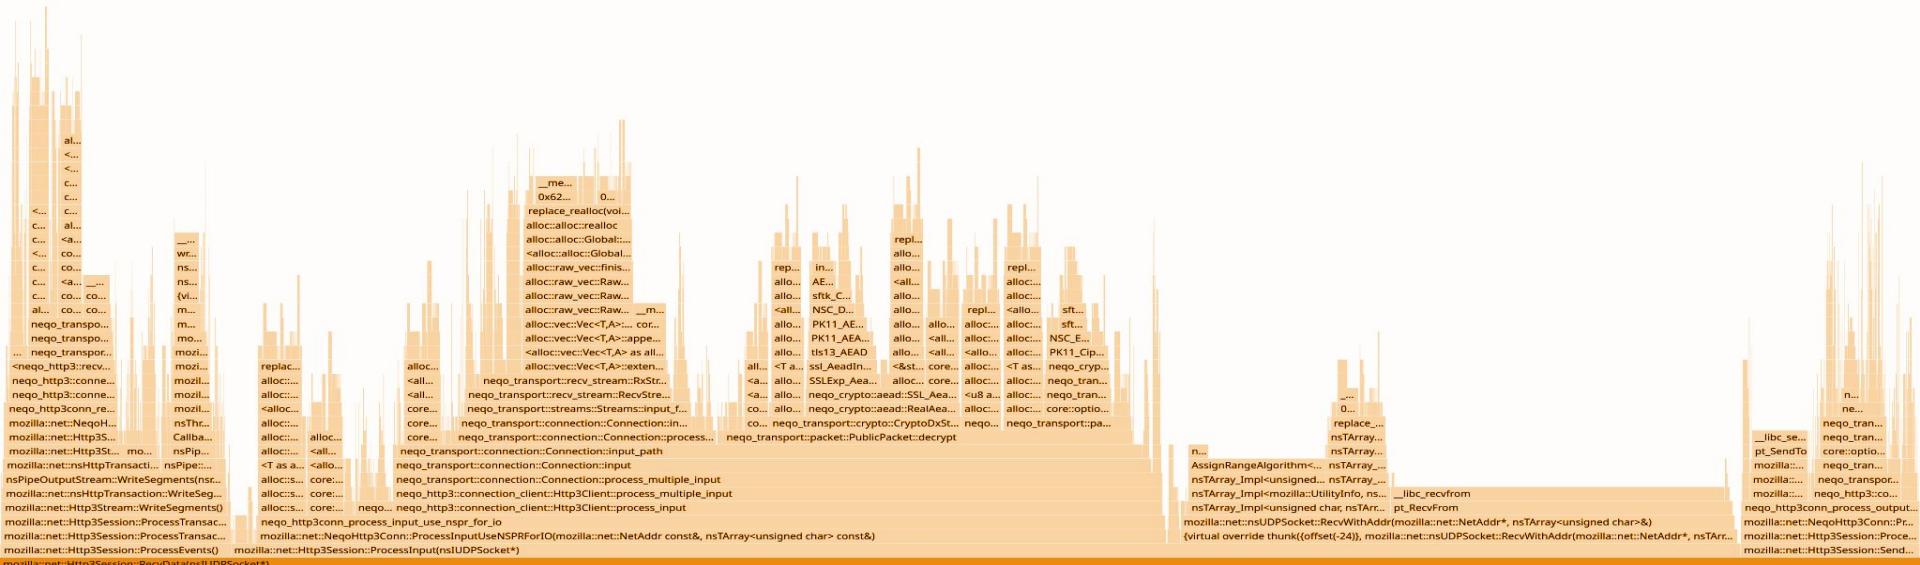


Fast UDP makes QUIC quicker

optimizing Firefox's HTTP3 IO stack

MIR^3 2025 - Max Inden

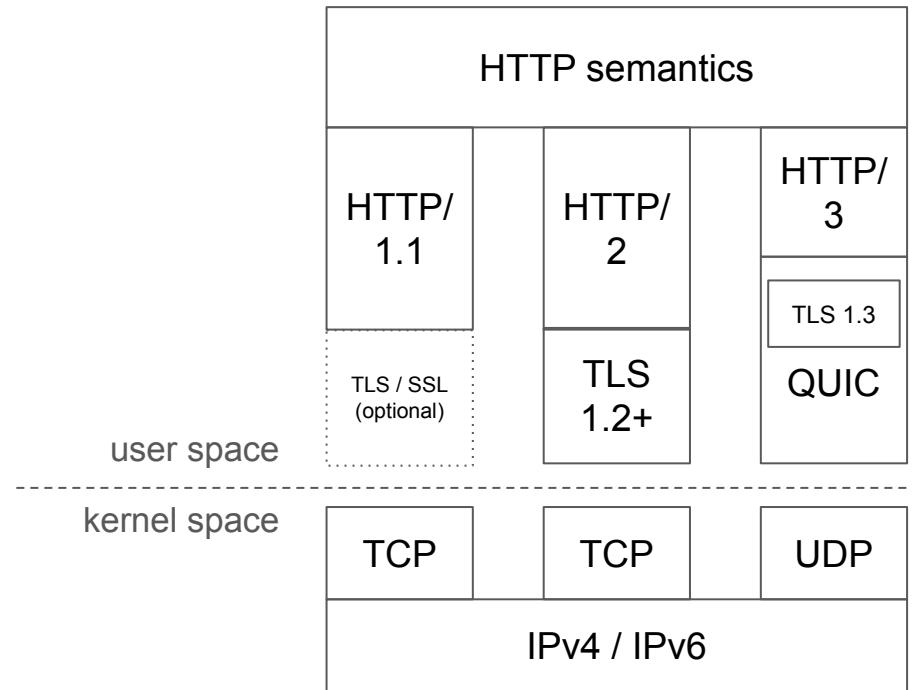


Max Inden

- Software engineer at Mozilla
- Working on HTTP3 and QUIC in Firefox
- mail@max-inden.de
- @mxinden
- Previously p2p, Kubernetes and Prometheus



QUIC runs on top of UDP

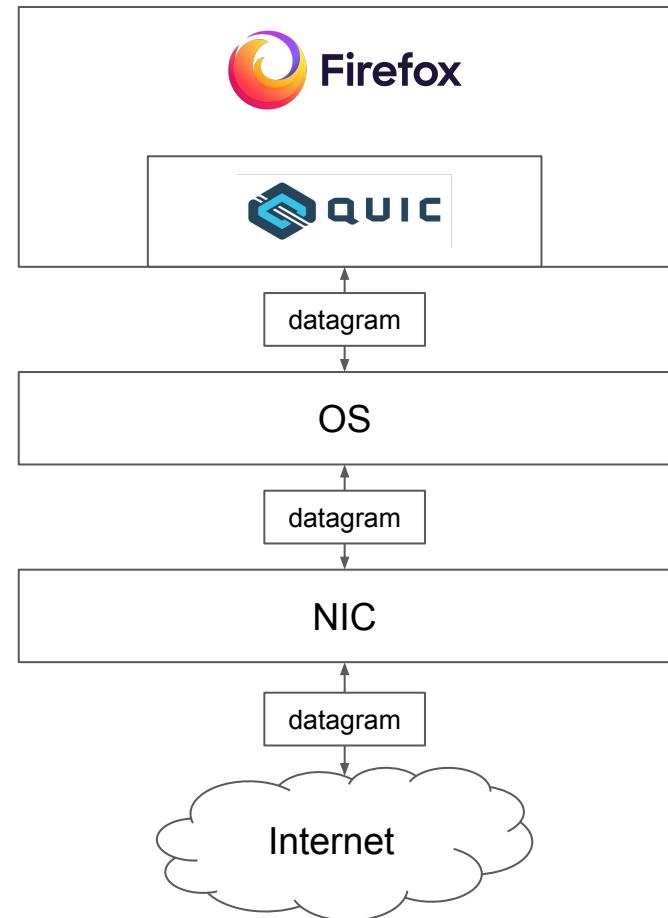


QUIC in userspace

Unsurprisingly an un-optimized userspace transport protocol on top of UDP is not as fast as heavily optimized TCP in kernel space + NICs.

Userspace QUIC might do:

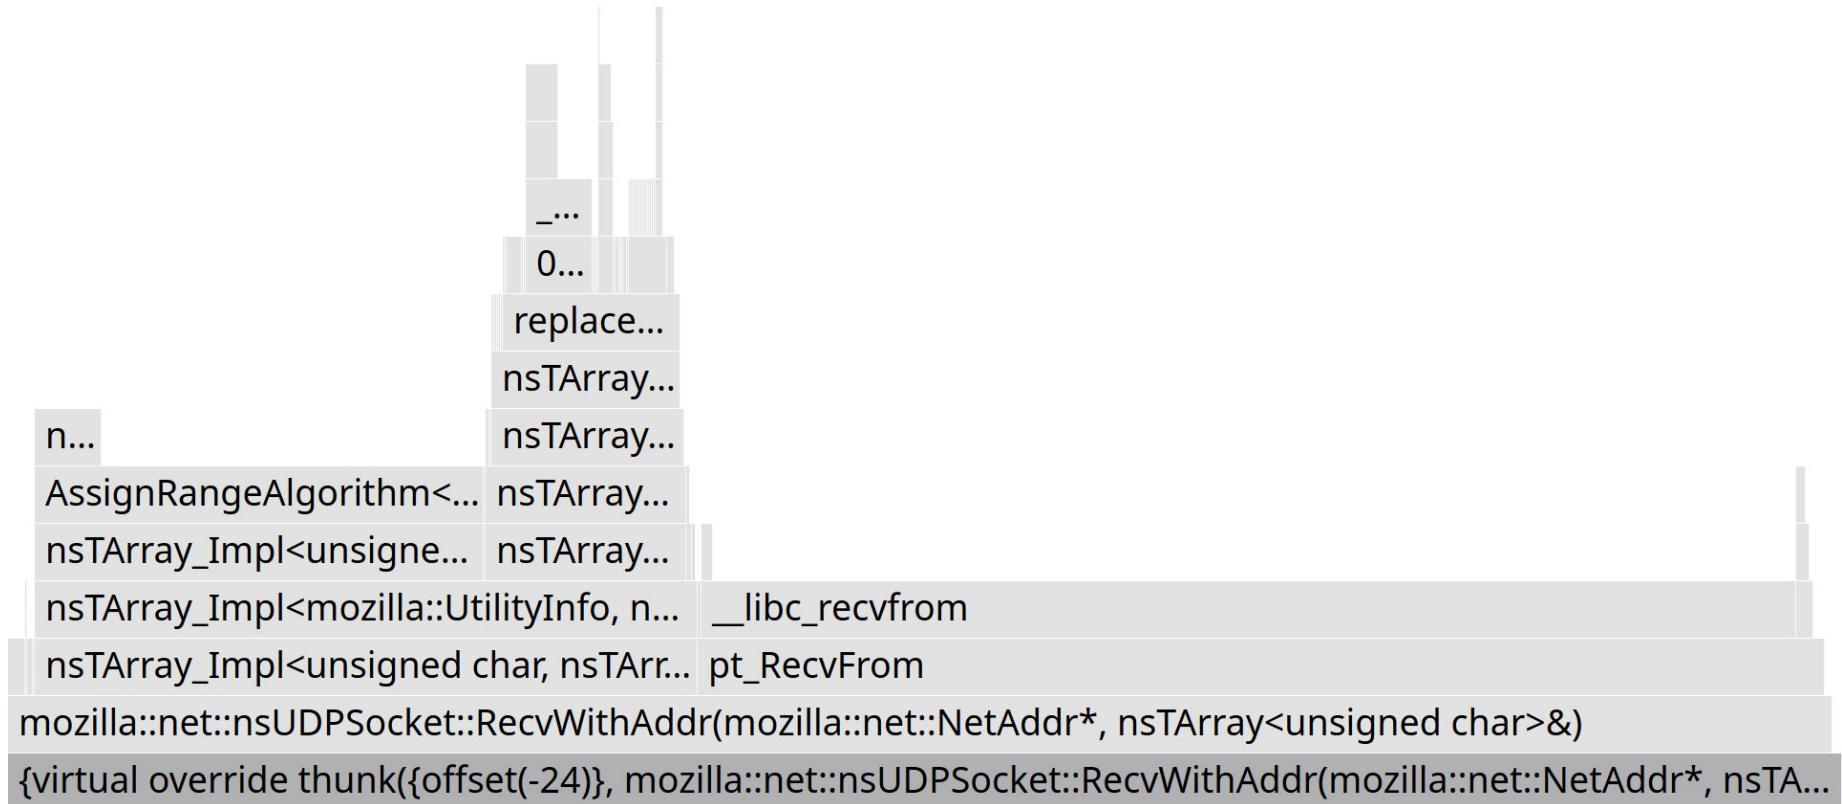
- one syscall per UDP datagram
- <= 1500 byte MTU for Internet traffic
- ACK of up to every second packet
- de Bruijn et al. “3.5x the CPU cycles per byte”



Un-optimized Firefox QUIC UDP IO

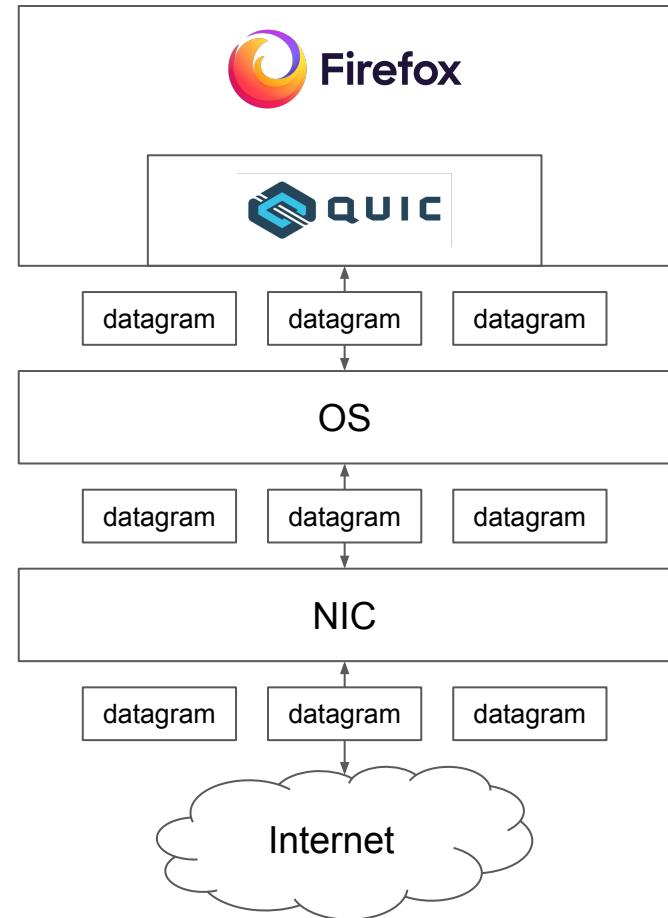


Un-optimized Firefox QUIC UDP IO



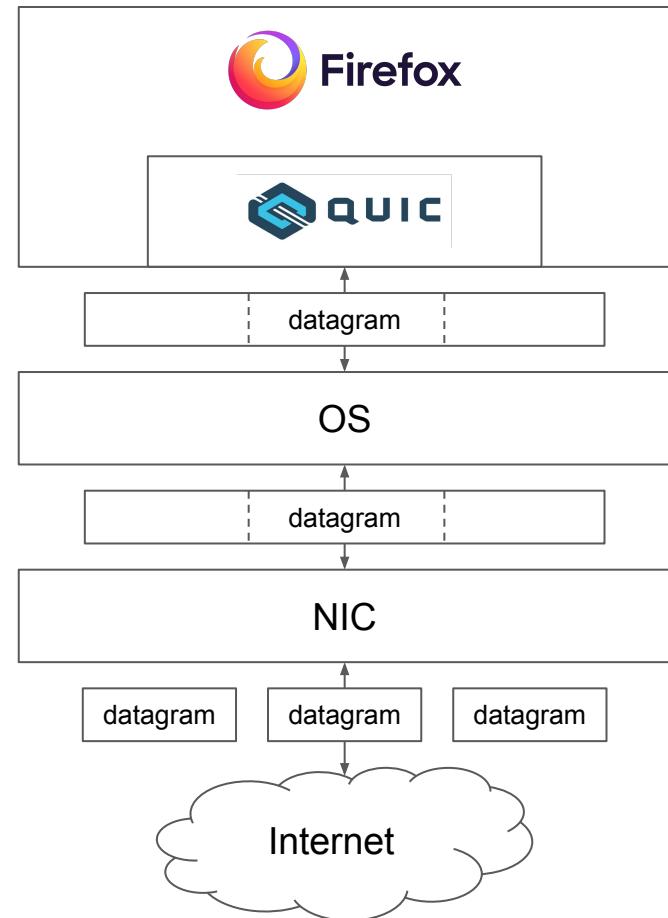
Multi-message syscalls

- MacOS sendmsg_x and recvmsg_x
- Linux's sendmmsg and recvmmsg
- 11% performance improvement on loopback transfer



Segmentation offloading

- Linux's GSO/GRO
- Windows' USO/URO
- measurements on Firefox Nightly Linux GRO
 - 75th of read syscalls read [2 or more packets](#), 95th read 10 or more packets.
 - 75th of read syscalls read [2.4 KiB](#) total, 95th read 12 KiB total.
- close to the 4 Gbit/s on loopback benchmark



Windows

- WSARecvMsg and WSASendMsg
- > fosstodon.org doesn't load with network.http.http3.use_nspr_for_io=false on ARM64 Windows
- URO disabled for now
 - On some drivers no segment size
- USO disabled for now
 - High packet loss
 - Crashes some drivers

MacOS

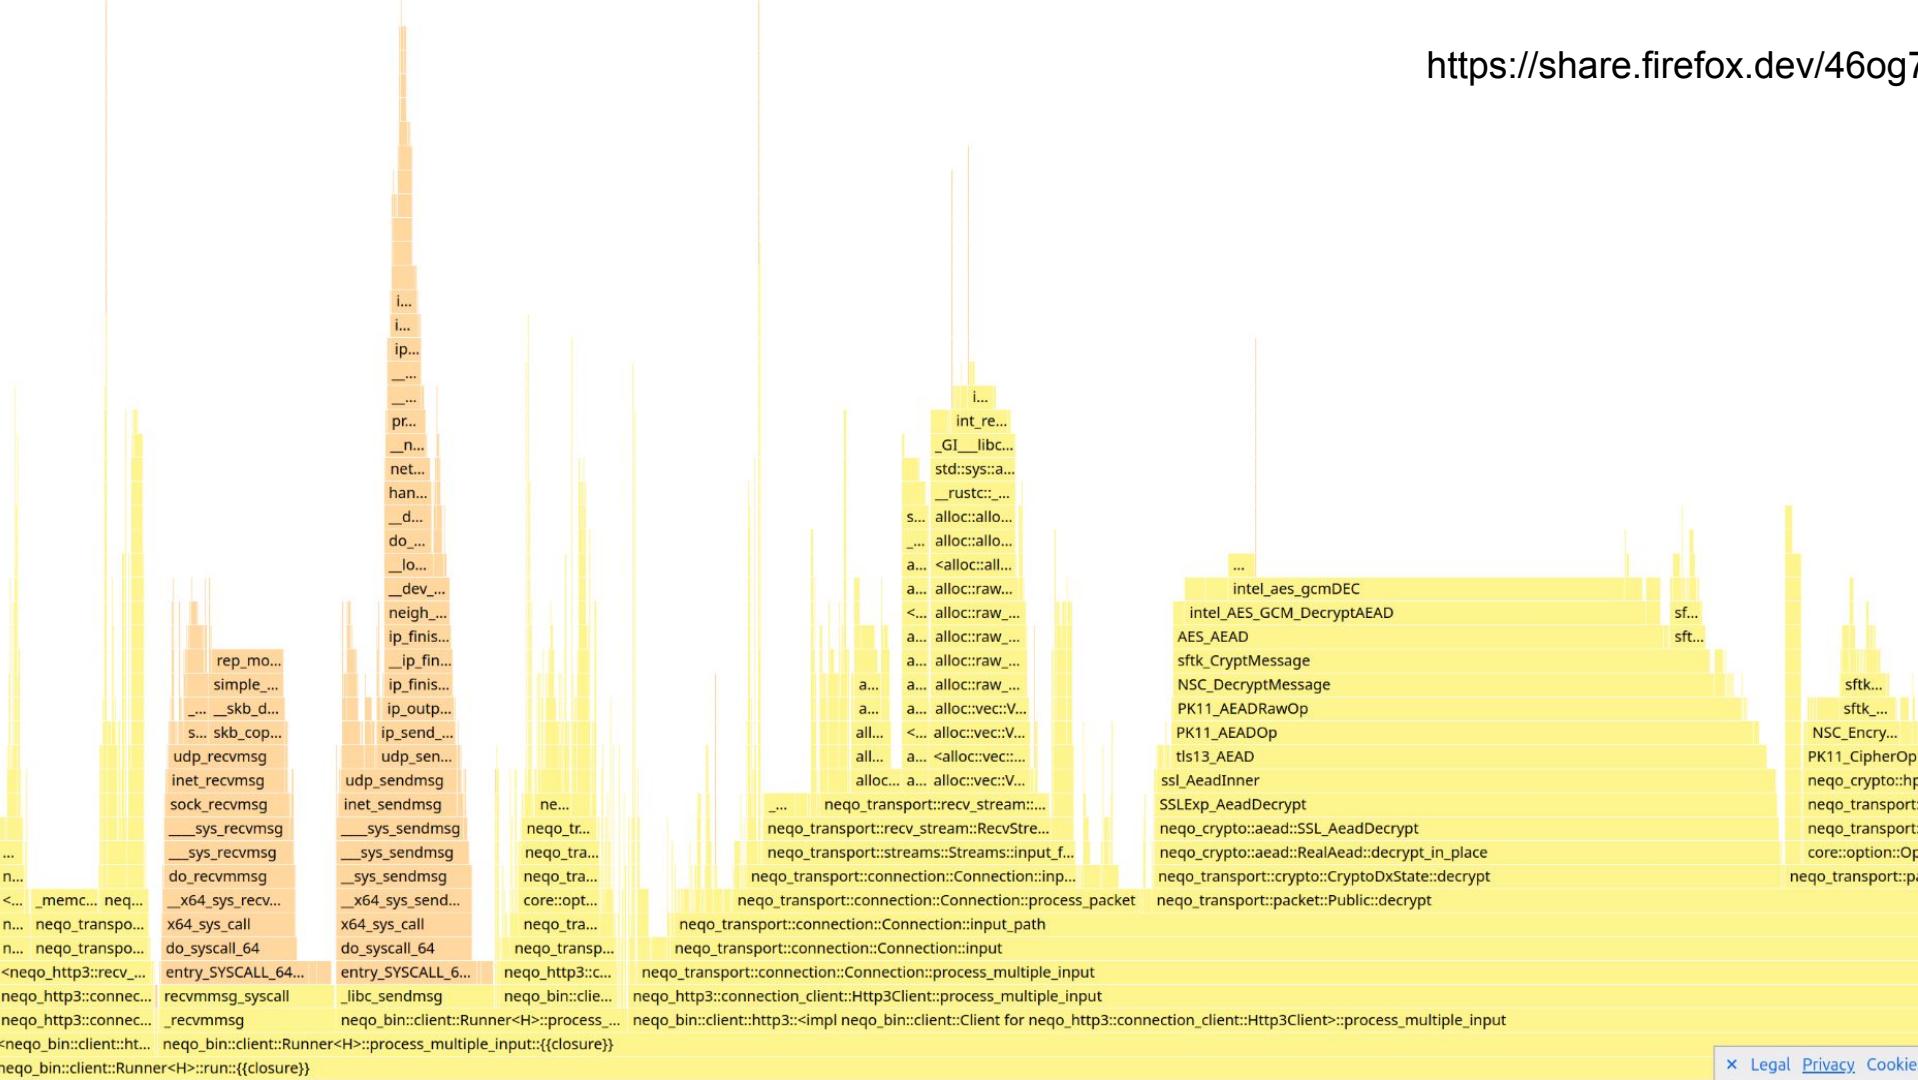
- sendmsg and recvmsg
- sendmsg_x recvmsg_x
- Apple network framework

Linux

- sendmsg and recvmsg
- No mmsg system calls
 - One socket per connection
 - Segmentation offloading superior to mmsg
- GSO (max 10 datagrams) and GRO (64k buffer)
-

Android

- sendmsg and recvmsg
- No mmsg system calls
 - One socket per connection
 - Segmentation offloading superior to mmsg
- GSO (max 10 datagrams) and GRO (64k buffer)
- But Android is not Linux



ECN

- Mark and report ECN
- Groundwork for L4S
- ~50% of paths ECN capable
- 75th percentile of QUIC connections see [>= 0.6% CE](#) marks on receive path.

Additional wins

- QUIC UDP IO in Rust using [quinn-udp](#)
- optimized memory management
 - we use a single long-lived 64k receive buffer for all QUIC connections
 - soundness check at compile time via Rust's borrow checker
 - Inplace en-/decryption
 - [significant reduction](#) in CPU time on large transfers

Next steps

- USO and URO on Windows
- Rollout on Android

max-inden.de/post/fast-udp-io-in-firefox/

