




QUIC and HTTP/3: Too big to fail!?

Robin Marx - [@programmingart](#)
PhD researcher – Hasselt University



UHASSELT

EDM

<https://quic.edm.uhasselt.be>
Symposium on the Future of HTTP – March 2019

QUIC and HTTP/3 are going to change the world!



Lucas Pardue @SimmerVigor · Mar 13



Replying to @alagoutte

In the next 10 years:

HTTP will go to a yearly release cycle. So we will have HTTP/2019 through to HTTP/2029.

QUIC wil replace everything, even payment systems and **5G**.



2



1



QUIC and HTTP/3 **might** change the world!

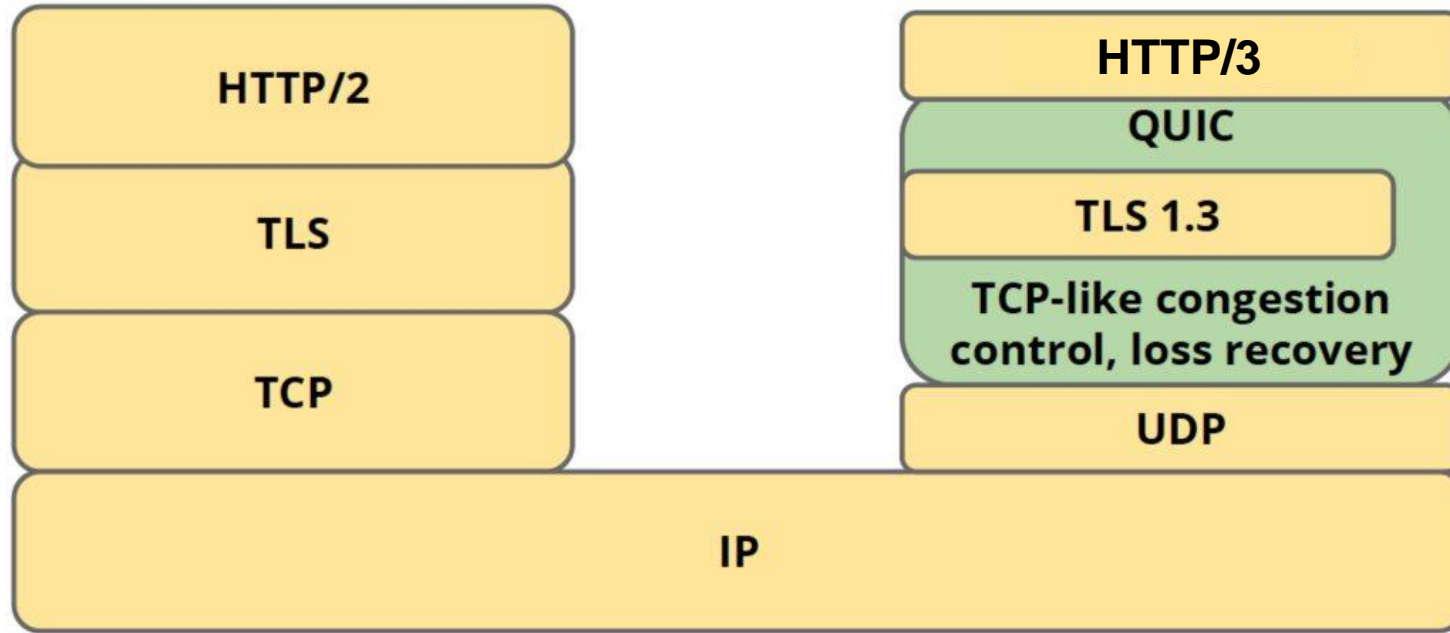


No one will need more than 637Kb
of memory for a personal computer

— *Bill Gates* —


AZ QUOTES

QUIC is special



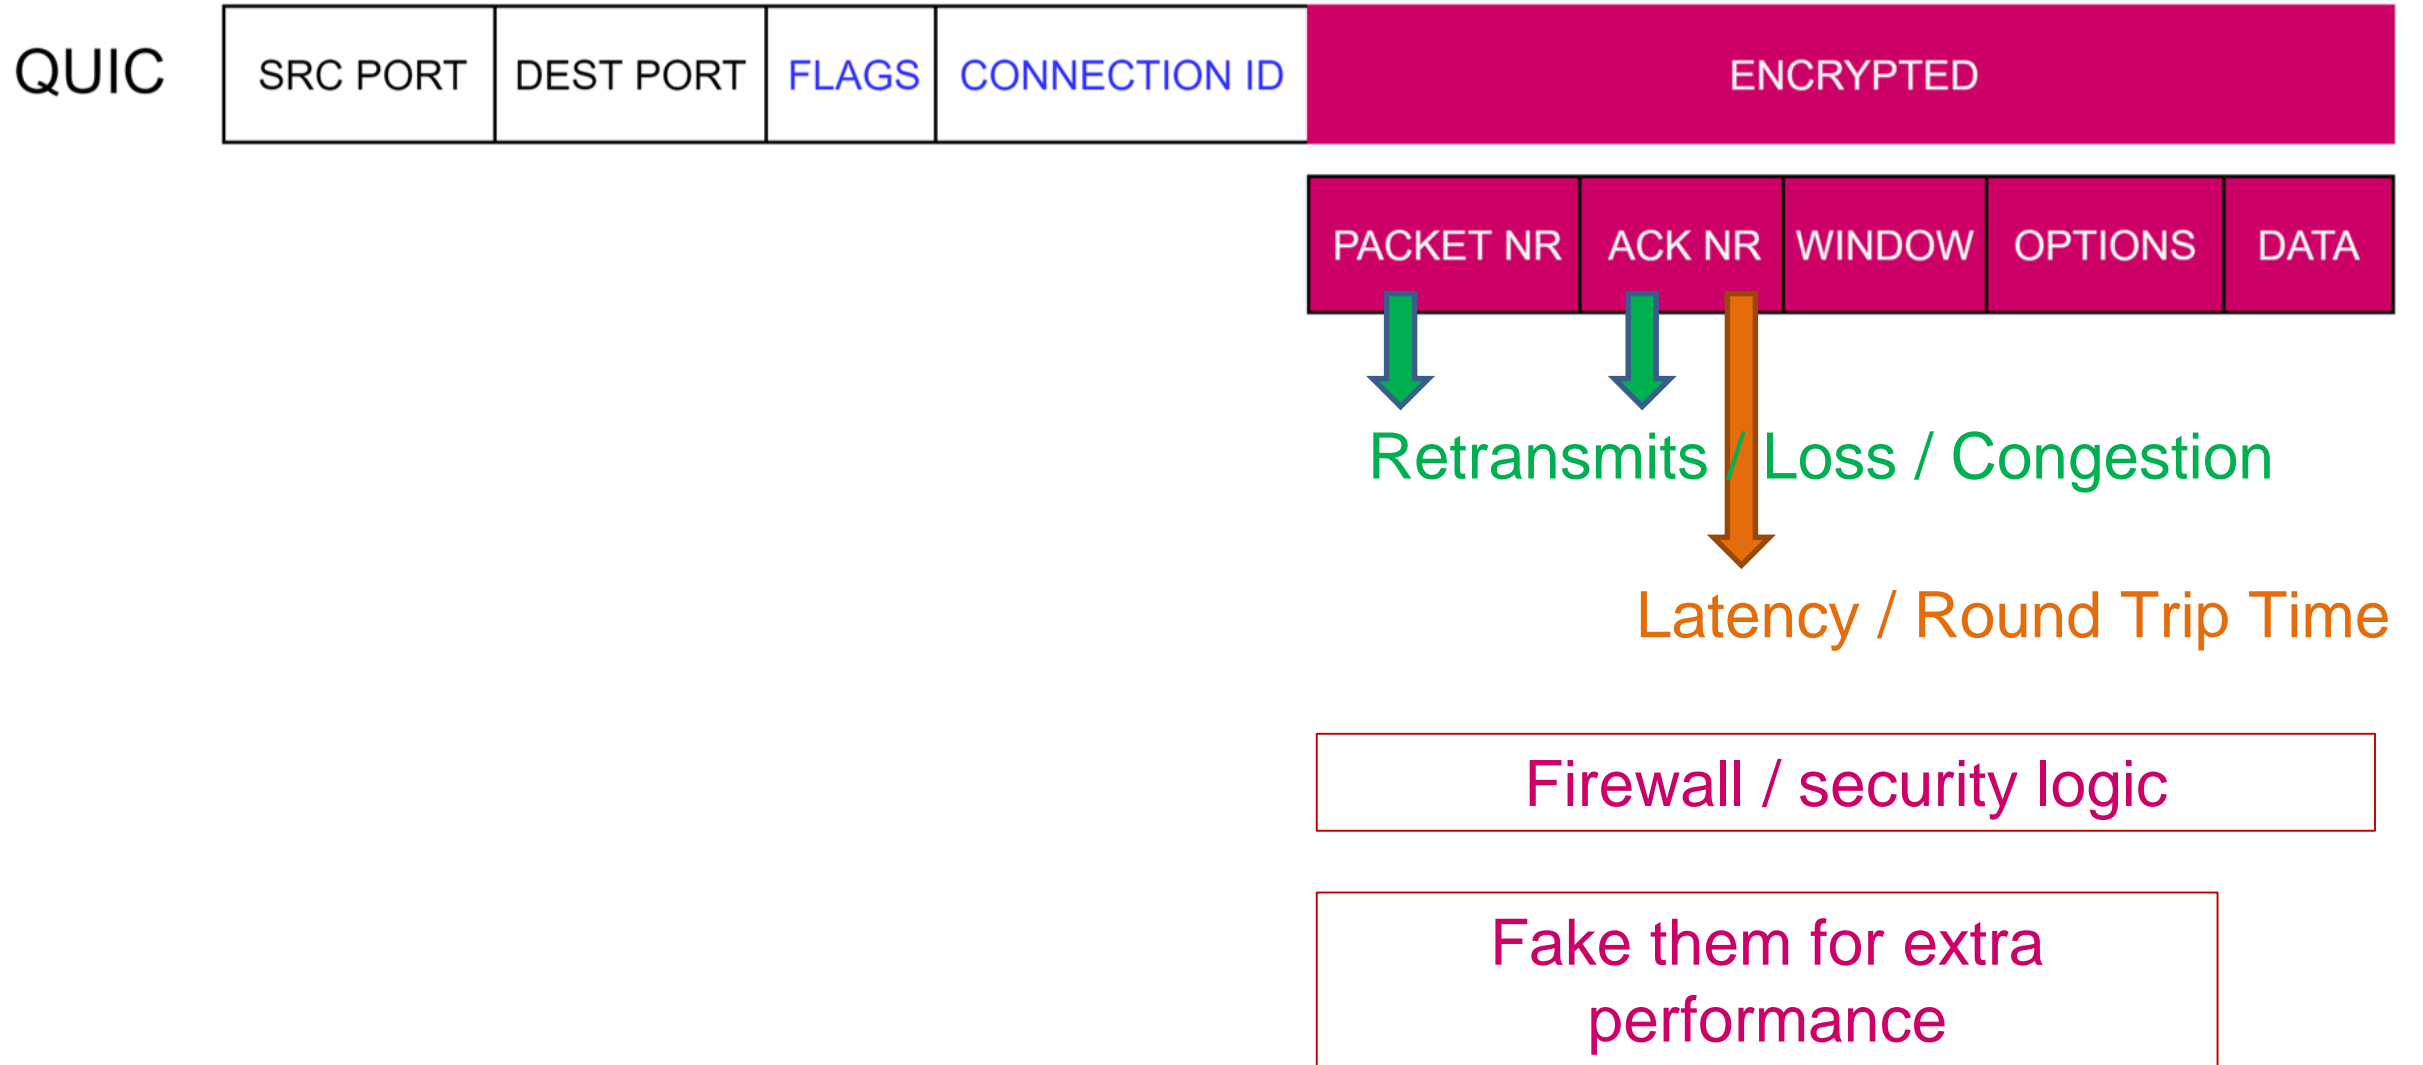
Re-implement:

- Reliability
- Ordering
- Congestion Control
- Flow Control
- ...

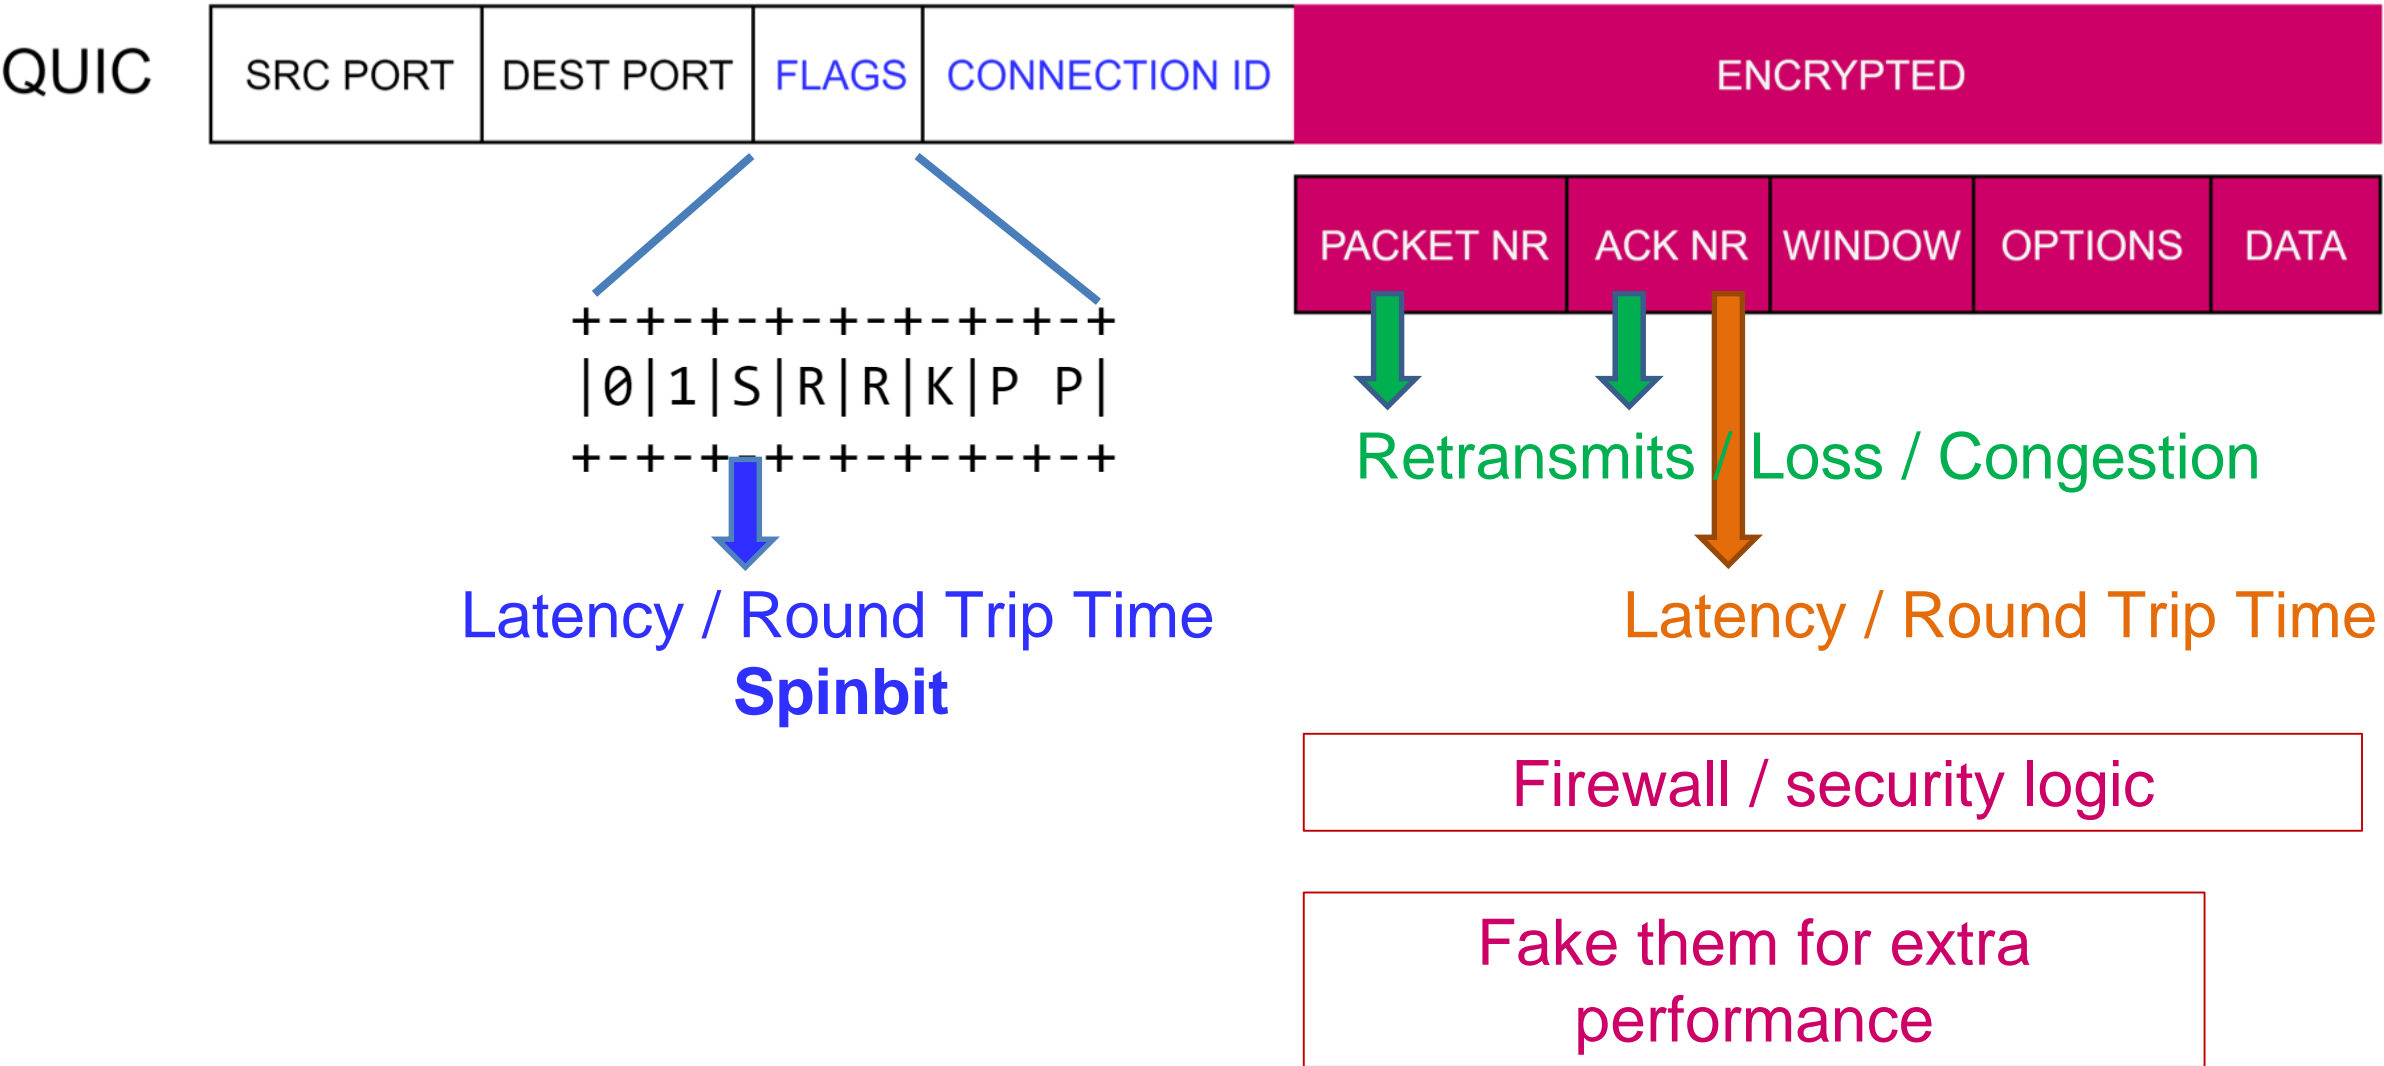


Middlebox
“Ossification”
prevention

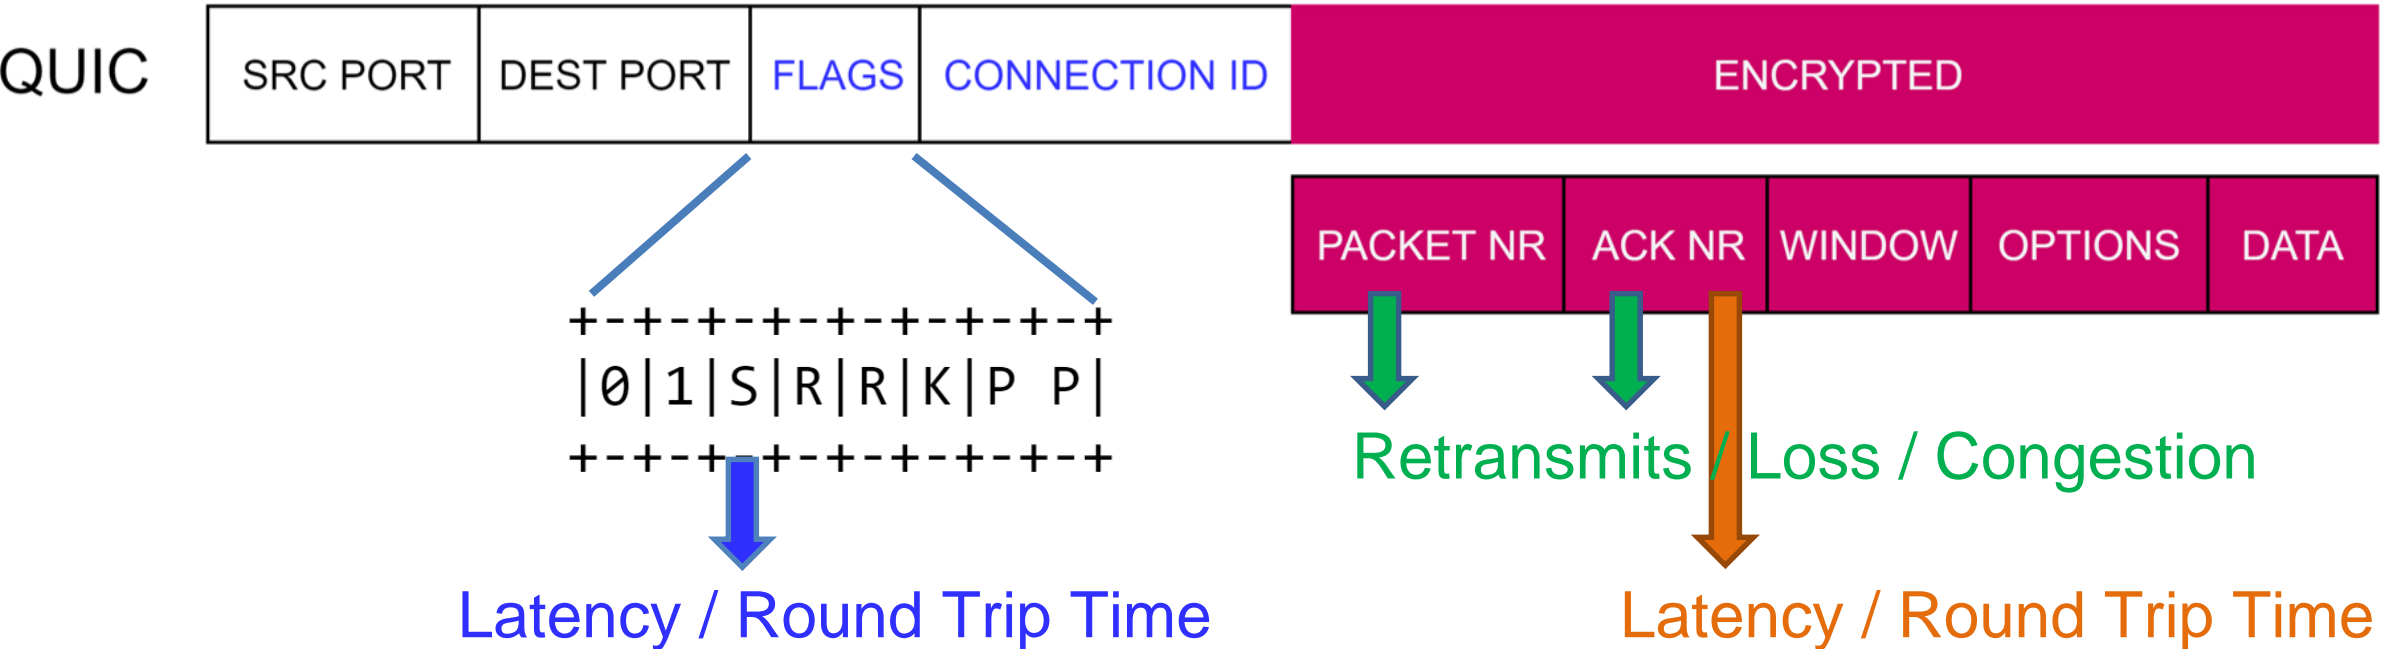
QUIC is end-to-end encrypted



QUIC is end-to-end encrypted



QUIC is end-to-end encrypted



Firewall / security logic

Fake them for extra performance

QUIC is end-to-end encrypted




QUIC is end-to-end encrypted

3 – 5% of networks



QUIC is end-to-end encrypted



~~3—5% of networks~~

20+% of networks

QUIC is end-to-end encrypted

~~3—5% of networks~~
20+% of networks



TCP fallback



QUIC is end-to-end encrypted: counterarguments

- Block QUIC = block big players (Google, FB, ...)
- QUIC doesn't need performance enhancing middleboxes
 - But... **satellites**
- They have no reason to block QUIC



QUIC is done in Userspace

“QUIC uses **only** 2x as an equivalent TCP + TLS stack”
- *Google engineers*

QUIC is done in Userspace


“QUIC uses **only** 2x as an equivalent TCP + TLS stack”

- *Google engineers*


“You need a sh*tload of **extra servers** to run QUIC”

- *What I'm reading*

QUIC is done in Userspace




(a) Varying object size, 1% Loss




(b) Varying object size, 112 ms RTT

Desktop

QUIC is done in Userspace



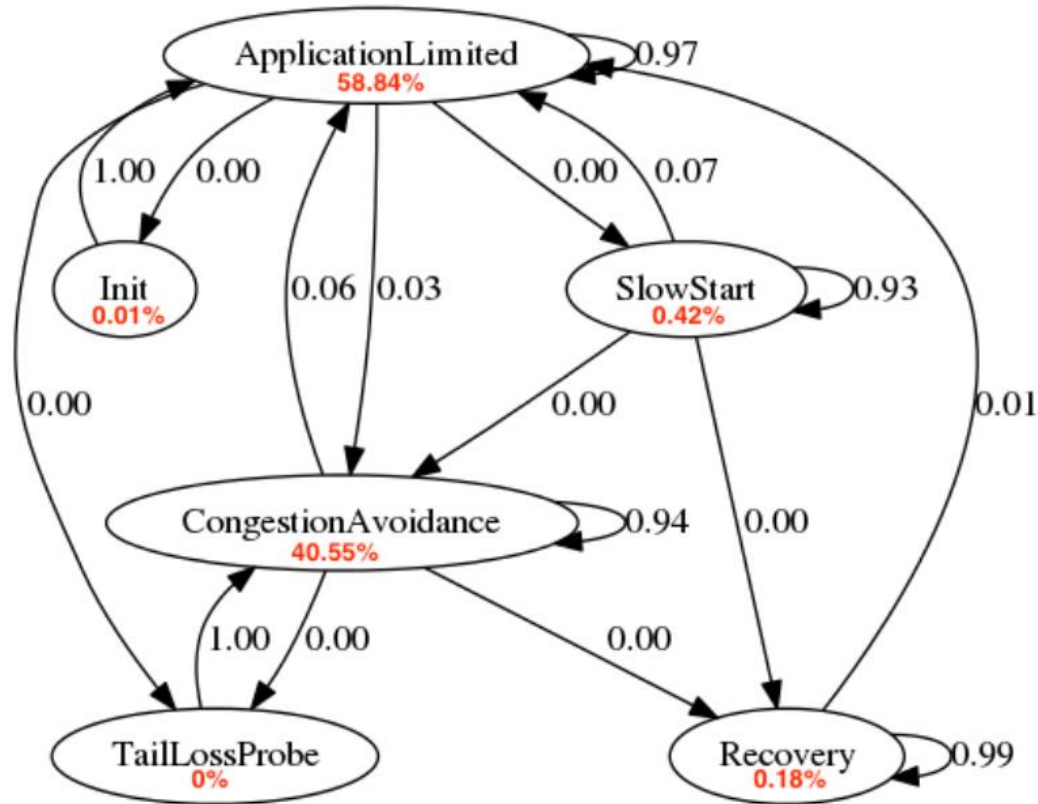
Desktop



Mobile


QUIC is done in Userspace

58.84%



(a) MotoG


7.05%




(b) Desktop

QUIC is done in Userspace: Counterarguments

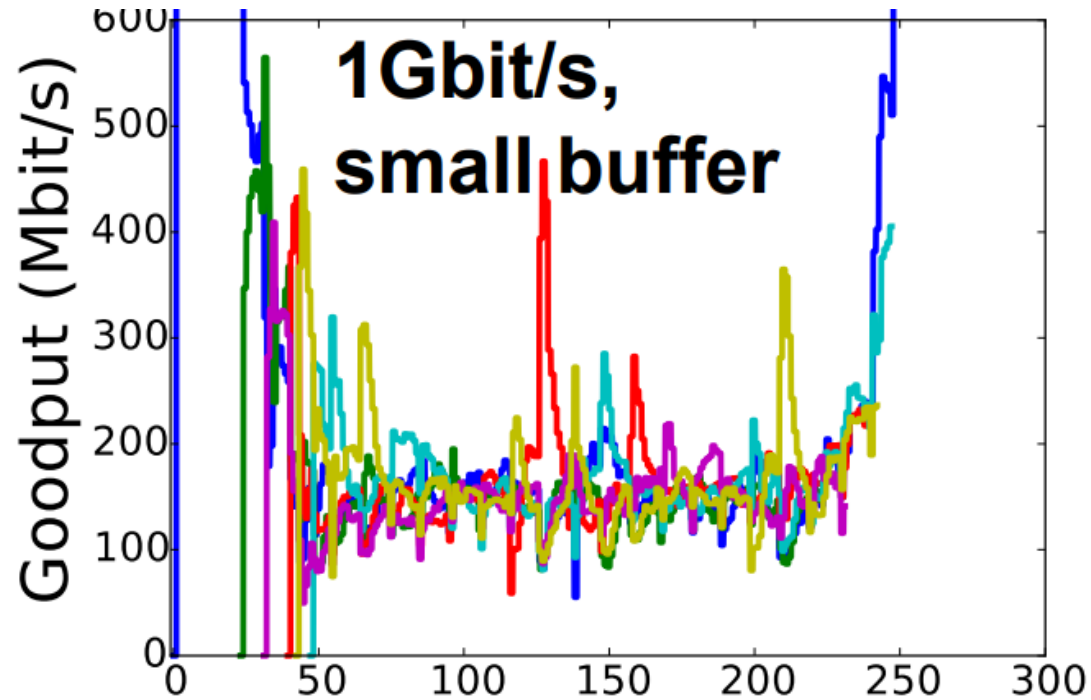
- QUIC will get hardware offload / move to kernel
 - But... **Variable-length encoding**
 - But... **ACK length ;)**
- Even with this overhead, **Google runs QUIC at scale**



QUIC is done in Userspace: Reprise




QUIC is done in Userspace: Reprise

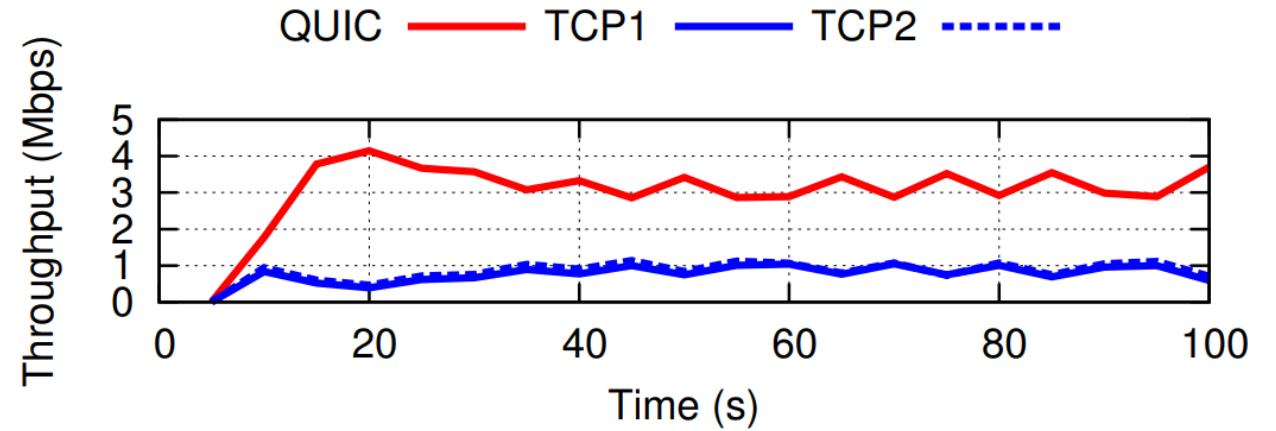


Fair

QUIC is done in Userspace: Reprise



Fair



(b) QUIC vs. two TCP flows


Not so fair

QUIC is done in Userspace: Reprise: Counterarguments

- Networks themselves will prevent abuse (AQM)
- Has been possible for ages, no real-world abuse noticed
 - But... **BBR**
 - But... 6 parallel TCP connections in HTTP/1.1




QUIC cuts down on latency with 0-RTT




TCP + TLS 1.2

2 RTT



TCP + TLS 1.3
Early Data

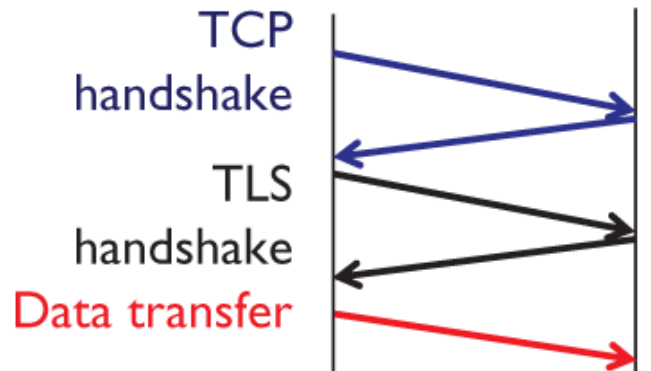
1 RTT



QUIC + TLS 1.3
Early Data


0 RTT

QUIC cuts down on latency with 0-RTT



TCP + TLS 1.2

2 RTT



TCP + TLS 1.3
Early Data

1 RTT



QUIC
Early Data

0 RTT

**TCP FAST OPEN +
TLS 1.3 Early Data**

QUIC cuts down on latency with 0-RTT



Poor Daniel

0-RTT HTTP POST
1. Pay Robin \$100 for his talk



Crafty Robin



QUIC cuts down on latency with 0-RTT




Poor Daniel

0-RTT HTTP POST
1. Pay Robin \$100 for his talk



Crafty Robin



Replay attack: can't just send anything

QUIC cuts down on latency with 0-RTT



Angry Daniel
1.1.1.1



Deserving Robin
2.2.2.2

0-RTT HTTP GET
"I am Robin at 2.2.2.2"

Send me one-gigabyte-file.json



Web Server

QUIC cuts down on latency with 0-RTT




Angry Daniel
1.1.1.1



Deserving Robin
2.2.2.2

0-RTT HTTP GET
"I am Robin at 2.2.2.2"

Send me one-gigabyte-file.json



one-gigabyte-file.json

UDP Amplification attack: can't send too much

QUIC cuts down on latency with 0-RTT



Angry Daniel
1.1.1.1




Deserving Robin
2.2.2.2

0-RTT HTTP GET
"I am Robin at 2.2.2.2"

Send me one-gigabyte-file.json



First 3-6 packets of
one-gigabyte-file.json



UDP Amplification attack: can't send too much

QUIC cuts down on latency with 0-RTT: counterarguments

- TCP Fast Open isn't feasible on real networks
 - But... just **right now**
- Clients can send **9000+** 0-RTT packets filled with padding
 - 1 0-RTT GET + 29 filled with zeroes => 90 packets response data!



QUIC has version negotiation

QUIC v3.5.66.6.8.55-Facebook

QUIC has version negotiation

QUIC v3.5.66.6.8.55-Facebook



QUIC **used to have** version negotiation

B.1. Since [draft-ietf-quic-transport-18](#)

- o Removed version negotiation; version negotiation, including authentication of the result, will be addressed in the next version of QUIC (#1773, #2313)

QUIC **used to have** version negotiation

B.1. Since [draft-ietf-quic-transport-18](#)

- o Removed version negotiation; version negotiation, including authentication of the result, will be addressed in the next version of QUIC (#1773, #2313)



QUIC **used to have** version negotiation: counterarguments

- We still have transport parameters and extension frames
- v2 will become main version and v1 will disappear quickly
- Clients will cache versions
 - But... **Caching is 1 of the 3 big problems in CS**



ACKCHYUALLY

▶▶ **UHASSETL EDM**

QUIC **used to have** version negotiation: counterarguments

- We still have transport parameters and extension frames
- v2 will become main version and v1 will disappear quickly
- Clients will cache versions
 - But... **Caching is 1 of the 3 big problems in CS**

- **2. Agreeing on the Spinbit**
- **3. Not logging plaintext passwords**
- **4. Off-by-one errors**



ACKCHYUALLY

QUIC is at exactly the right complexity

▪ V1

- Congestion control + loss detection
- Flow control
- Encryption and integrity protection
- Connection migration
- 0-RTT support
- Independent streams
- Low overhead
- DoS prevention
- Stateless Retry
- ...
- Not even talking about HTTP/3 features here


▪ Postponed to V2

- Multipath
- Forward error correction
- Unreliable data transfer
- Support for other crypto besides TLS 1.3
- ...
- Most non-HTTP/3 applications are being postponed to V2
- IoT, realtime media, ...

QUIC is at exactly the right complexity



QUIC is at exactly the right complexity



QUIC is at exactly the right complexity

- V1 is too complex
 - Will have **deployment issues and bugs** for a long time
 - Could lead to people holding off on usage
- V1 is not complex enough
 - Tougher to convince things like IoT/games to switch later on



QUIC is at exactly the right complexity: counterarguments

- HTTP/2 has been buggy for years, still used
- QUIC can evolve very rapidly: V2 will be here soon
- QUIC is meant for the long run
 - But... **uptake momentum** is important too



ACKCHYUALLY

▶▶ **UHASSELT EDM**


<https://github.com/andydavies/http2-prioritization-issues>
<https://twitter.com/AndyDavies/status/1065916677408346112>
<https://blog.cloudflare.com/http-2-prioritization-with-nginx/>
<https://speeder.edm.uhasselt.be/www18/>

QUIC uses TLS 1.3, so it's secure

- TLS 1.3 in itself seems valid enough
 - But QUIC uses it in **new ways**
- Lots of discussion at the IETF this week
 - **Key updates**, version negotiation, amplification prevention, ...
- If attack is found, might need to disable QUIC completely
 - Luckily: **easy and fast to update**

QUIC uses TLS 1.3, so it's secure

- TLS 1.3 in itself seems valid enough
 - But QUIC uses it in **new ways**
- Lots of discussion
 - Key updates, version
- If attack is found
 - Luckily: **easy and**



Summary of **CONFIRMED QUIC FACTS**

- QUIC pisses off network and firewall operators
- QUIC is slow and destroys batteries
- QUIC traffic will drown out all TCP flows
- QUIC's 0-RTT is completely useless
- QUIC will incur version negotiation every single time
- QUIC is too complex and not complex enough at the same time.
- QUIC is unsafe and **will lead to Trump's re-election**

Prediction

“QUIC will
become the
major internet
transport in 5
years”

- *Bill Gates (probably)*



bit.ly/quicsurvey