Multipath QUIC

Quentin De Coninck

https://multipath-quic.org

UCLouvain

fnrs LA LIBERTÉ DE CHERCHER

Outline

- Designing Multipath for (G)QUIC
- Evaluating Multipath Benefits
- Adapting to IETF QUIC
- Open Challenges and Opportunities

Why do we want Multipath?

Why Multipath QUIC?

- **QUIC assumes a single-path flow**

Why Multipath QUIC?

- **QUIC assumes a single-path flow**

Why Multipath QUIC?

- **QUIC assumes a single-path flow**

Why Multipath QUIC?

- QUIC assumes a single-path flow

- Multipath QUIC

Why Multipath QUIC?

- **QUIC assumes a single-path flow**

- **Multipath QUIC**
  - Bandwidth aggregation

Why Multipath QUIC?

- **QUIC assumes a single-path flow**

- **Multipath QUIC**
  - Bandwidth aggregation
  - Seamless network handover
    - Can try new WiFi while keeping using LTE

Hem, connection migration?

Hem, connection migration?

- **Connection ID(s) to identify flow**
  - Resilient to 4-tuple change (IP, port)

Hem, connection migration?

- **Connection ID(s) to identify flow**
  - Resilient to 4-tuple change (IP, port)
- **IETF QUIC probes "paths"**
  - PATH_CHALLENGE / PATH_RESPONSE

Hem, connection migration?

- **Connection ID(s) to identify flow**
  - Resilient to 4-tuple change (IP, port)
- **IETF QUIC probes “paths”**
  - PATH_CHALLENGE / PATH_RESPONSE
- **Required mechanisms for multipath**
- **But no simultaneous usage of paths for data exchange**

Disclaimer

- In the remaining of this section, only (old) Google QUIC version is explained
- The IETF version has its dedicated section :-)

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

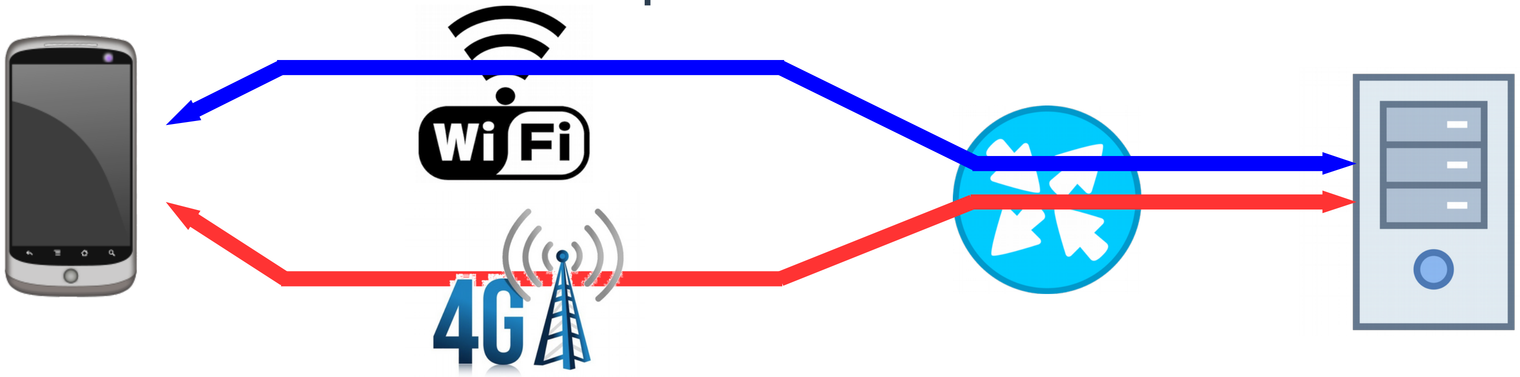

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

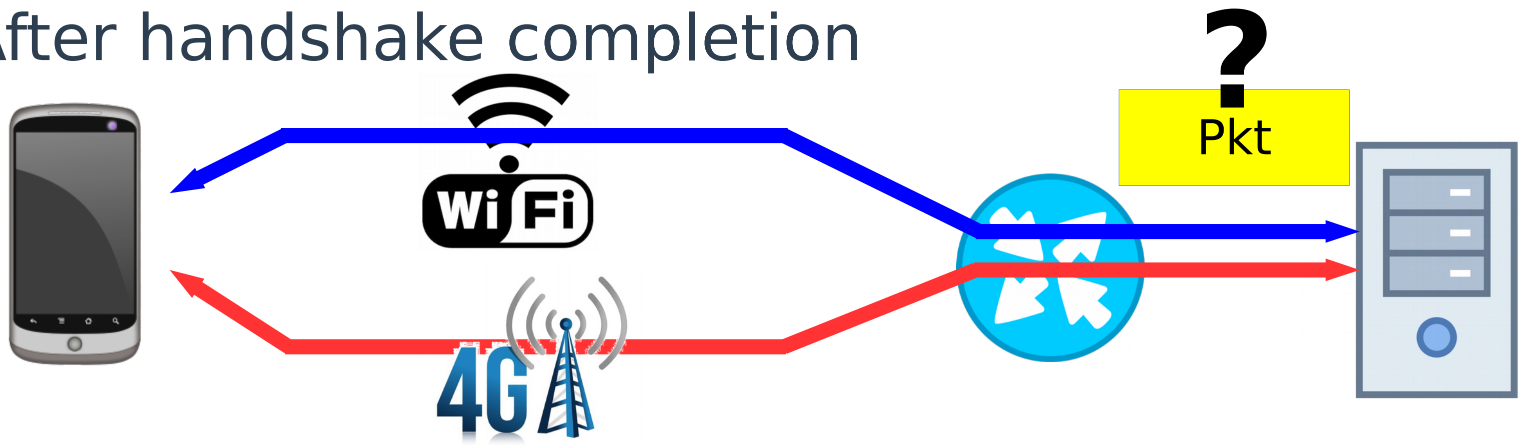

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

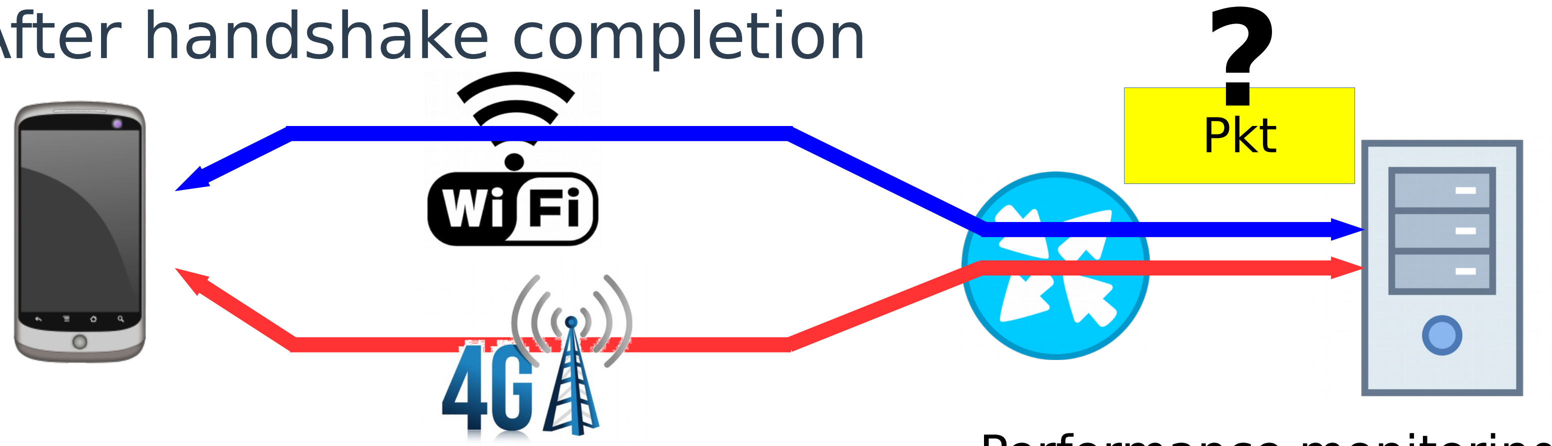

Performance monitoring?
Loss detection?
Path congestion control?

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

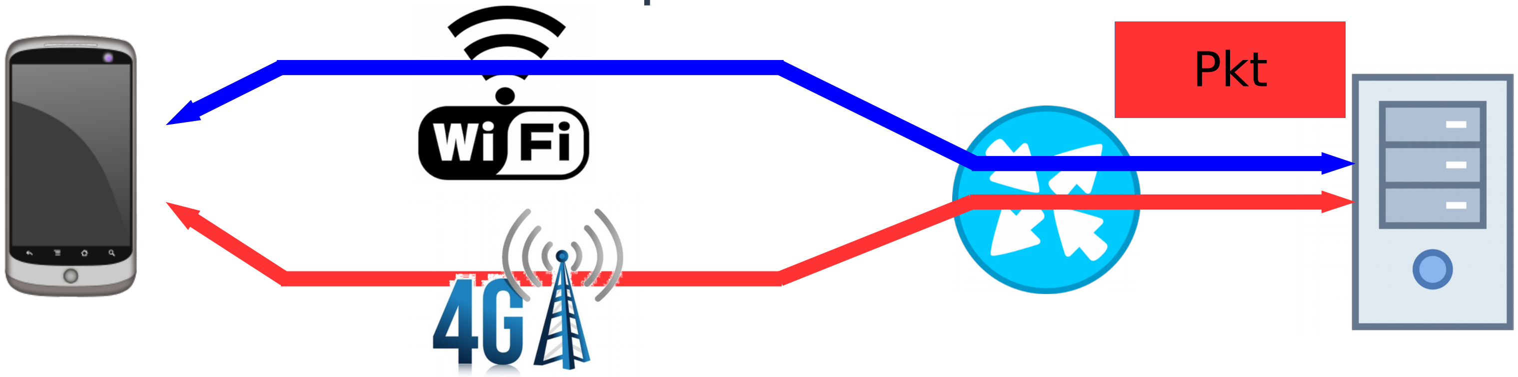

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

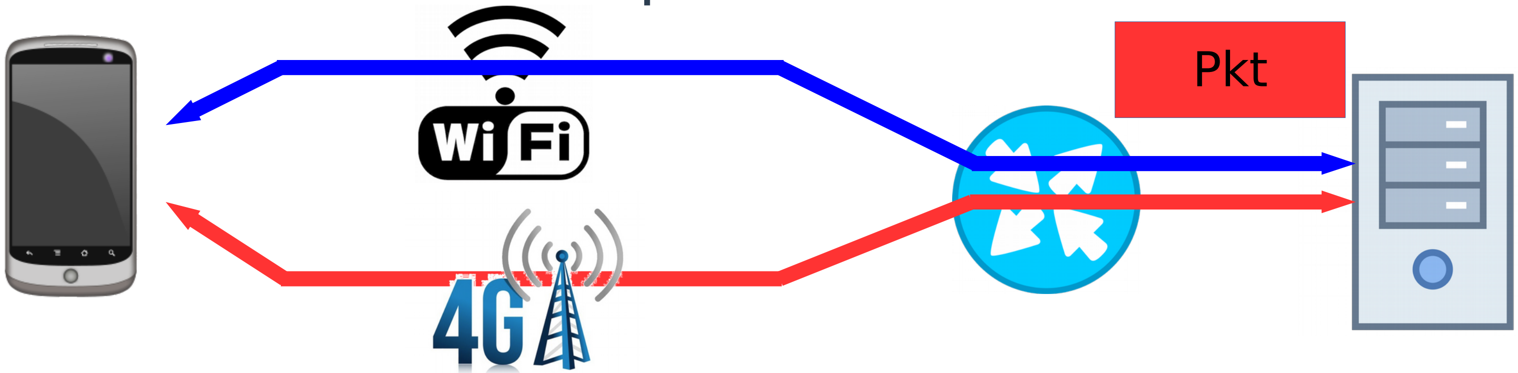

Flags	Connection ID	Path ID	Packet Number	Encrypted Payload...

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

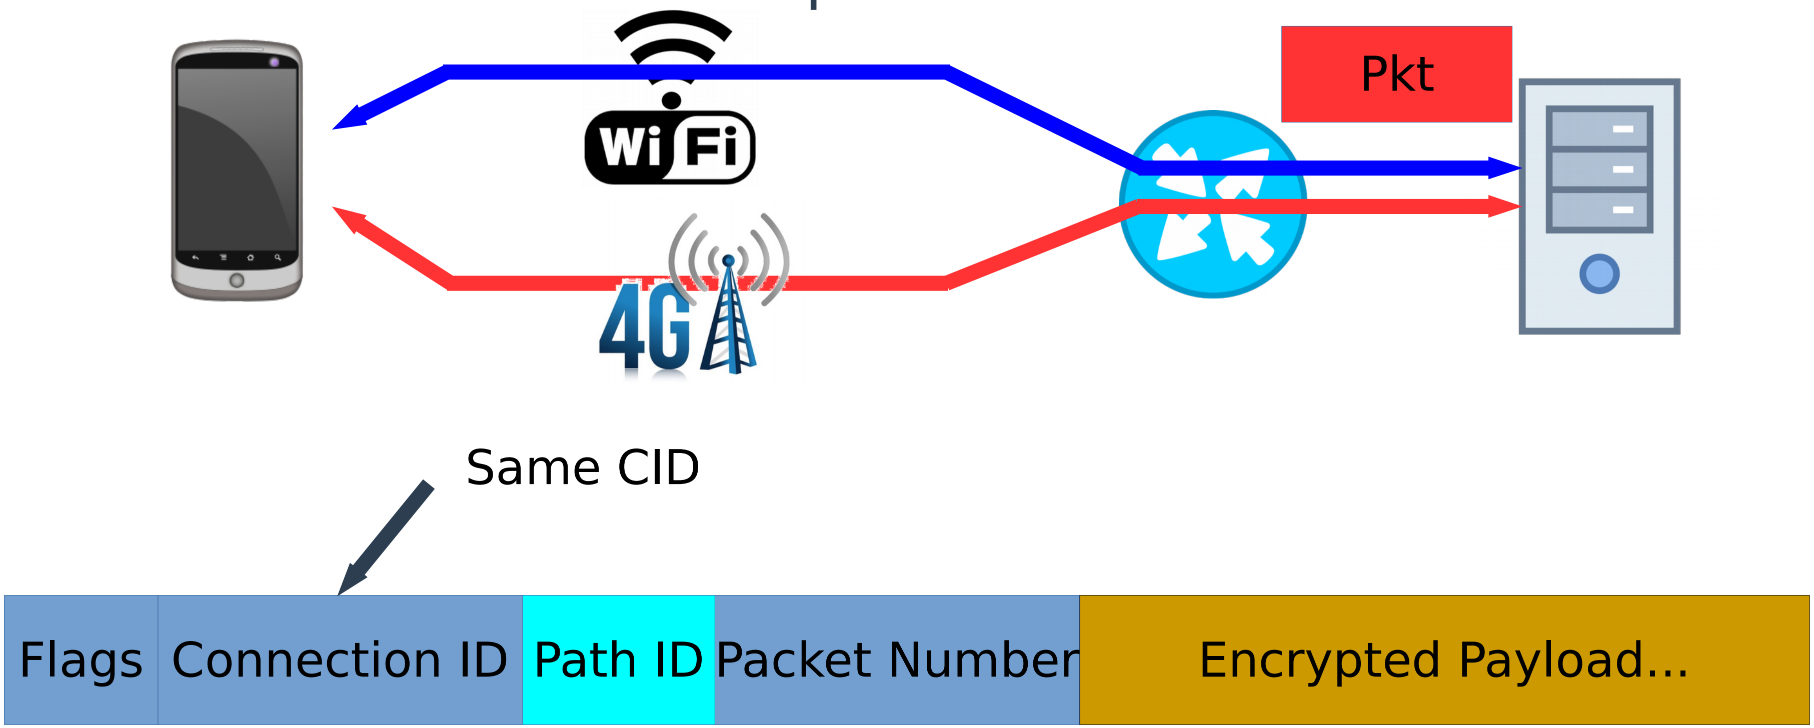

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

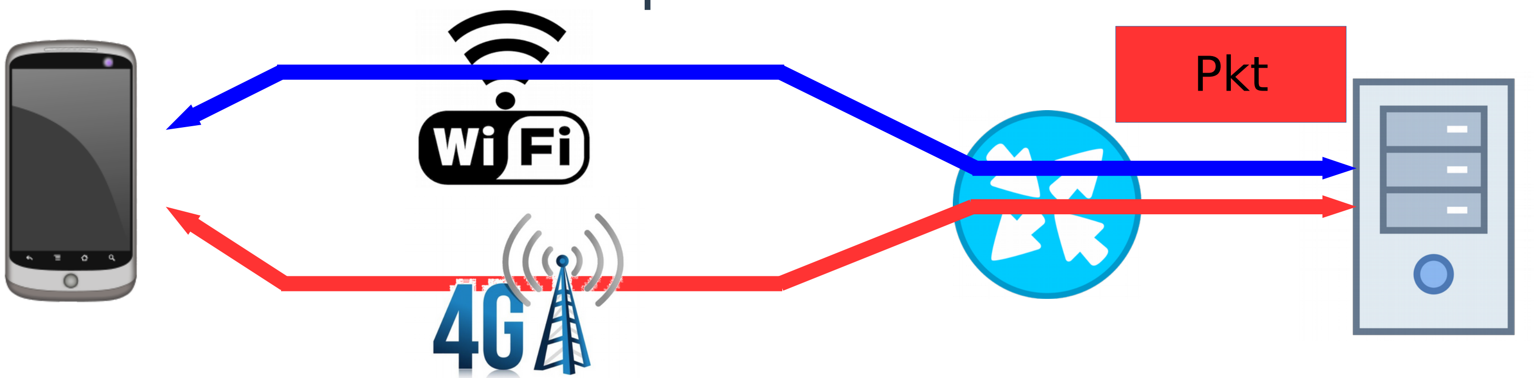

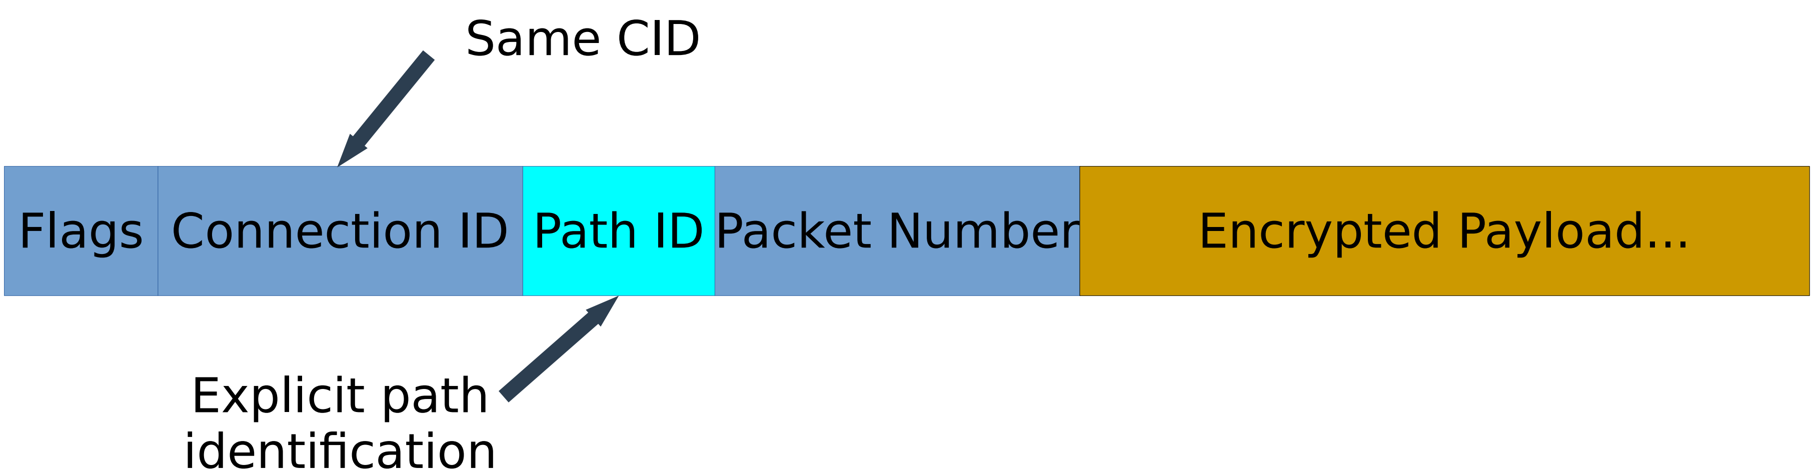

Designing Multipath (G)QUIC

- **Connection is composed of paths**
  - After handshake completion

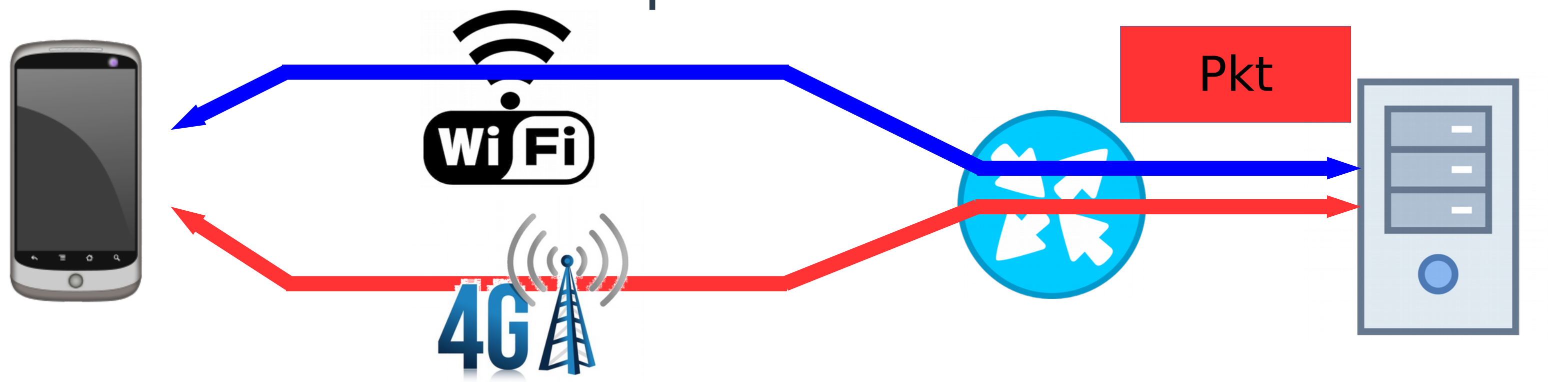

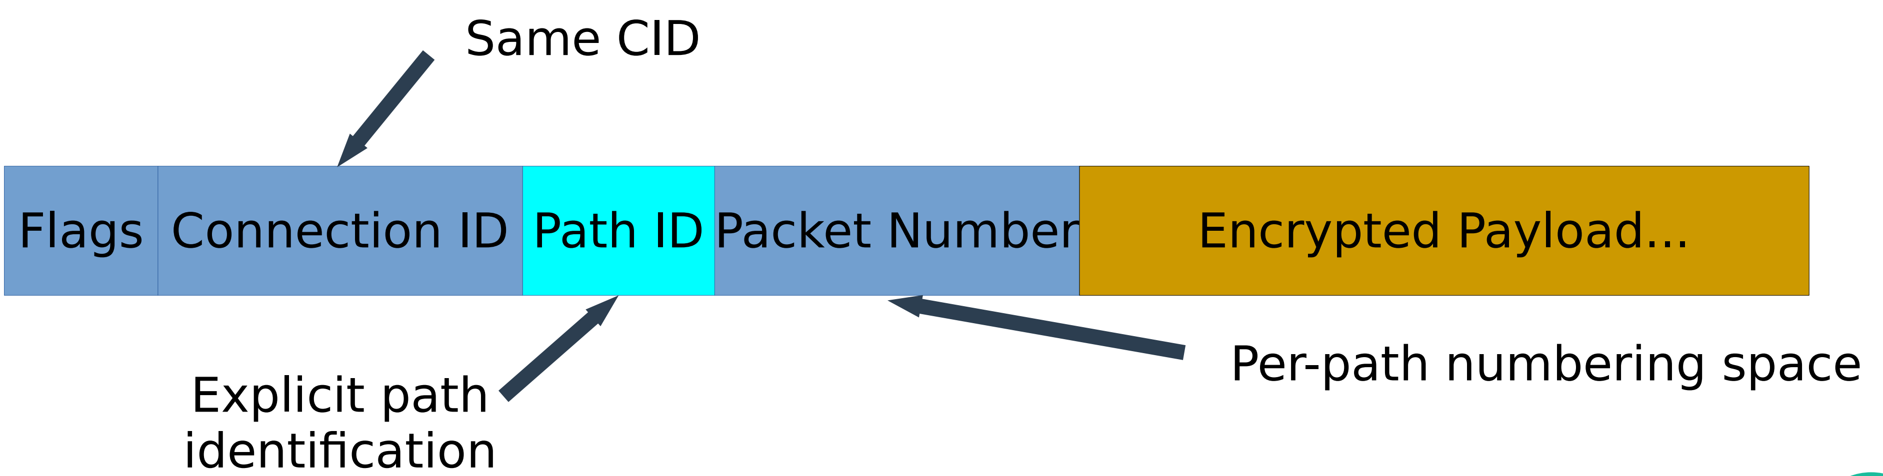

Architecture of Multipath (G)QUIC

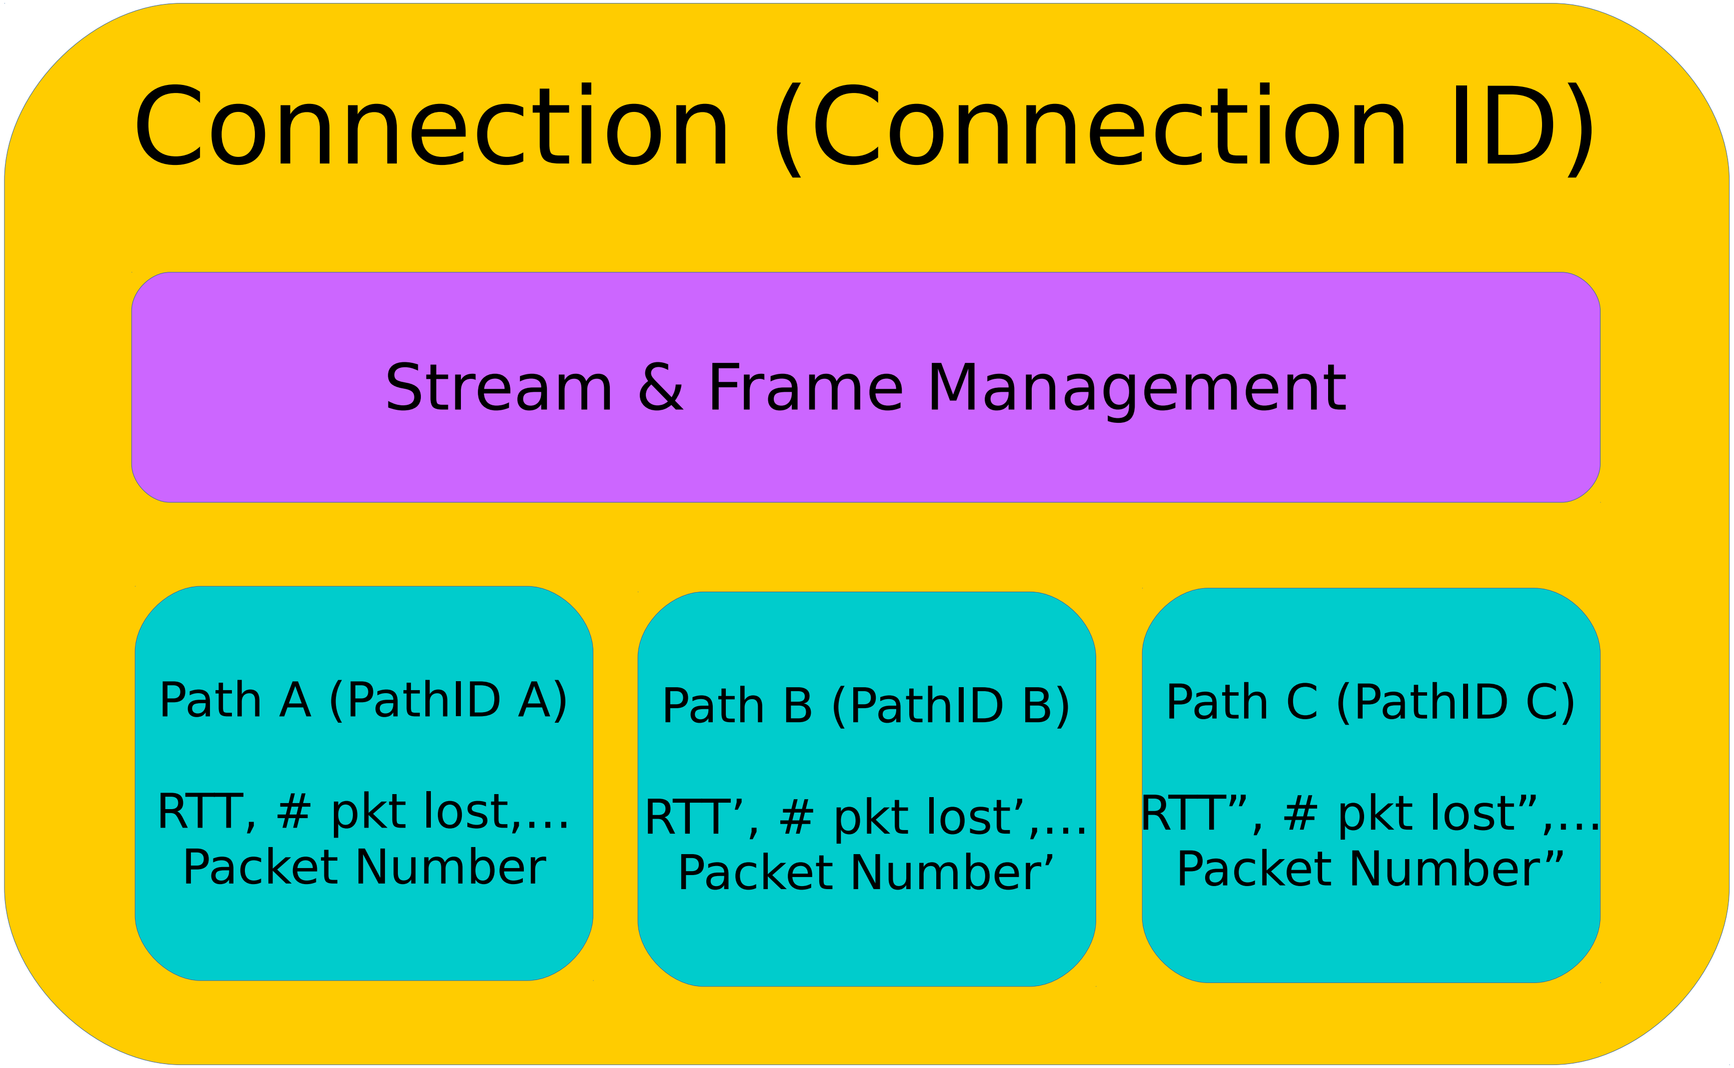

Multipath QUIC Data Transfer

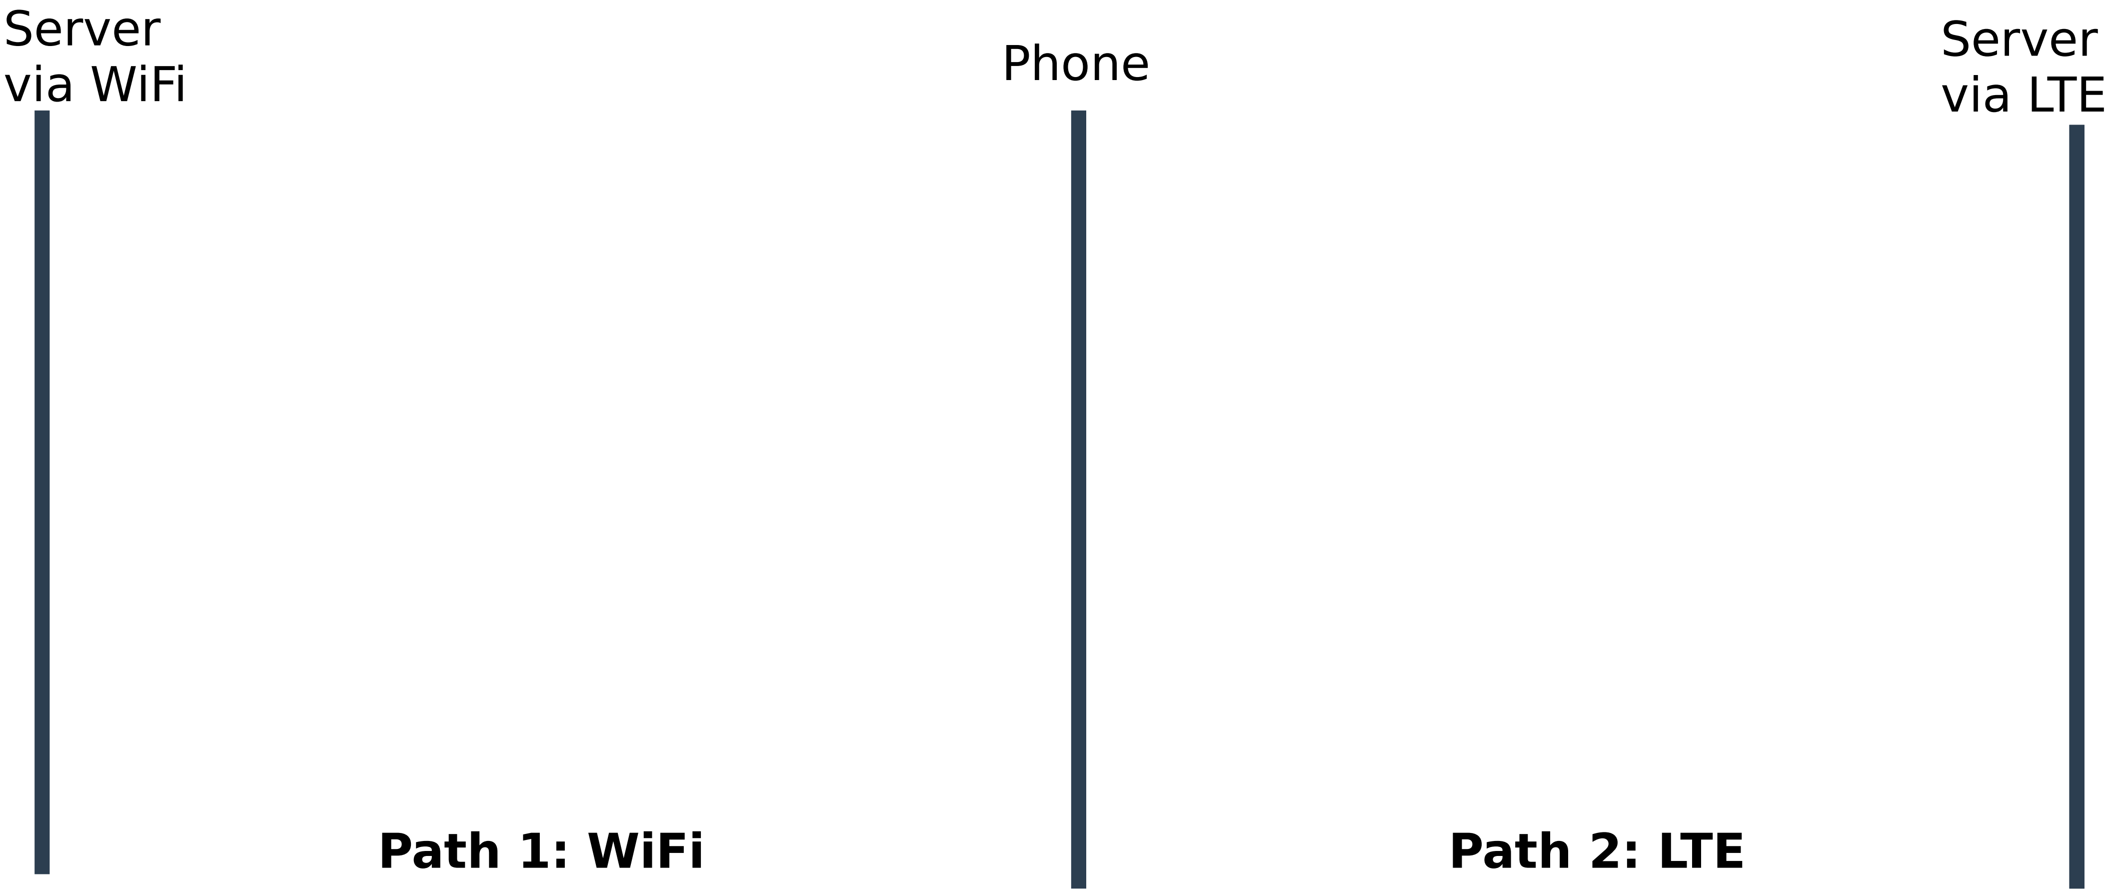

Multipath QUIC Data Transfer

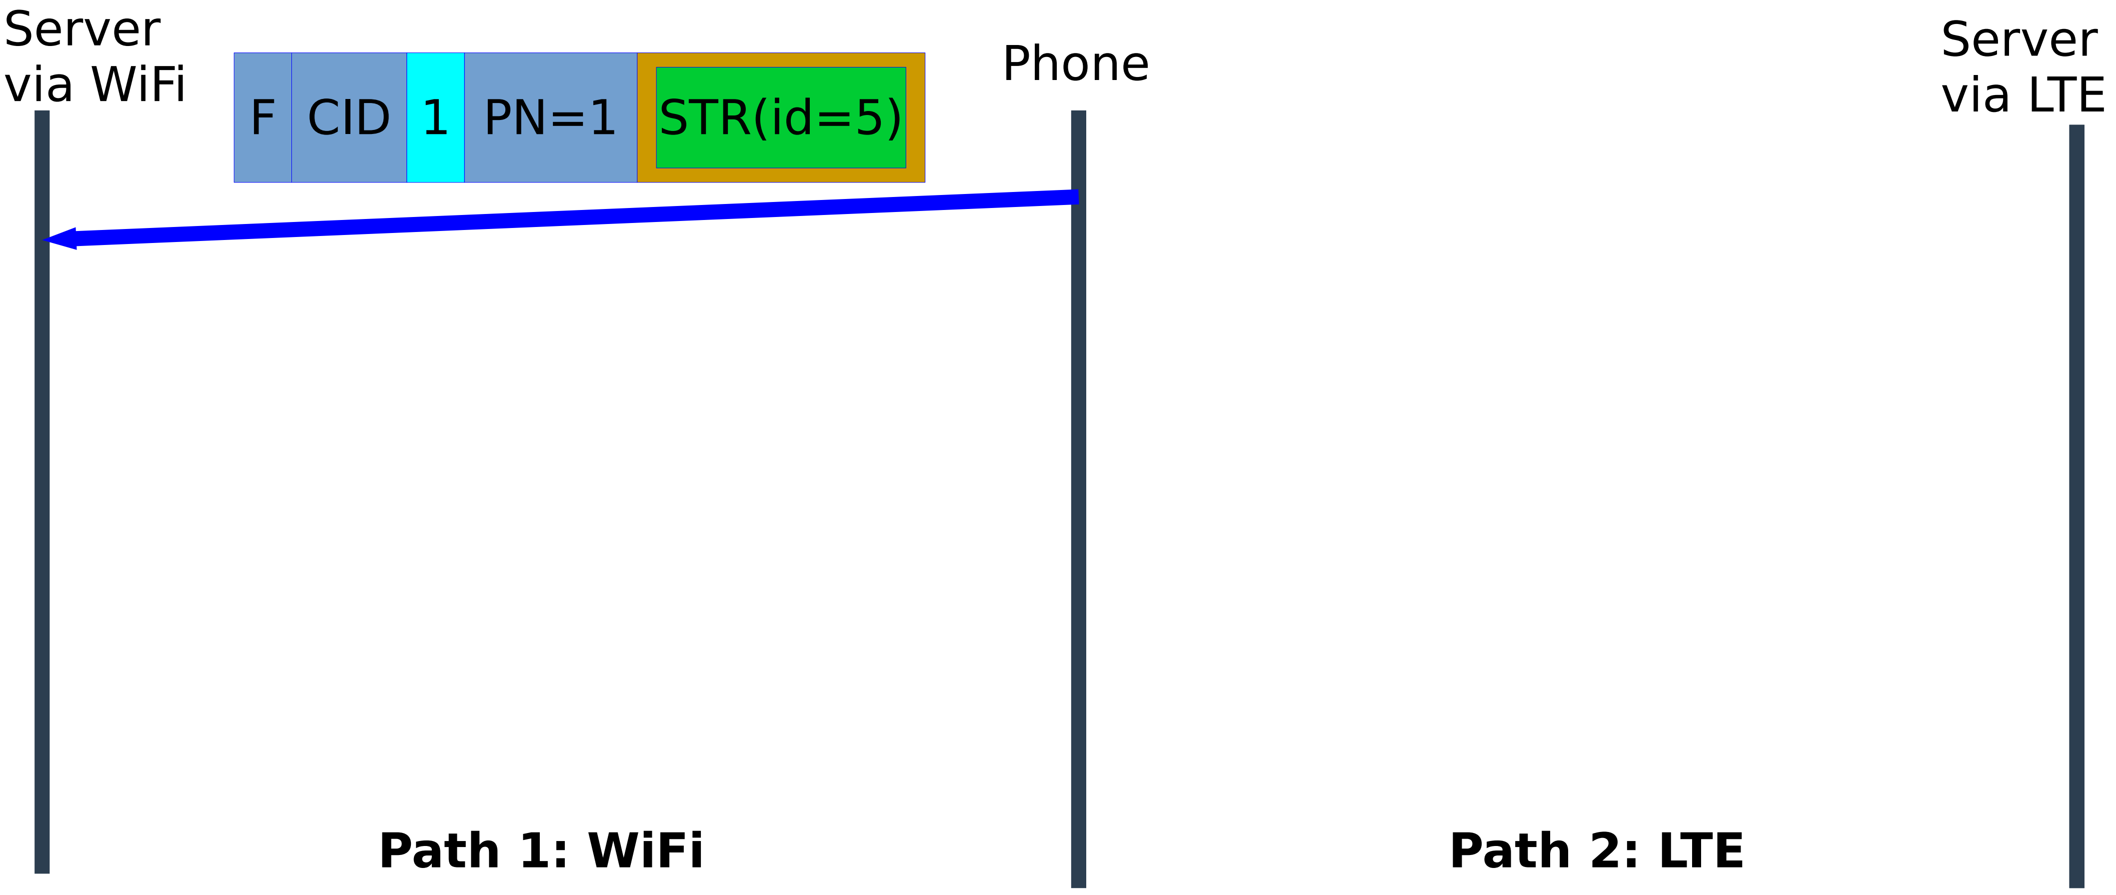

Multipath QUIC Data Transfer

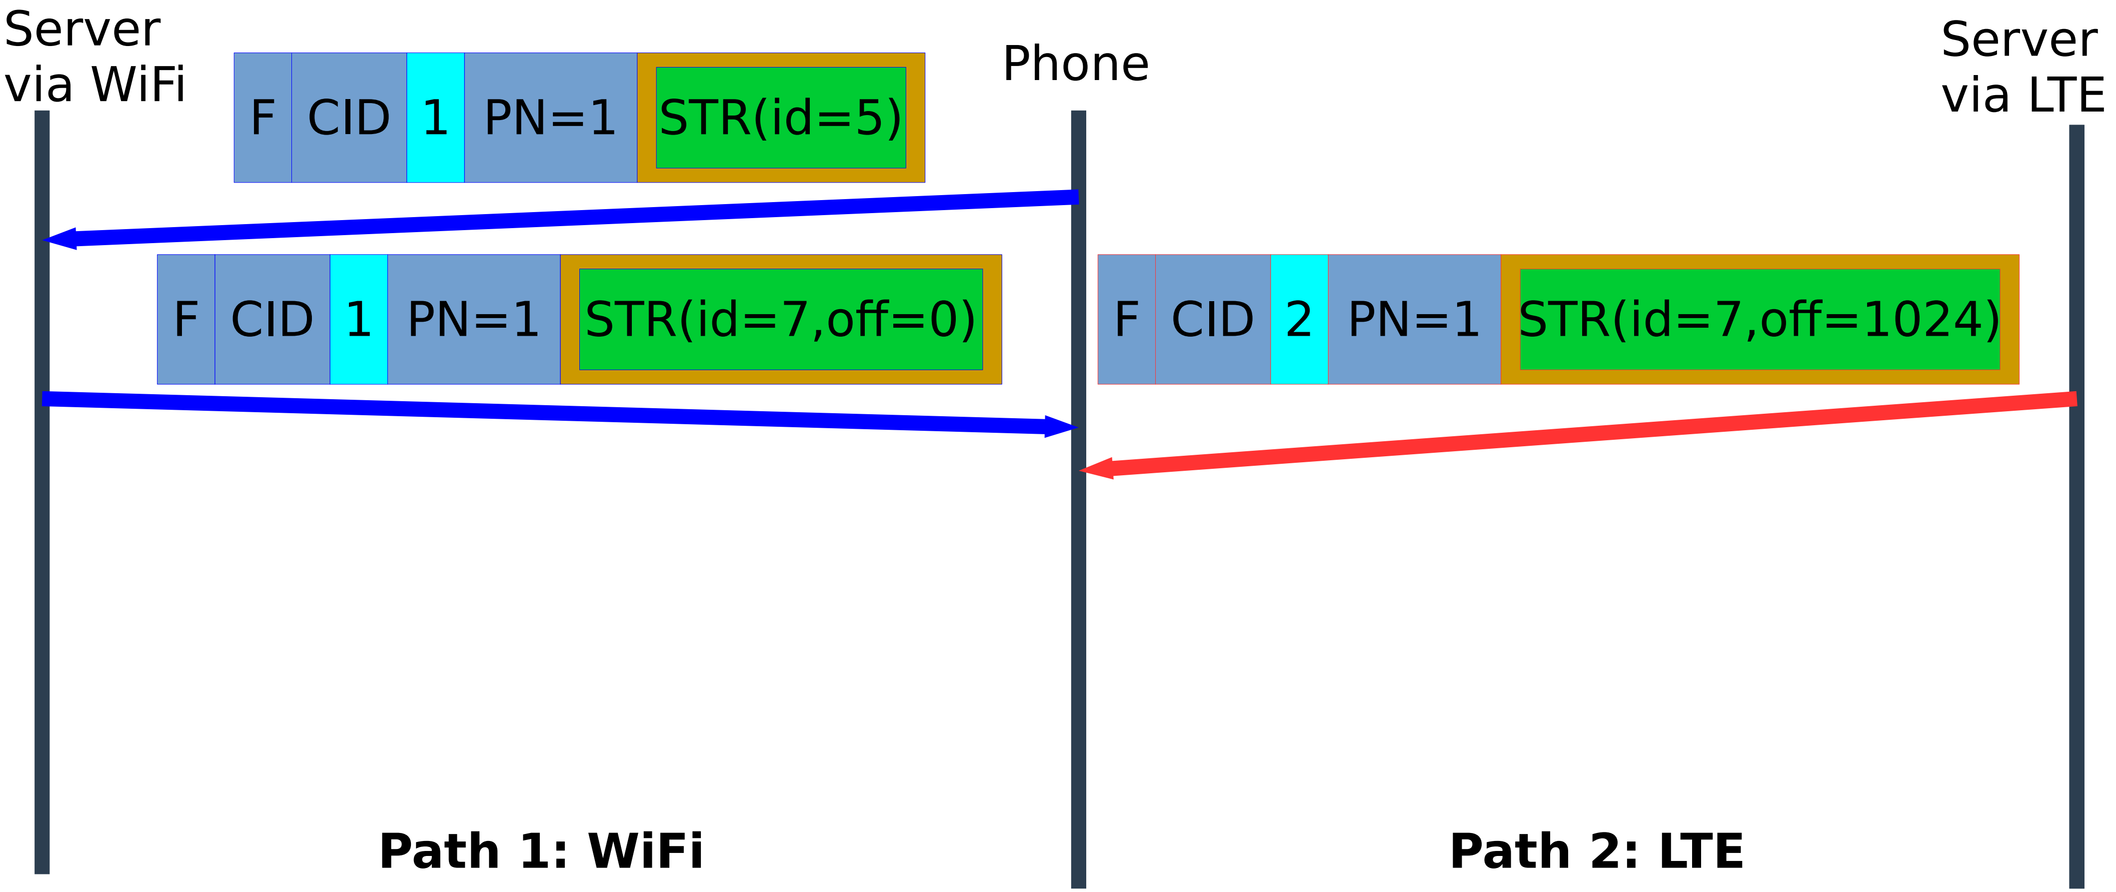

Multipath QUIC Data Transfer

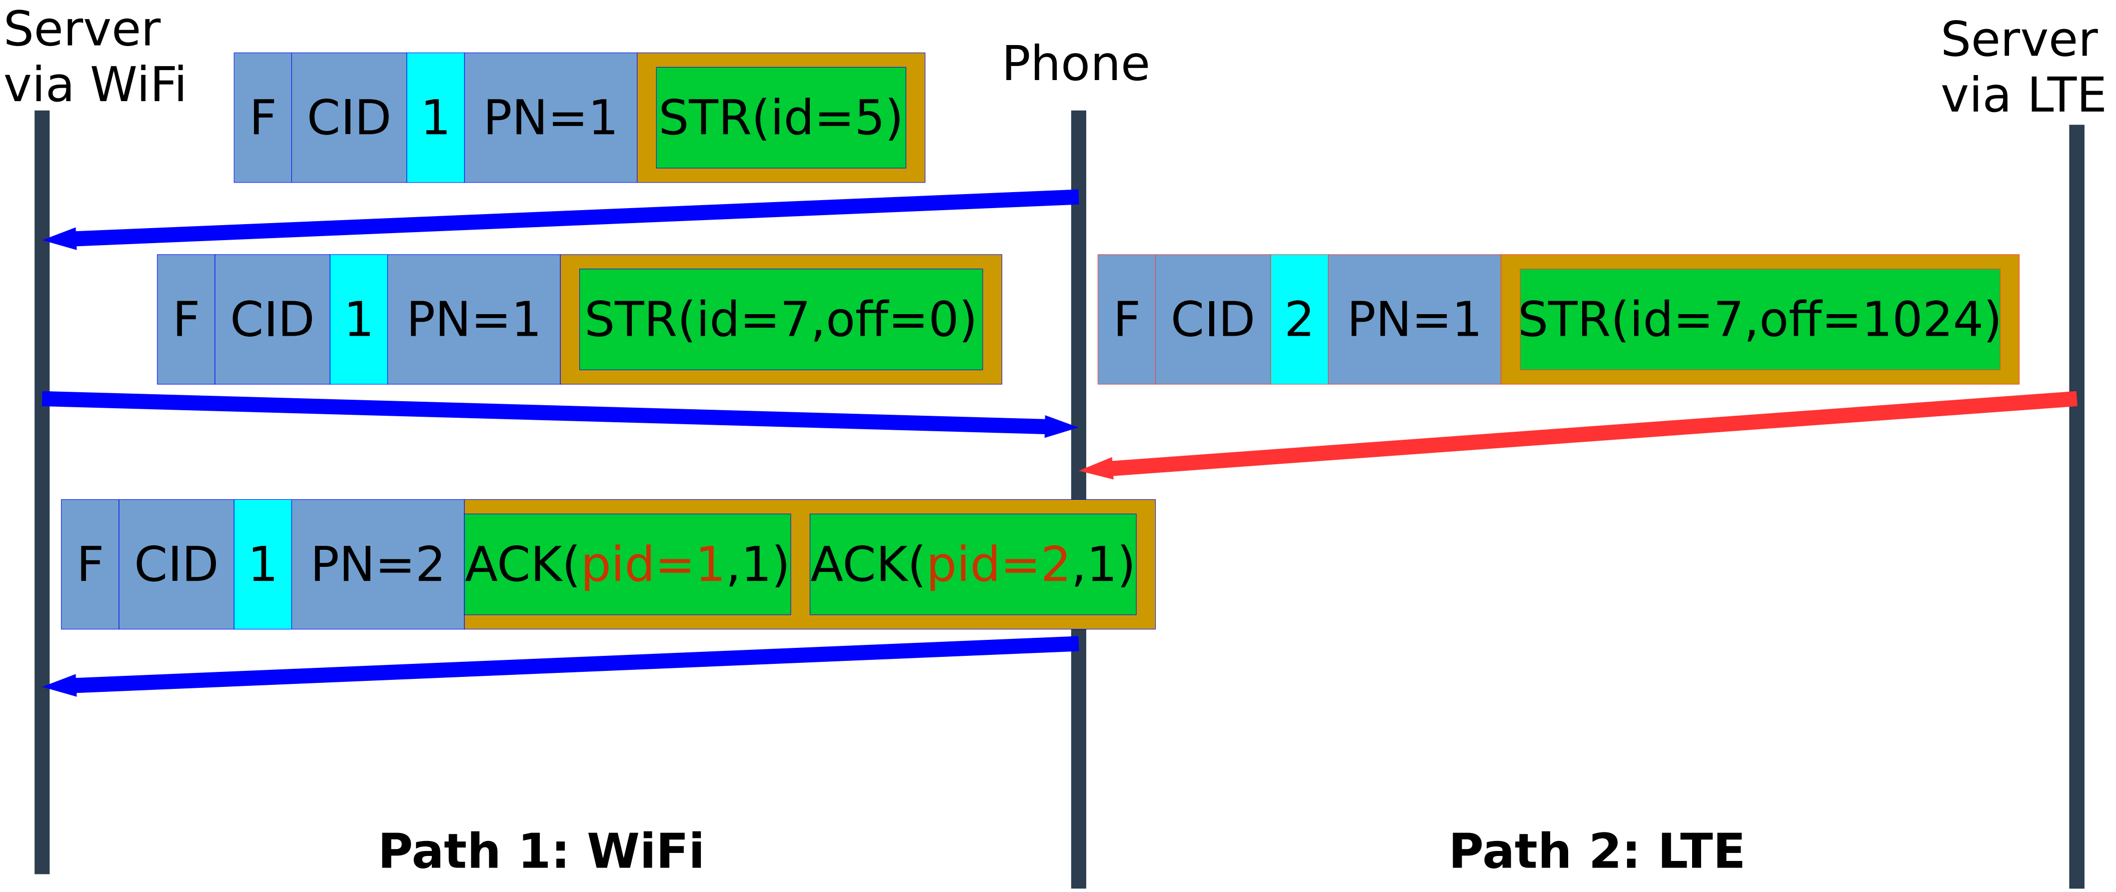

Multipath QUIC Data Transfer

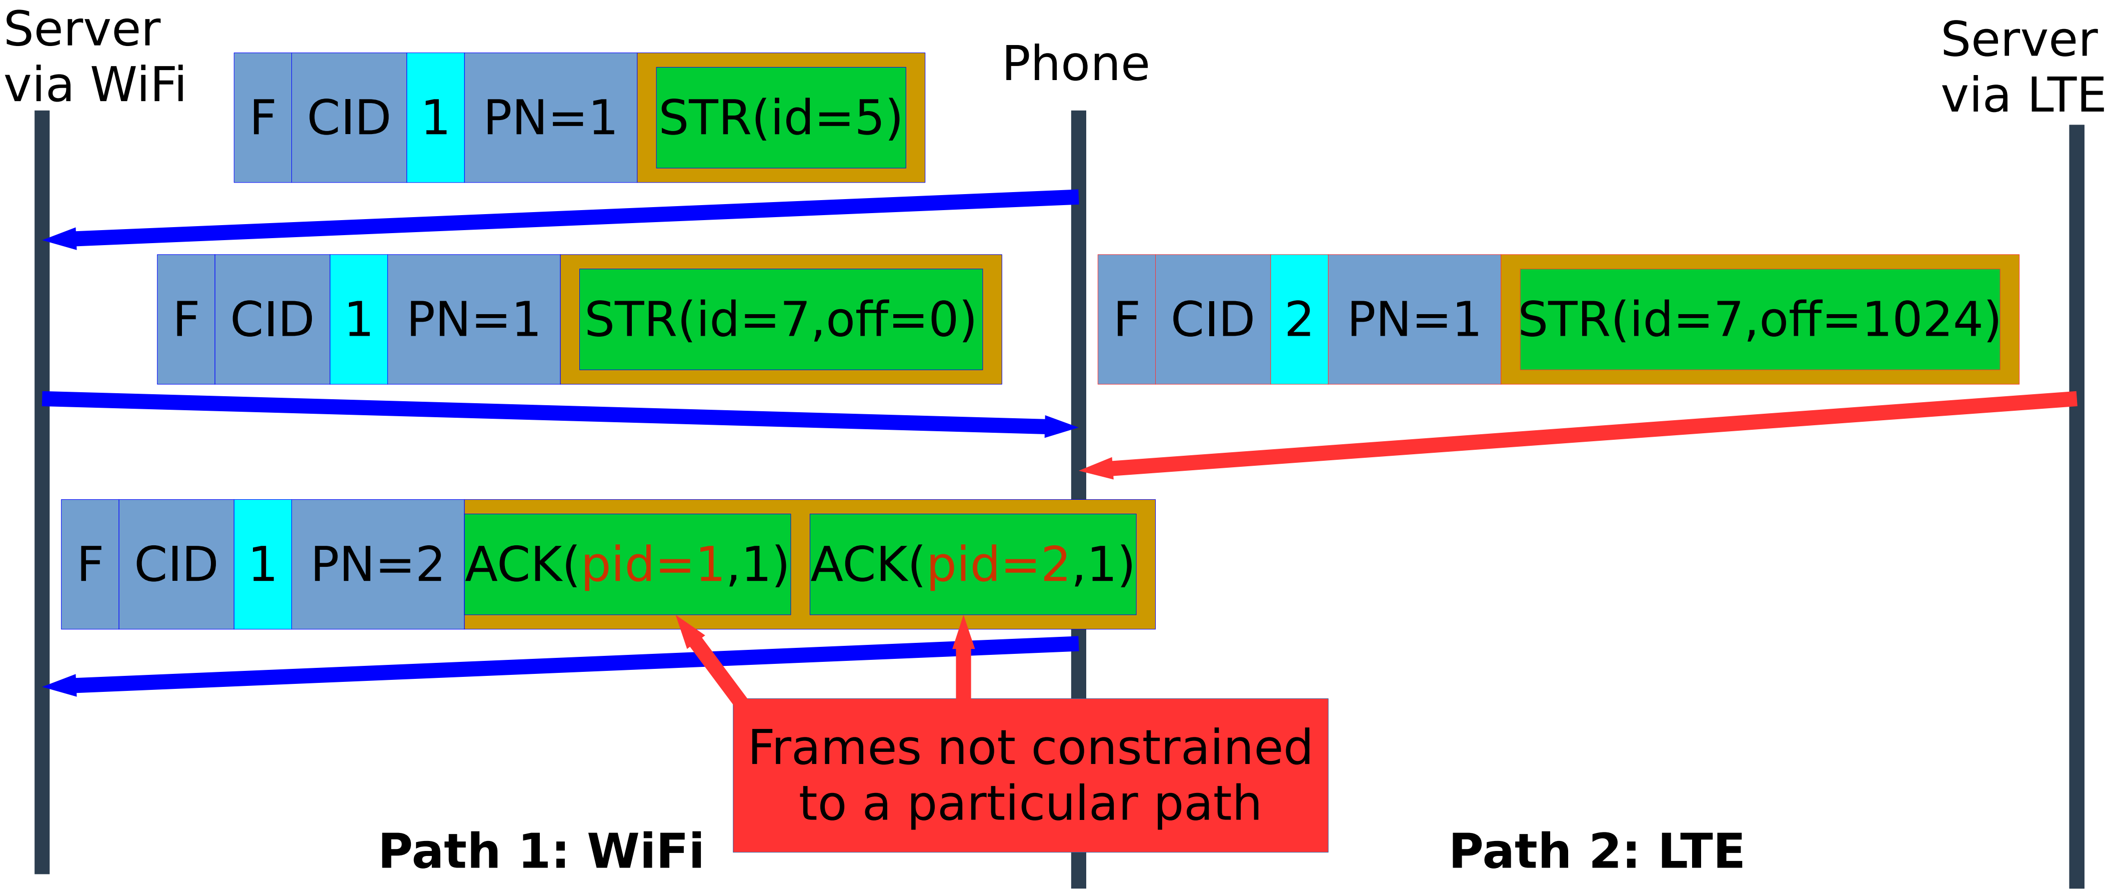

Multipath Mechanisms

- Path management
- Packet scheduling
- Congestion control scheme

Path Management

- **How and when paths are established?**

Path Management

- **How and when paths are established?**

Path Management

- **How and when paths are established?**

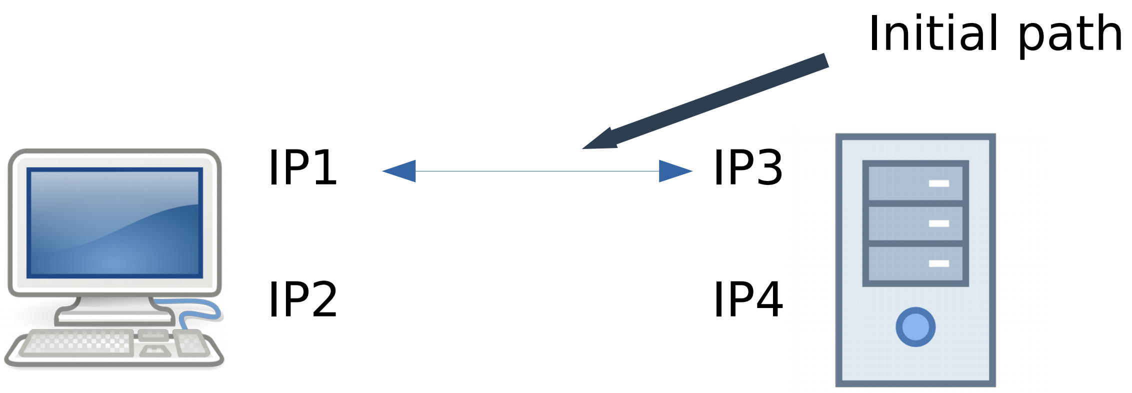

Path Management

- **How and when paths are established?**

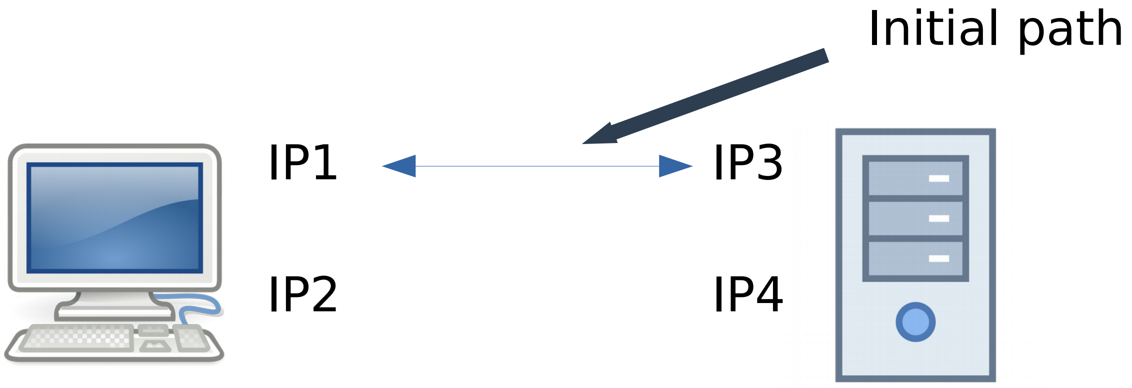

- **ADD_ADDRESS + REMOVE_ADDRESS frames**

Path Management

- **How and when paths are established?**

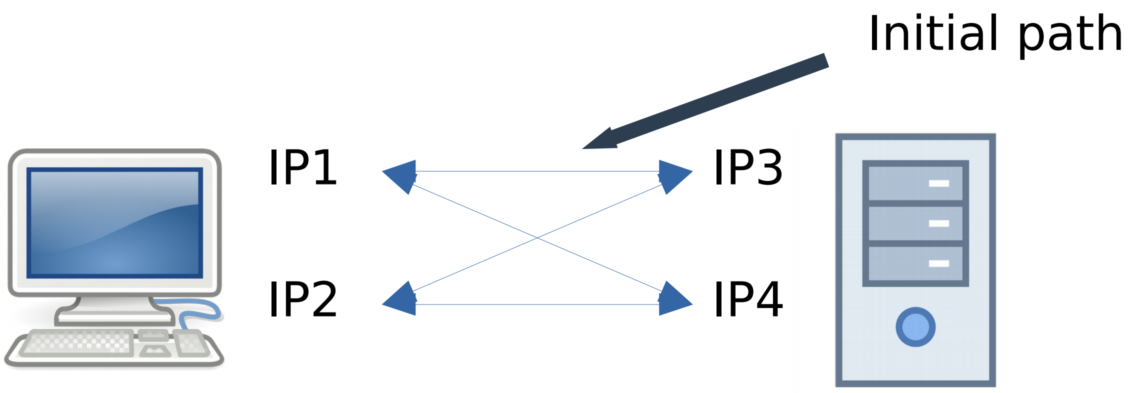

- **ADD_ADDRESS + REMOVE_ADDRESS frames**
- **Full-mesh fashion**

Packet Scheduling

- **Lowest-latency first**

Packet Scheduling

- **Lowest-latency first**

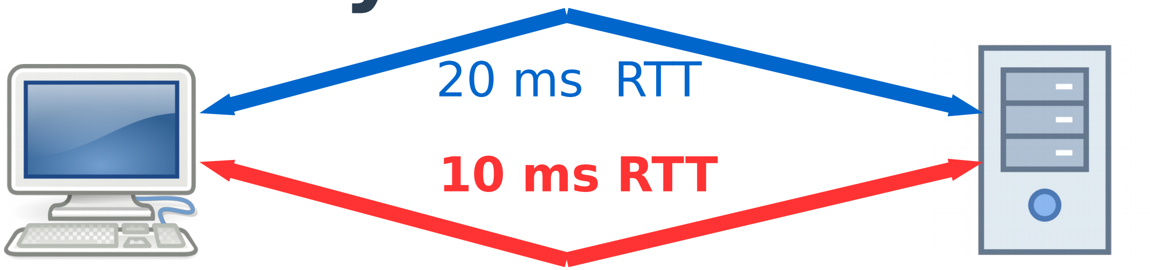

Packet Scheduling

- **Lowest-latency first**

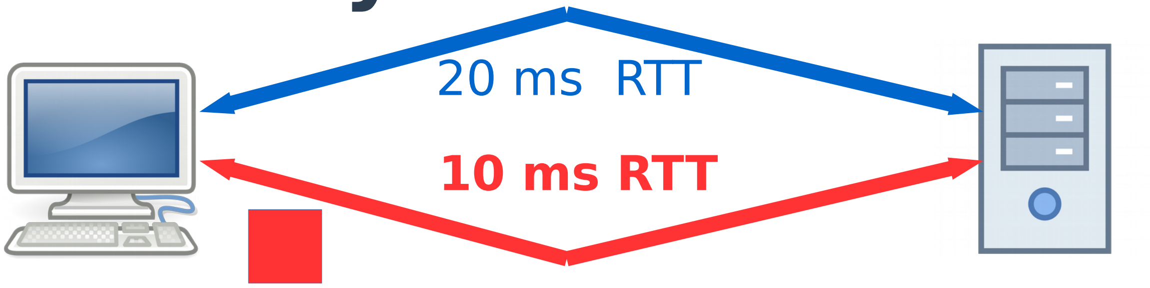

Packet Scheduling

- **Lowest-latency first**

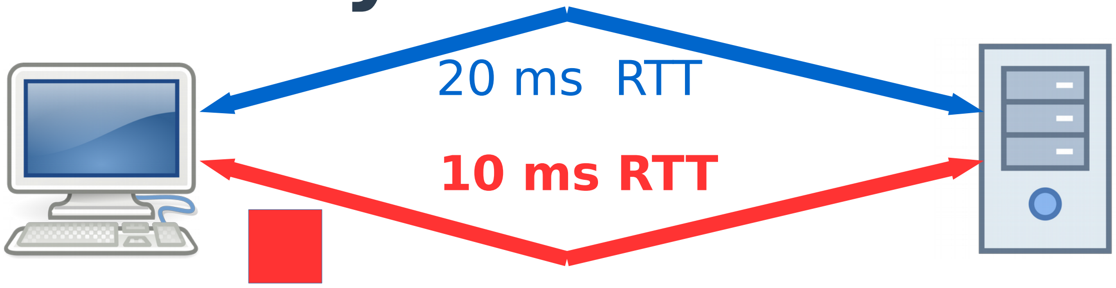

- **What if the path latency is unknown?**

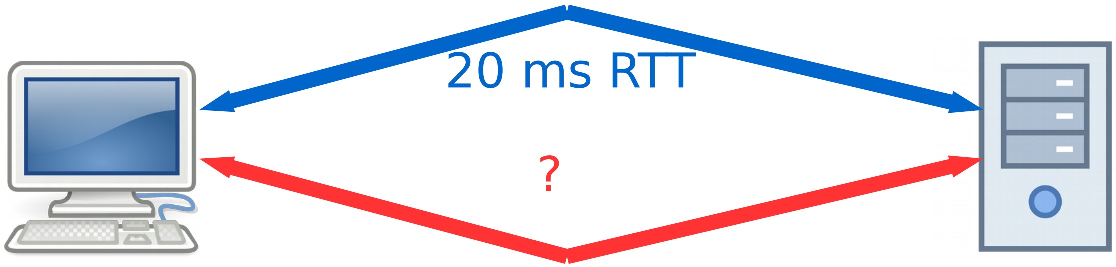

Packet Scheduling

- **Lowest-latency first**

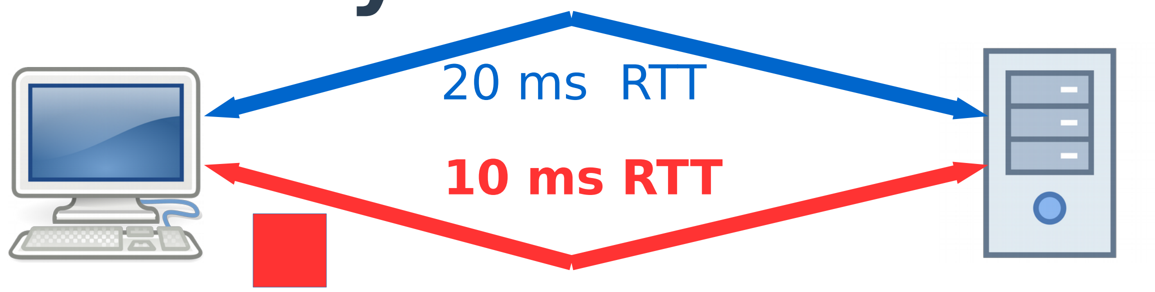

- **What if the path latency is unknown?**

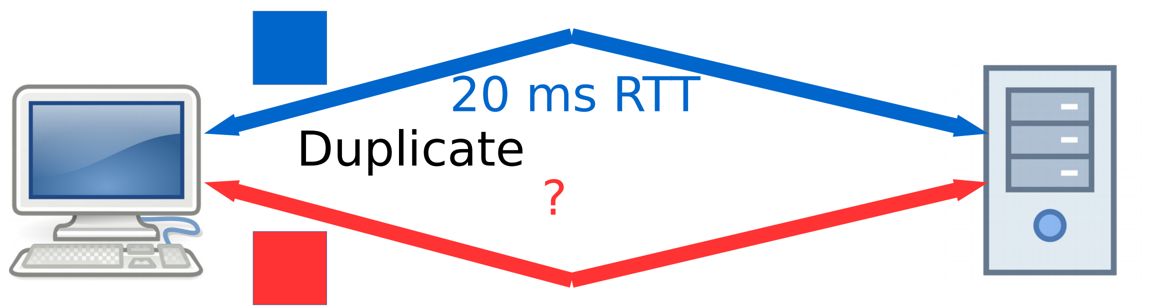

Packet Scheduling

- **Lowest-latency first**

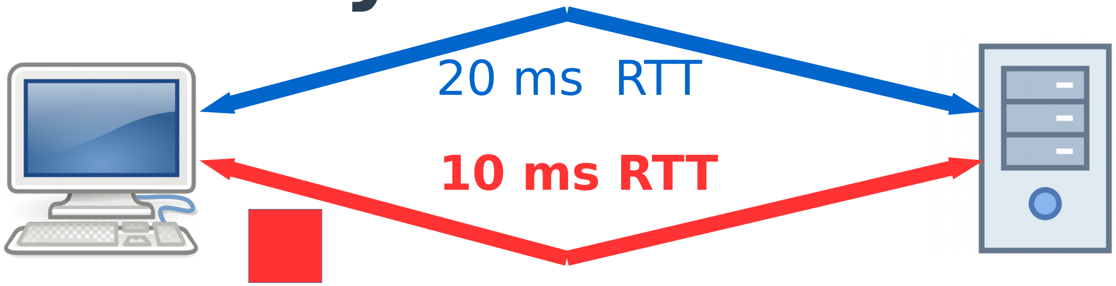

- **What if the path latency is unknown?**

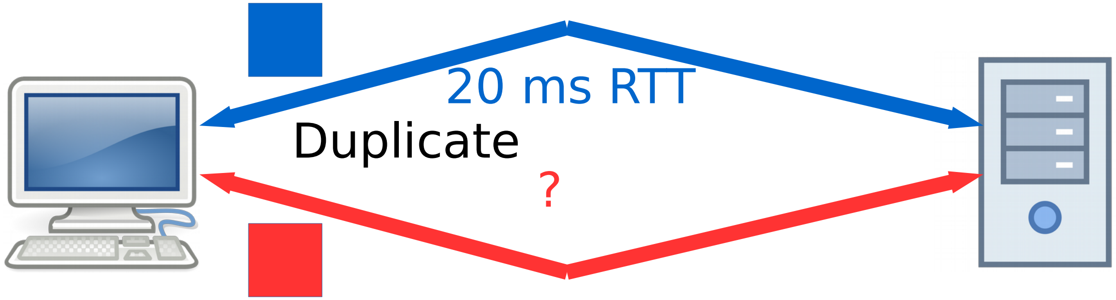

- **Schedule ALL frames (not only data)**

Congestion Control Scheme

- **Multipath = need for coupled CC**
  - CUBIC would be unfair
- **Opportunistic Linked Increase Algorithm**
  - MPTCP state-of-the-art

Evaluating Multipath Benefits

Evaluation of Multipath QUIC

CoNEXT'17

- **(Multipath) QUIC vs. (Multipath) TCP**
  - Multipath QUIC: based on *quic-go*
  - Linux MPTCP v0.91 with default settings
- **Mininet environment with 2 paths**

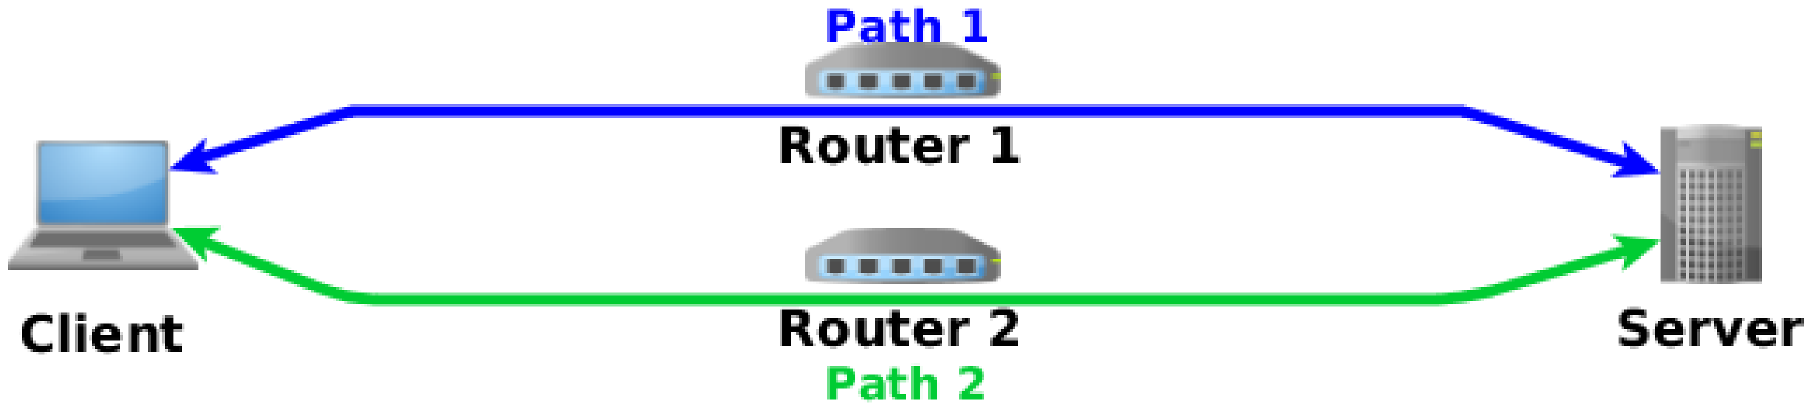

Evaluating Bandwidth Aggregation

- **20 MB Download**
  - Over a single stream
  - Collect the transfer time

Evaluating Bandwidth Aggregation

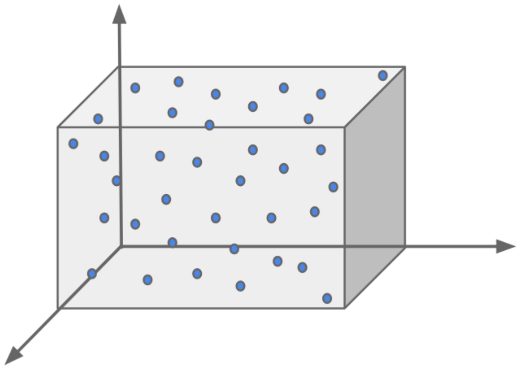

- **20 MB Download**
  - Over a single stream
  - Collect the transfer time
- **Experimental design, WSP algorithm**

Evaluating Bandwidth Aggregation

CoNEXT'17

- **20 MB Download**
  - Over a single stream
  - Collect the transfer time
- **Experimental design, WSP algorithm**
- **2x253 network scenarios**
  - Vary the initial path

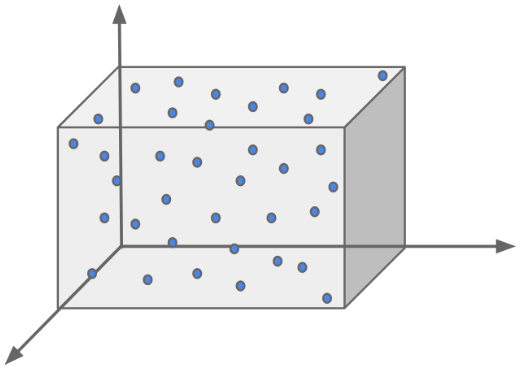

Factor	Minimum	Maximum
Capacity [Mbps]	0.1	100
Round-Trip-Time [ms]	0	50
Queuing Delay [ms]	0	100
Random Loss [%]	0	2.5

Actual Multipath Benefit

- **Experimental Aggregation Benefit**
  - Multipath QUIC/TCP vs. single-path QUIC/TCP

Actual Multipath Benefit

- **Experimental Aggregation Benefit**
  - Multipath QUIC/TCP vs. single-path QUIC/TCP

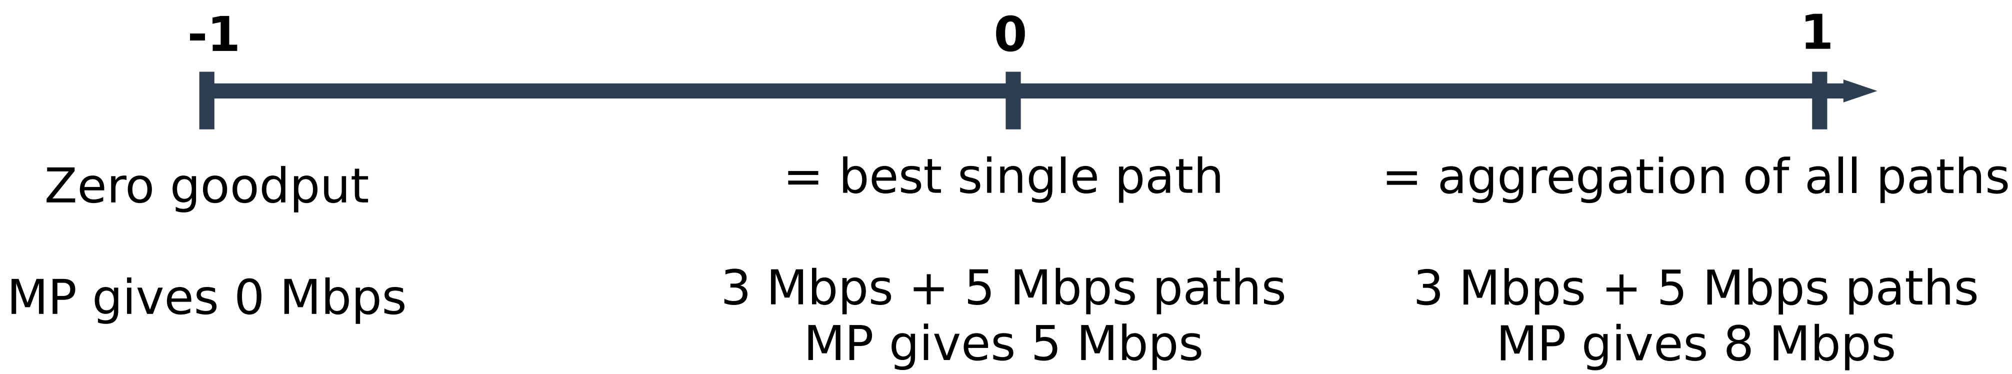

Actual Multipath Benefit

- **Experimental Aggregation Benefit**
  - Multipath QUIC/TCP vs. single-path QUIC/TCP

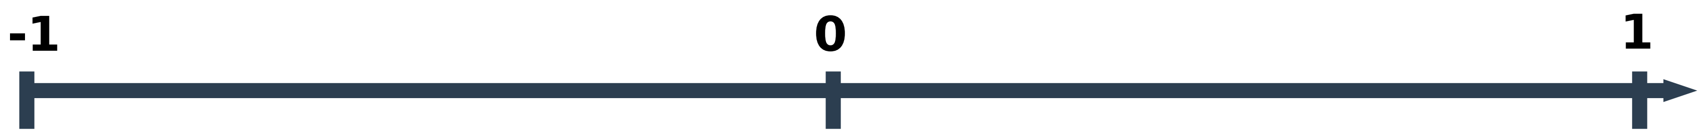

Zero goodput

MP gives 0 Mbps

= best single path

3 Mbps + 5 Mbps paths
MP gives 5 Mbps

= aggregation of all paths

3 Mbps + 5 Mbps paths
MP gives 8 Mbps

- **Results depends on the first path used**
  - Handshake latency over initial path

Benefits of Multipath – No Loss

CoNEXT'17

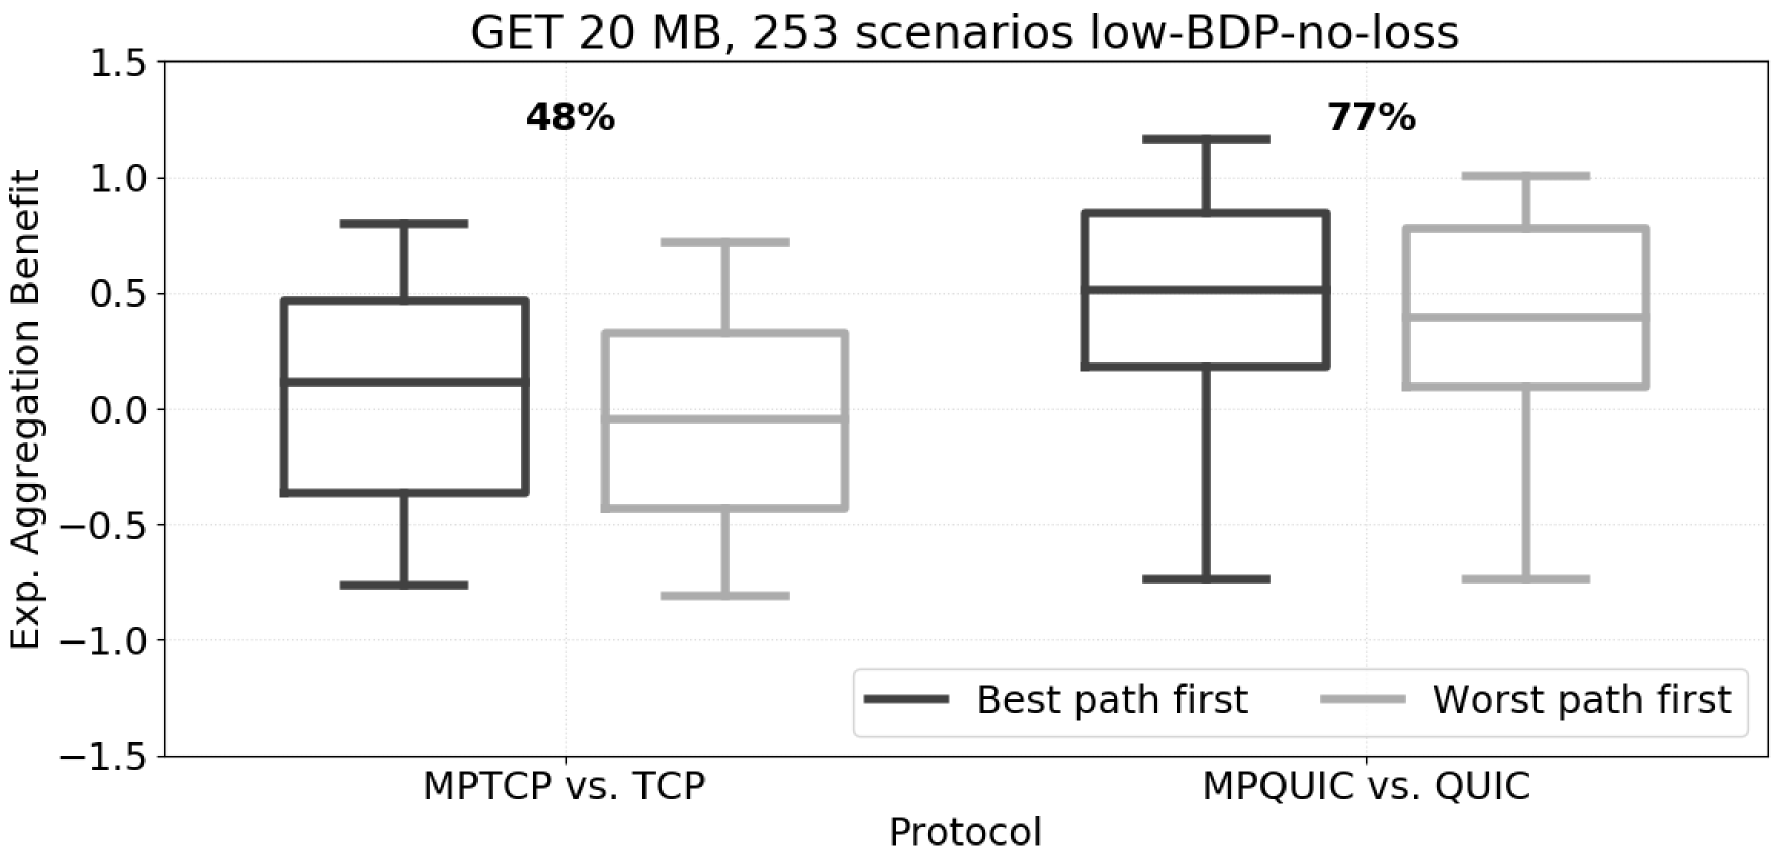

Benefits of Multipath – No Loss

CoNEXT'17

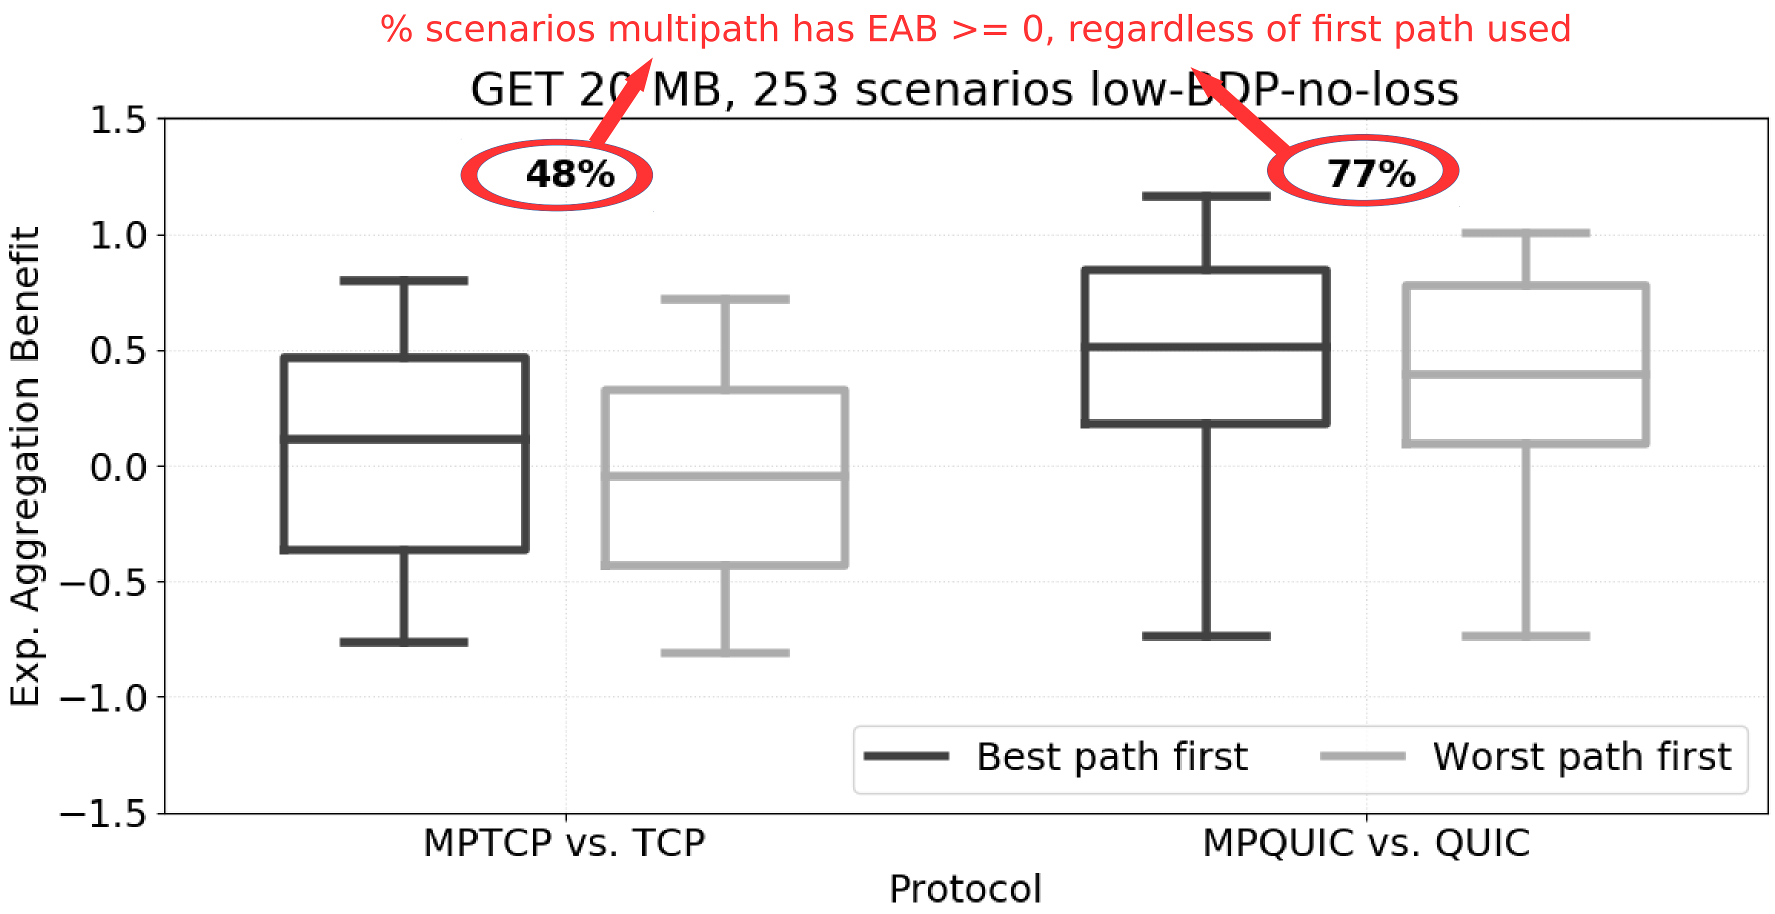

Benefits of Multipath – Losses

CoNEXT'17

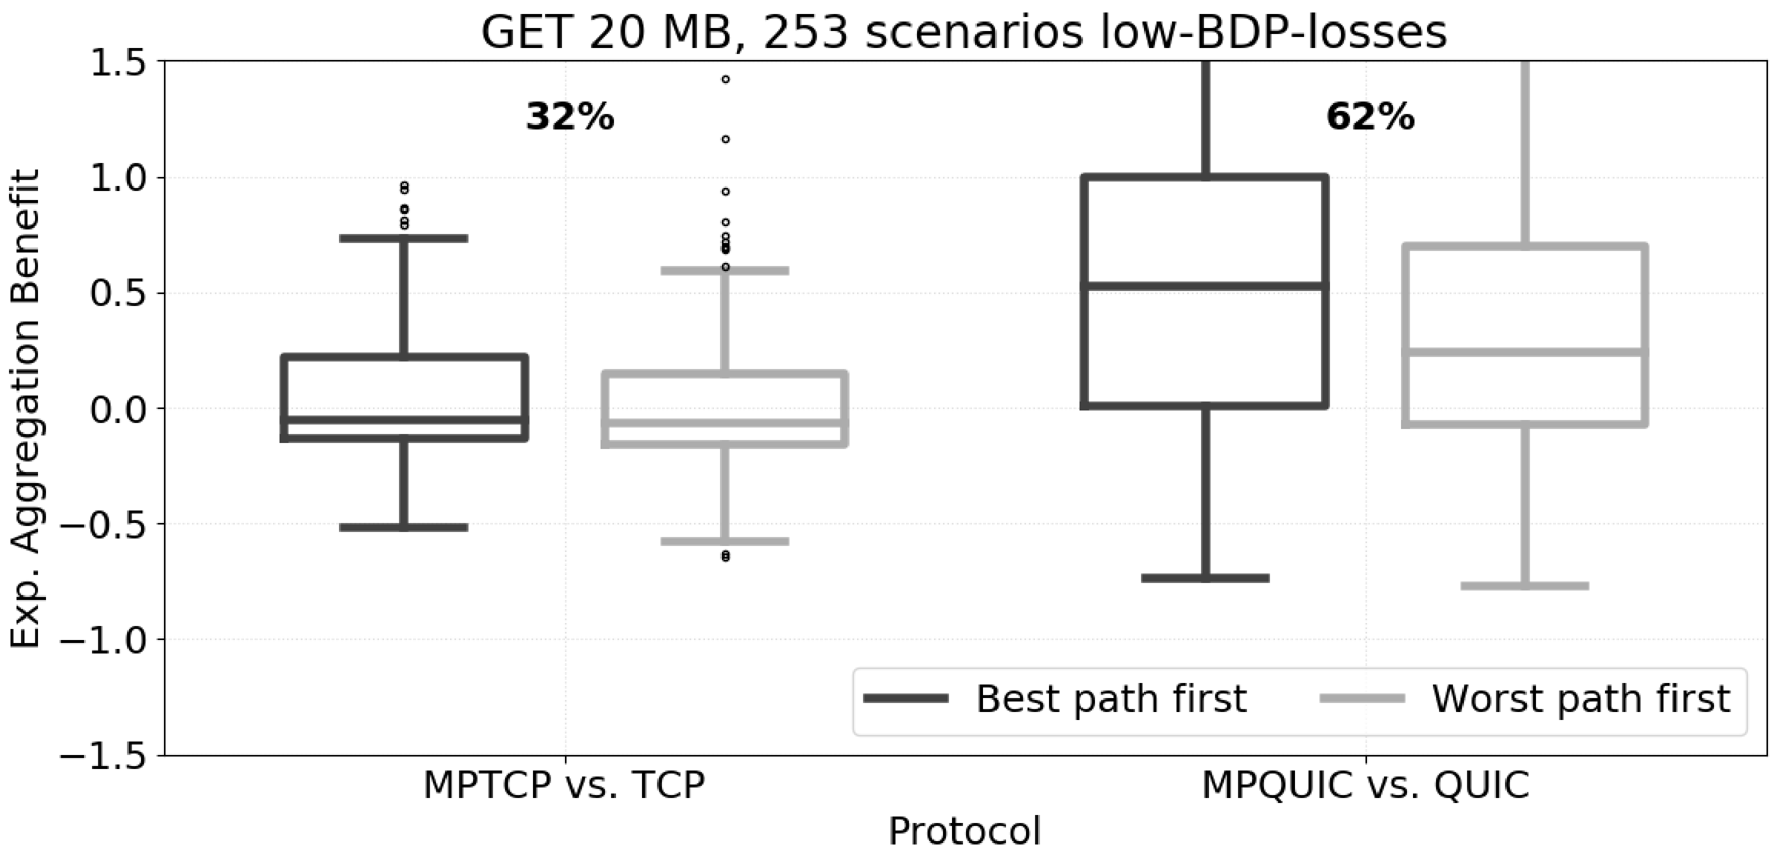

And with Real Networks?

ICC'18

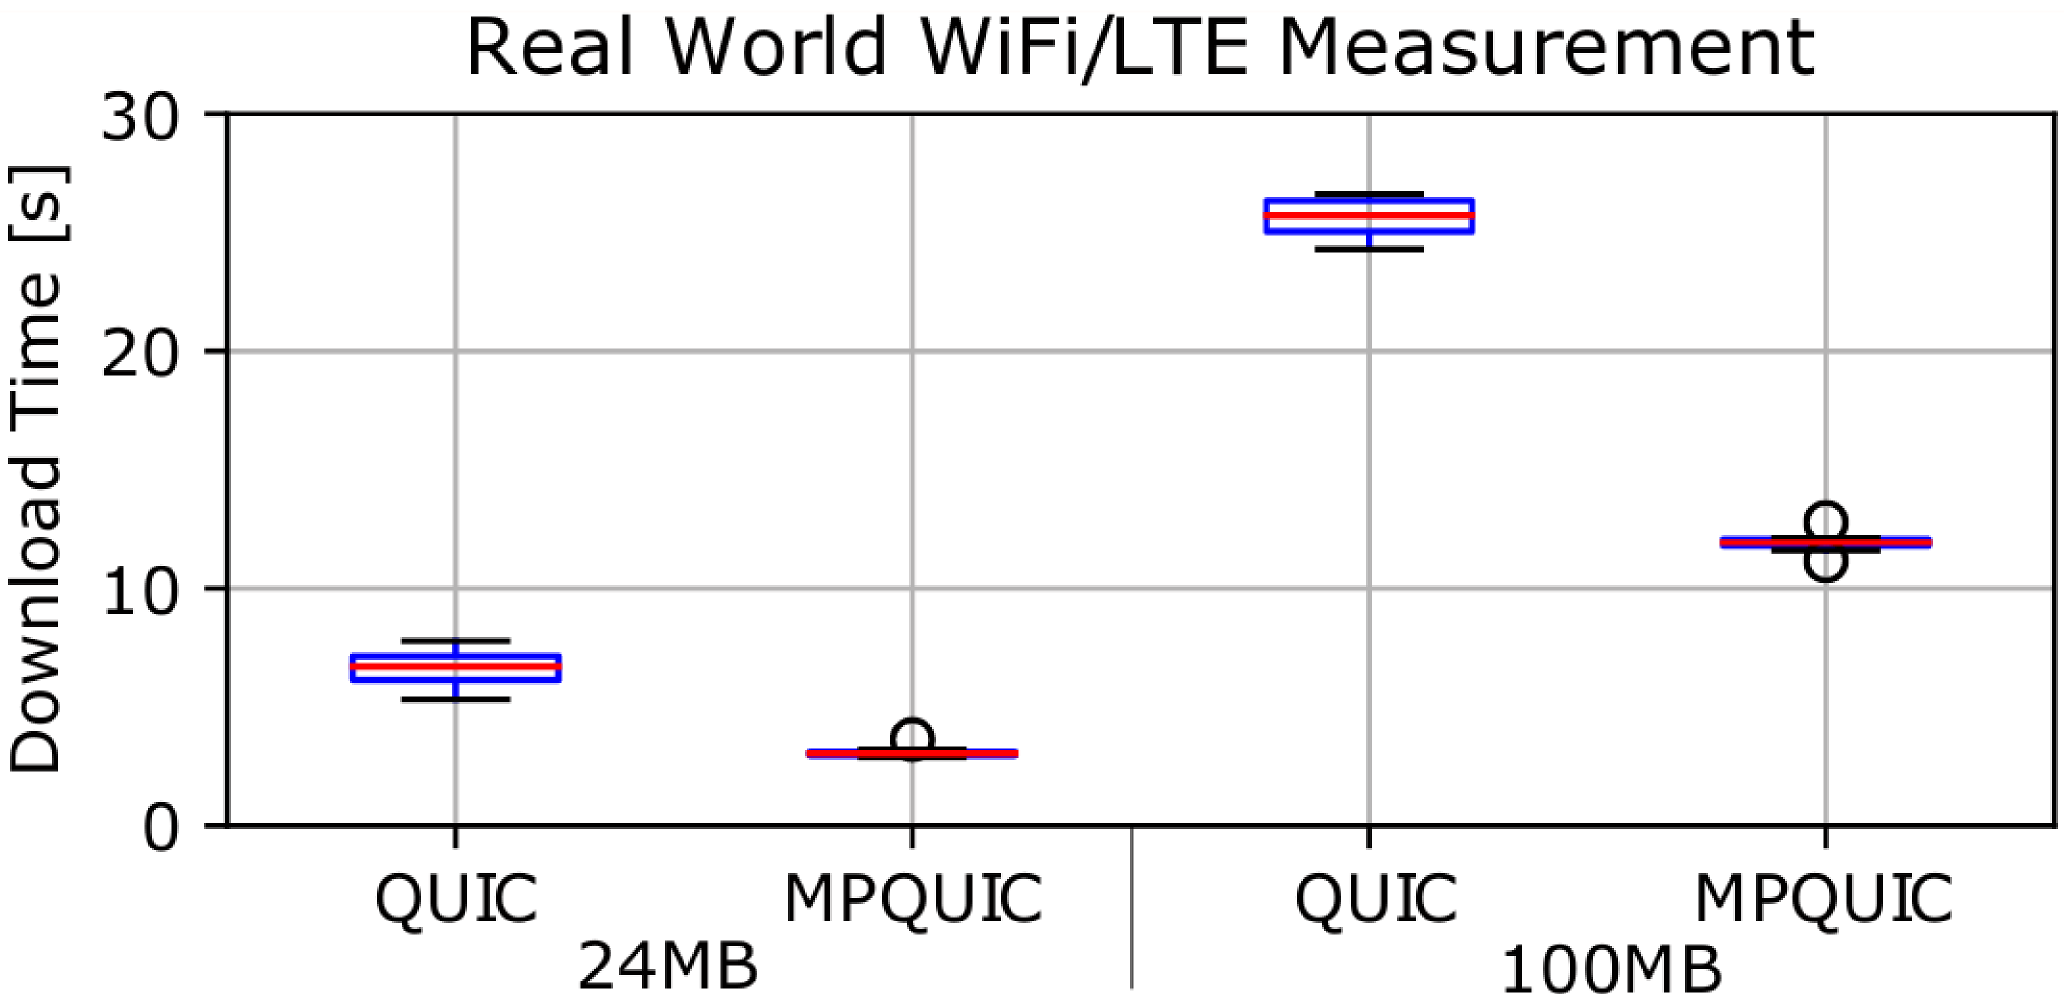

Network Handover Support

- **Apple MPTCP deployment mainly for handover**
  - Main use case for Siri

Network Handover Support

- **Apple MPTCP deployment mainly for handover**
  - Main use case for Siri
- **Request/Response traffic**
  - 750 bytes request/responses every 1/3 s
  - Measure delay seen by client

15ms RTT, 100% loss after 3 s

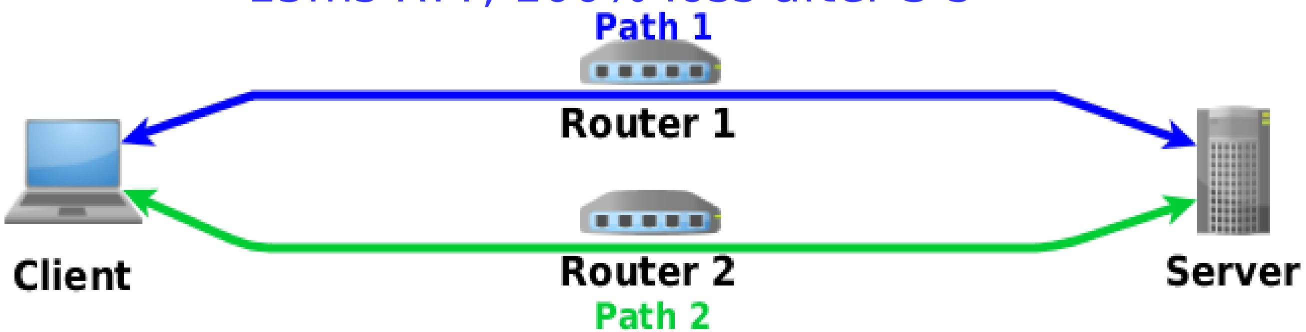

25ms RTT

Multipath TCP Handover

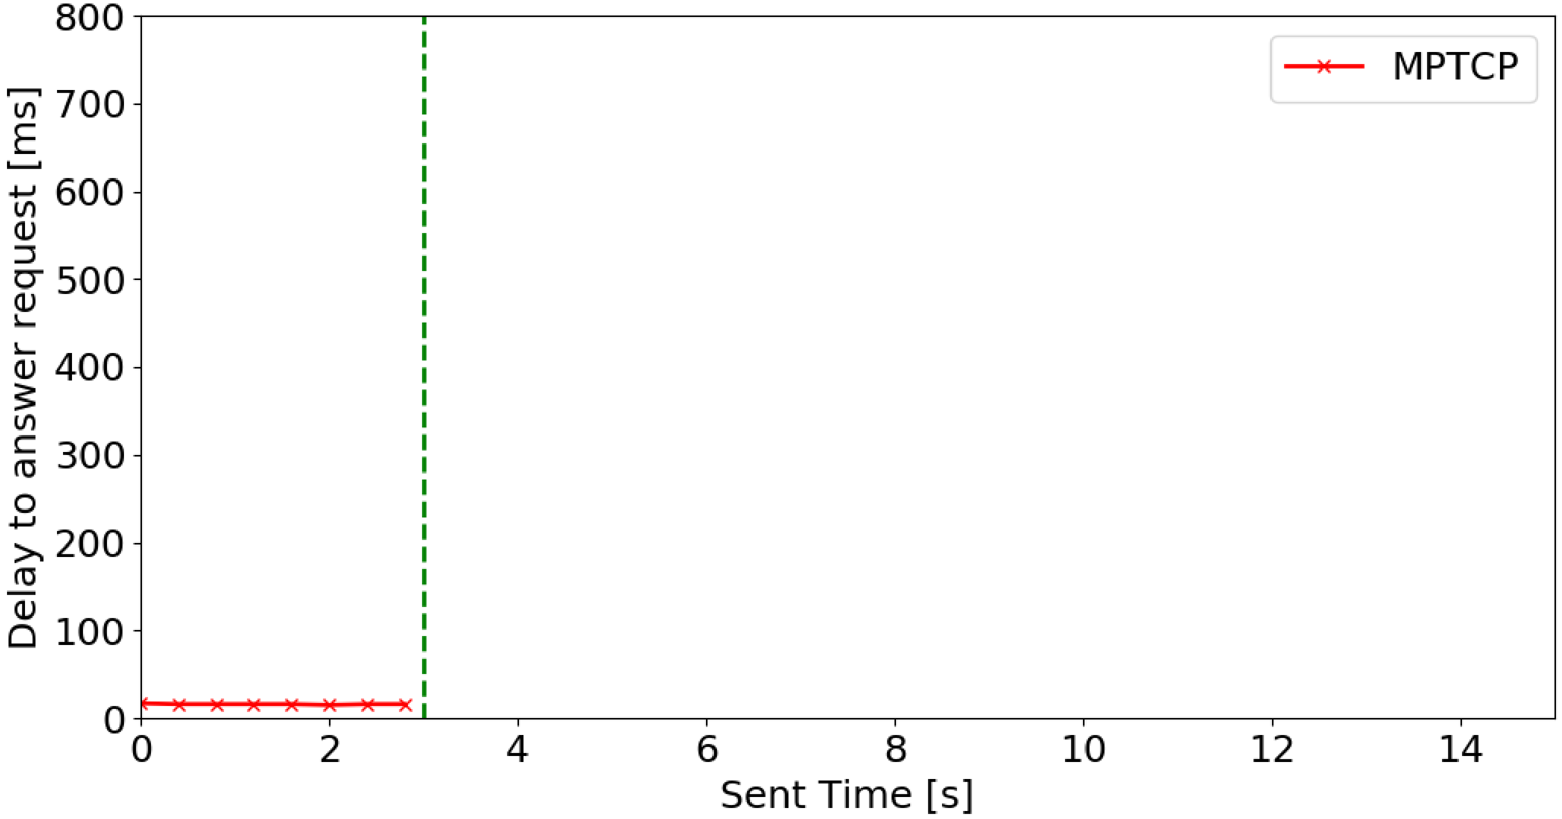

Multipath TCP Handover

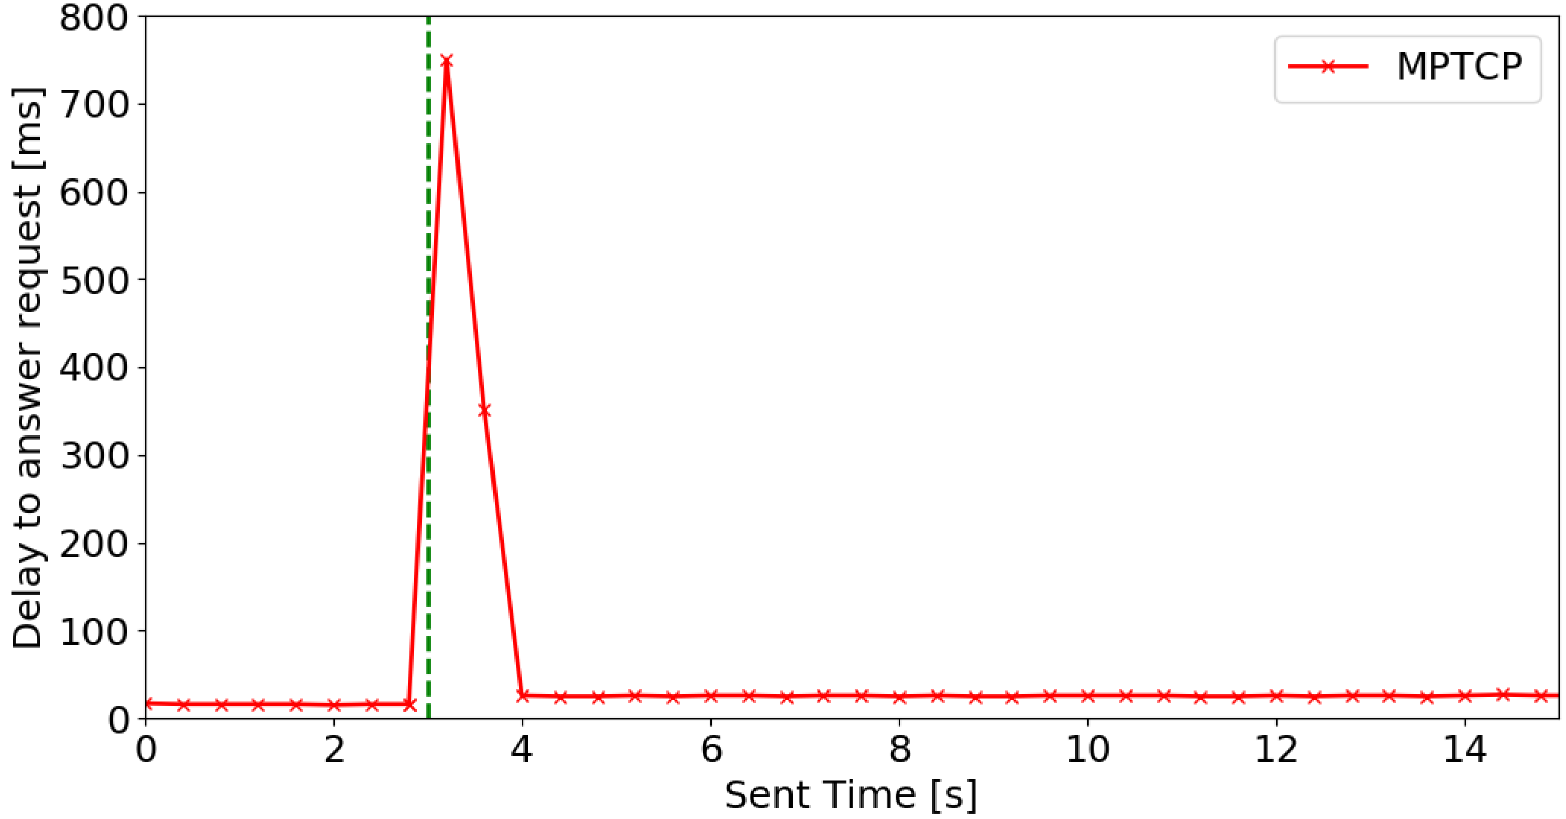

MPTCP Handover Explained

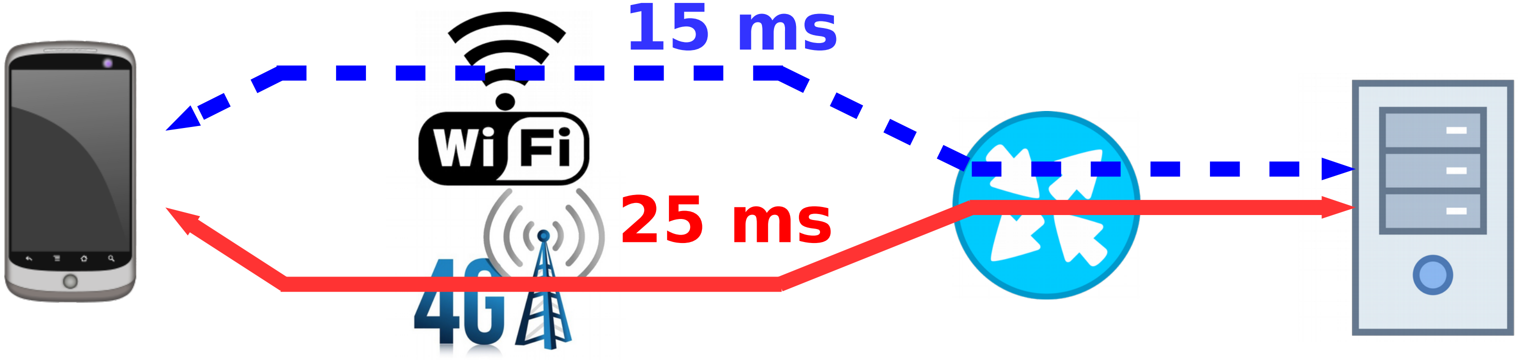

MPTCP Handover Explained

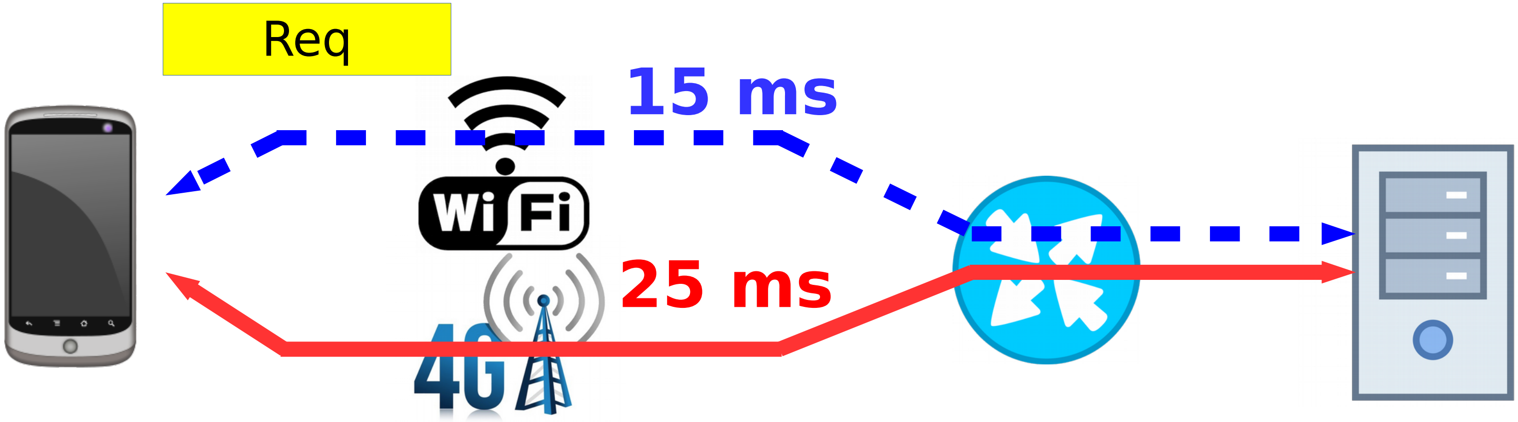

MPTCP Handover Explained

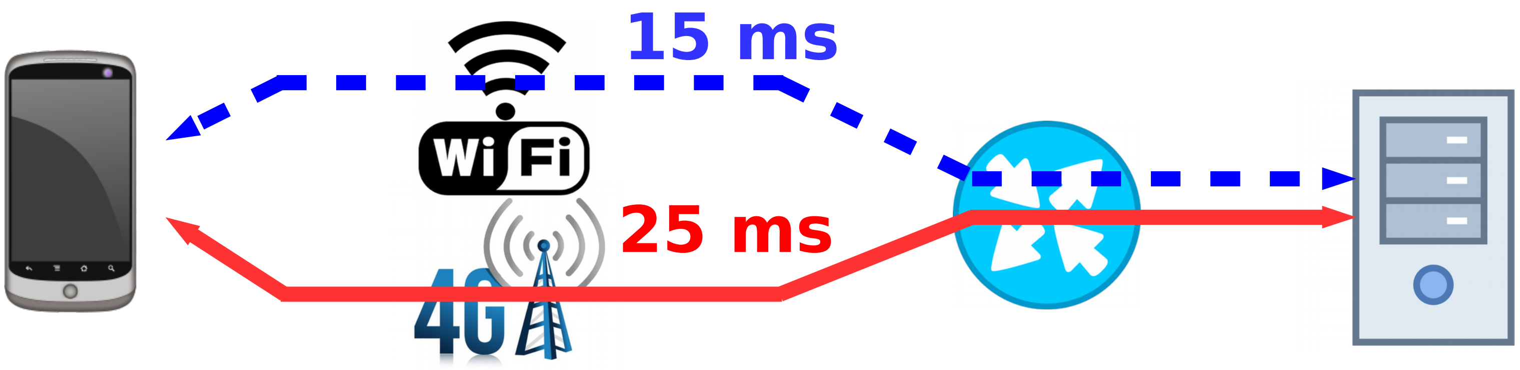

MPTCP Handover Explained

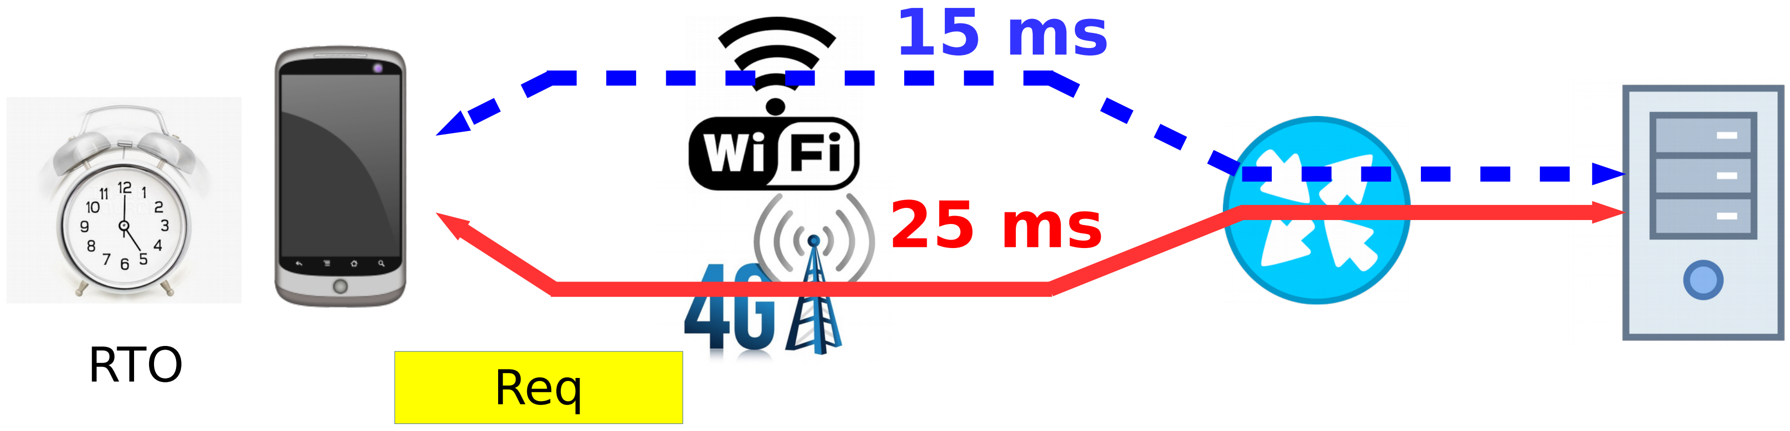

MPTCP Handover Explained

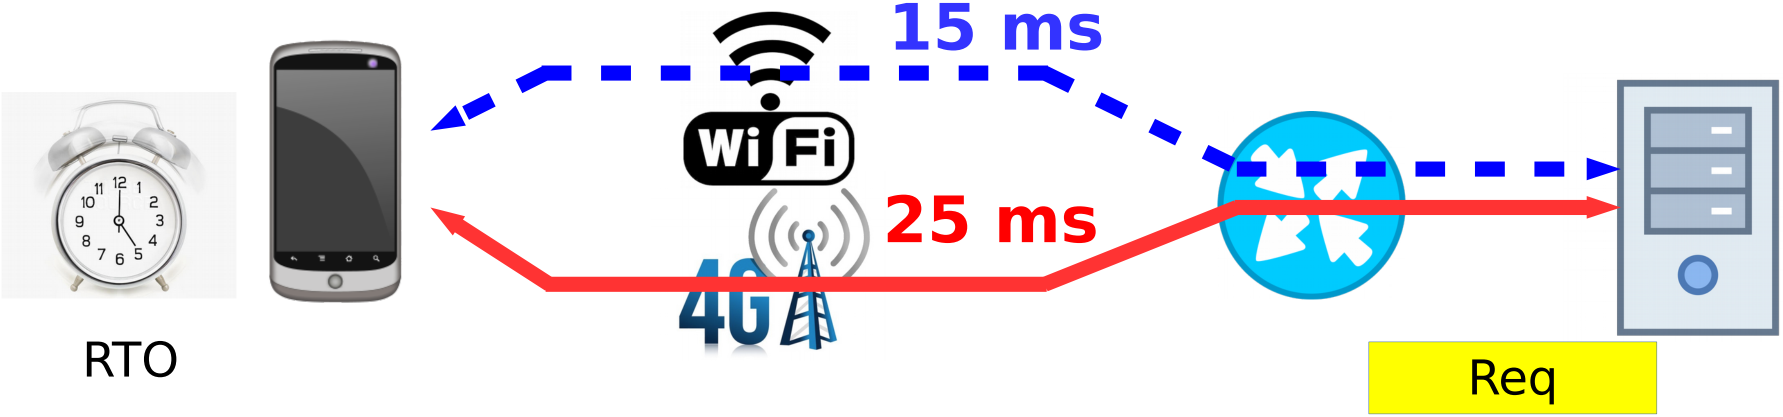

MPTCP Handover Explained

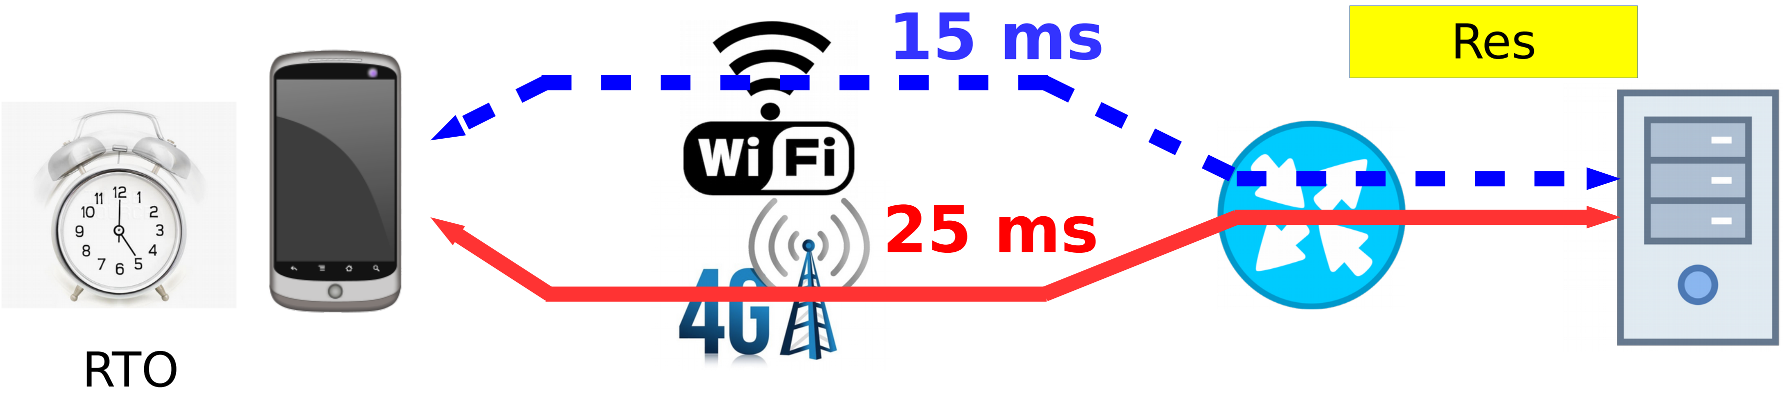

MPTCP Handover Explained

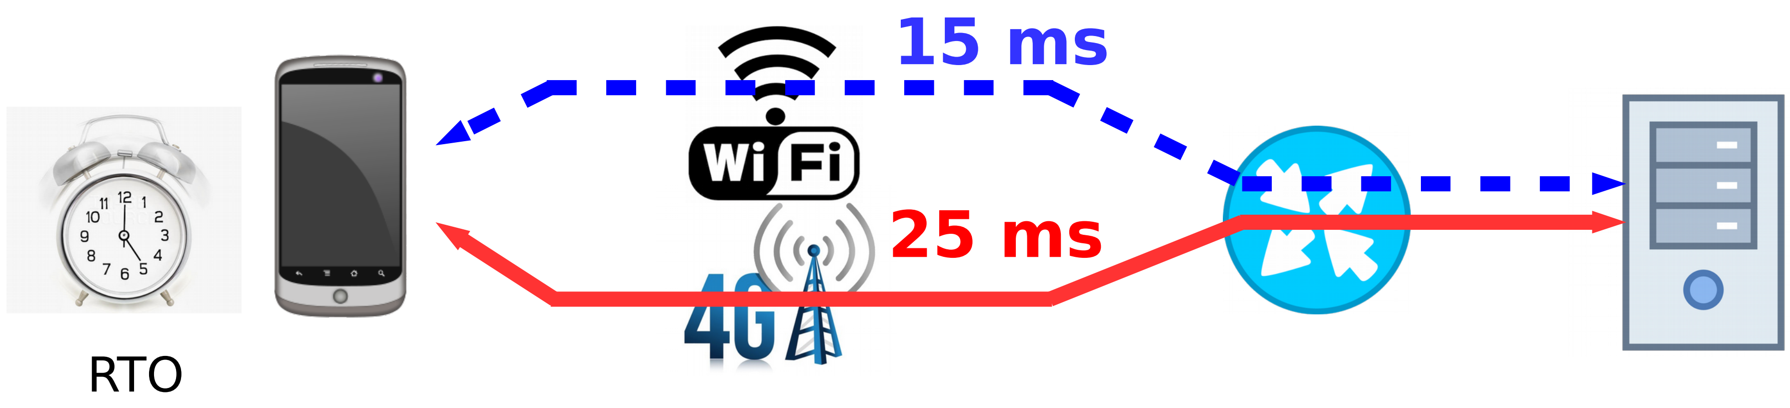

MPTCP Handover Explained

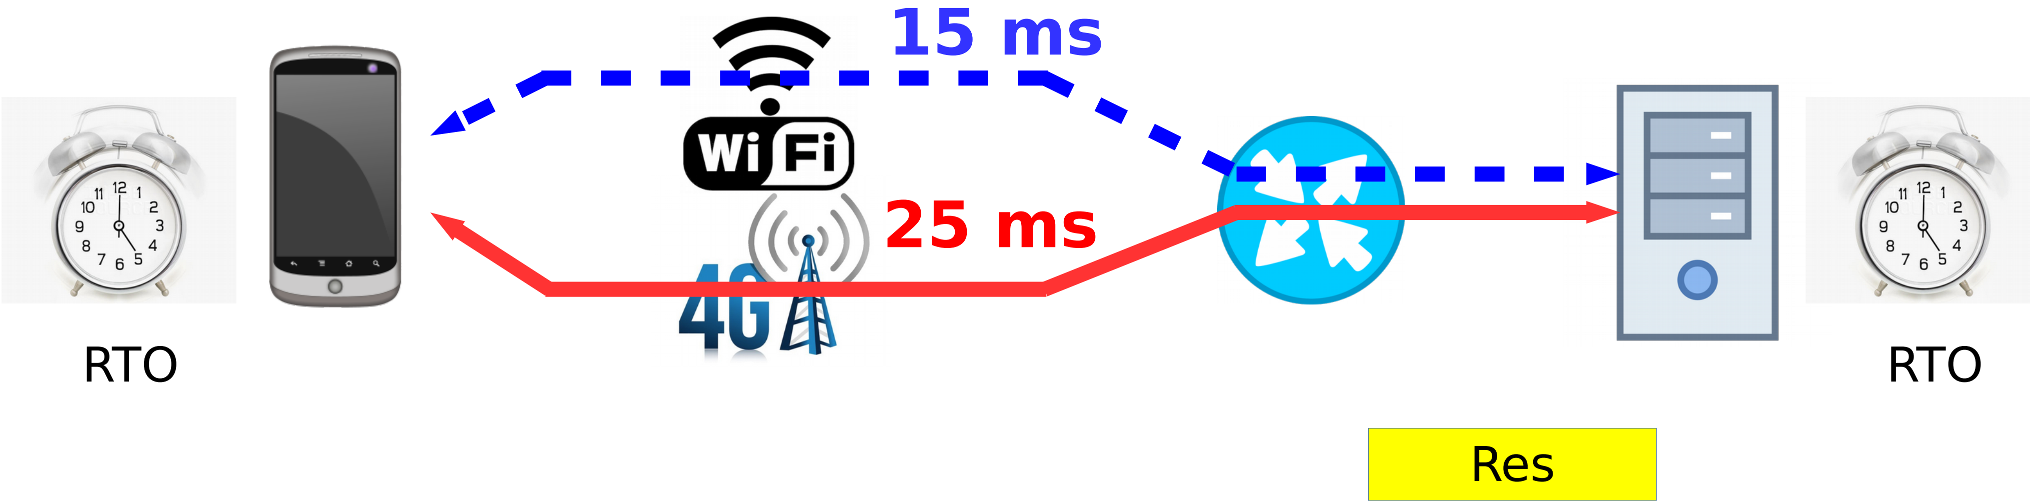

MPTCP Handover Explained

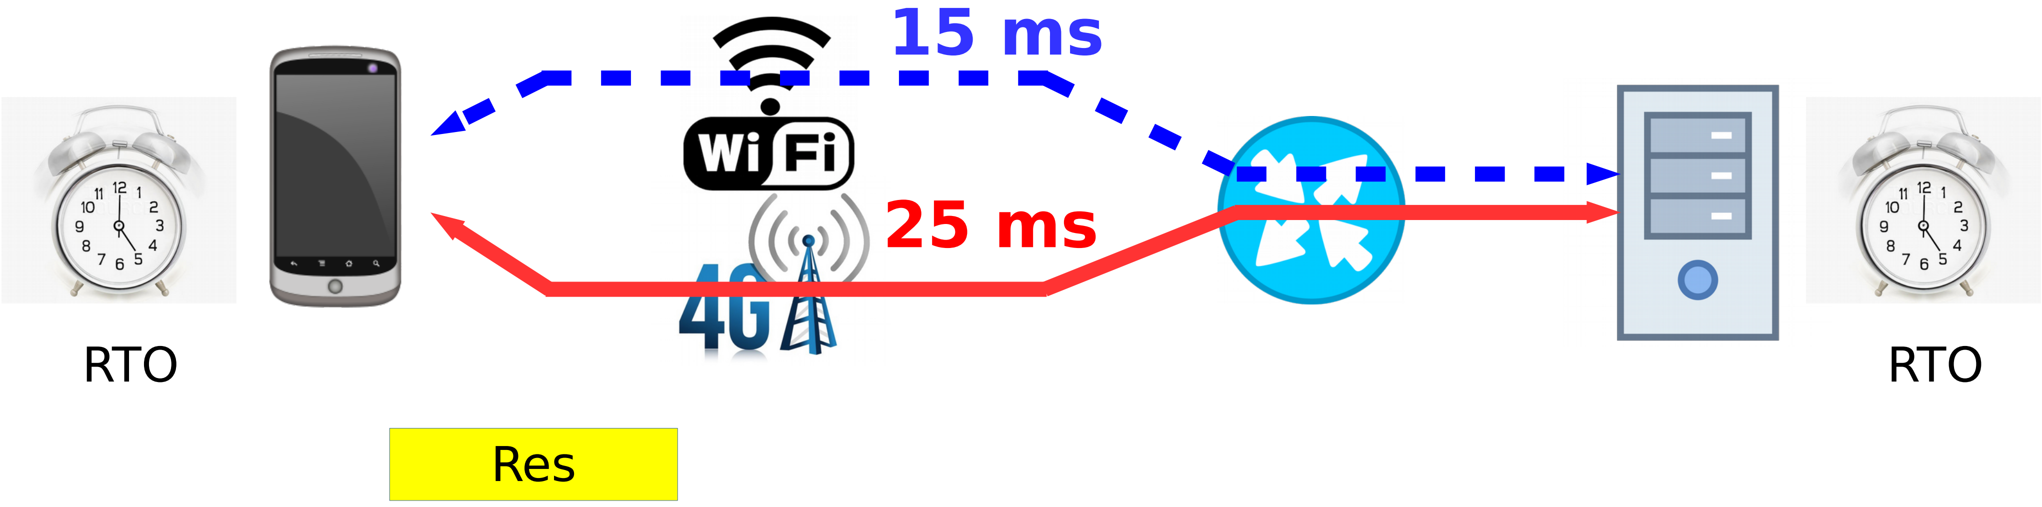

What about Multipath QUIC?

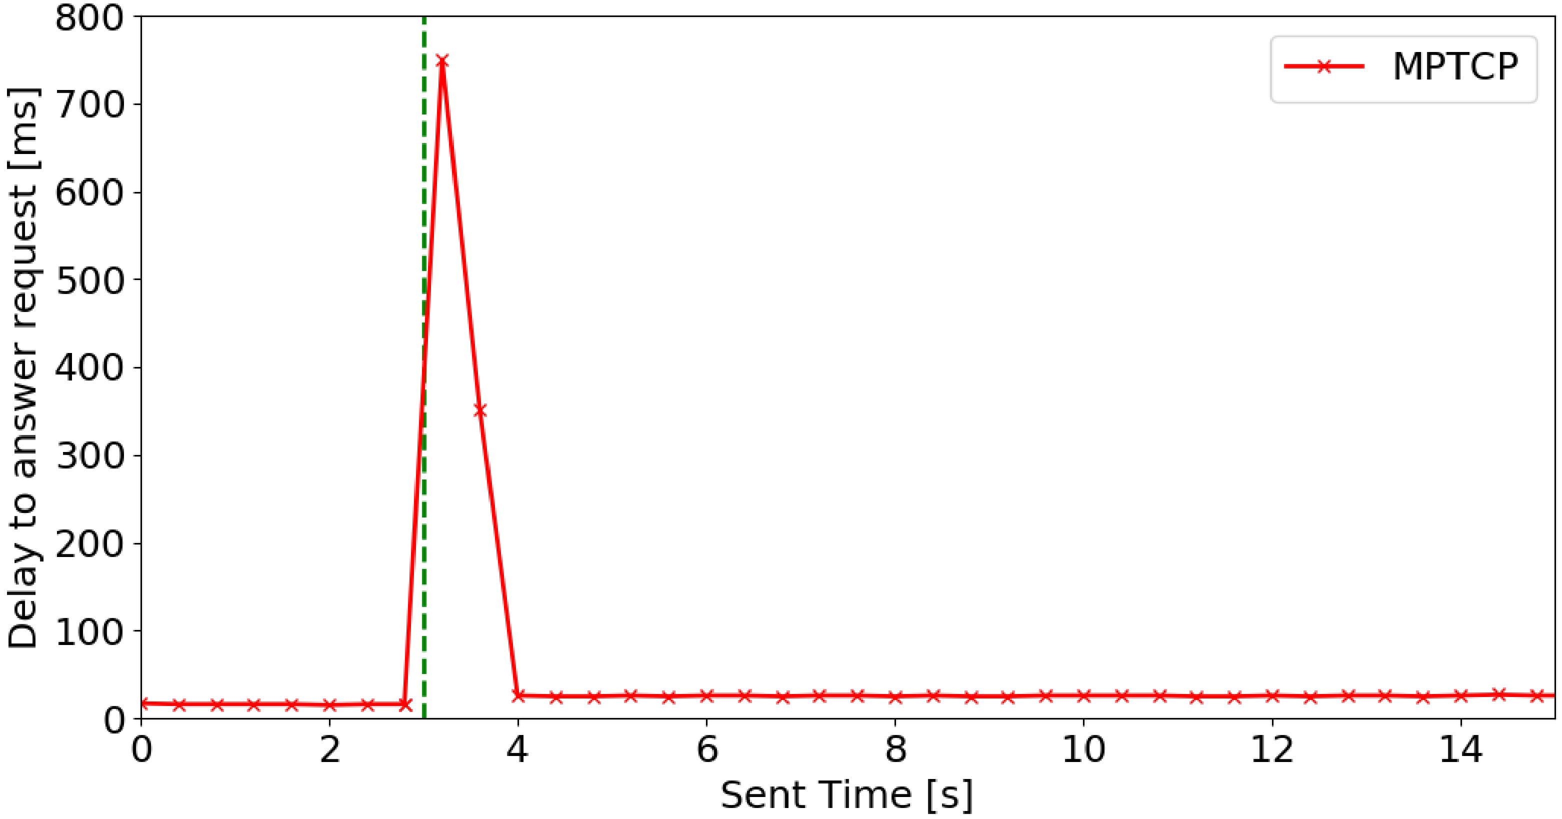

What about Multipath QUIC?

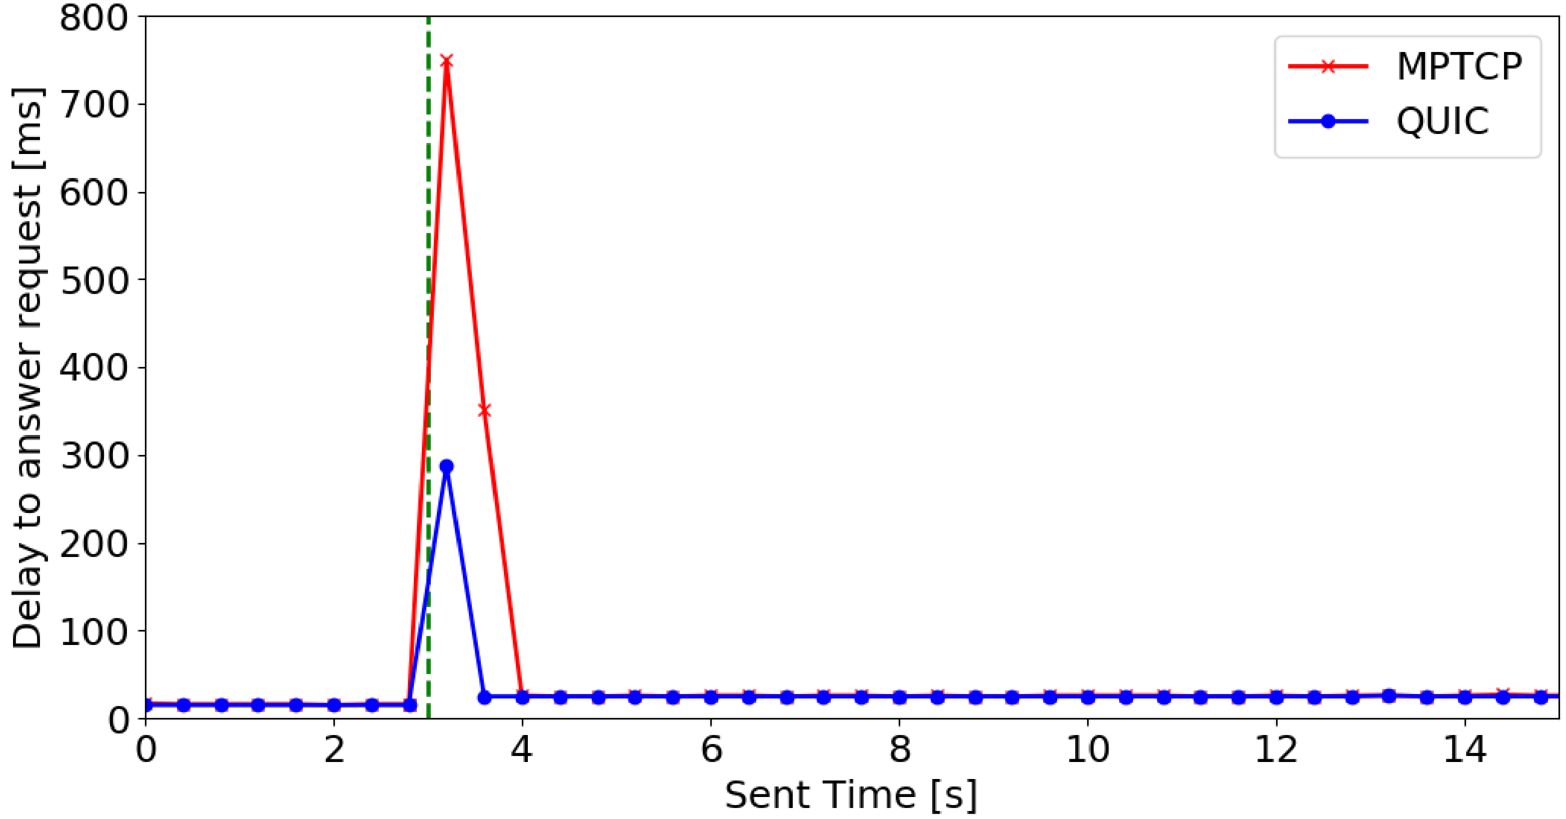

How is it possible?

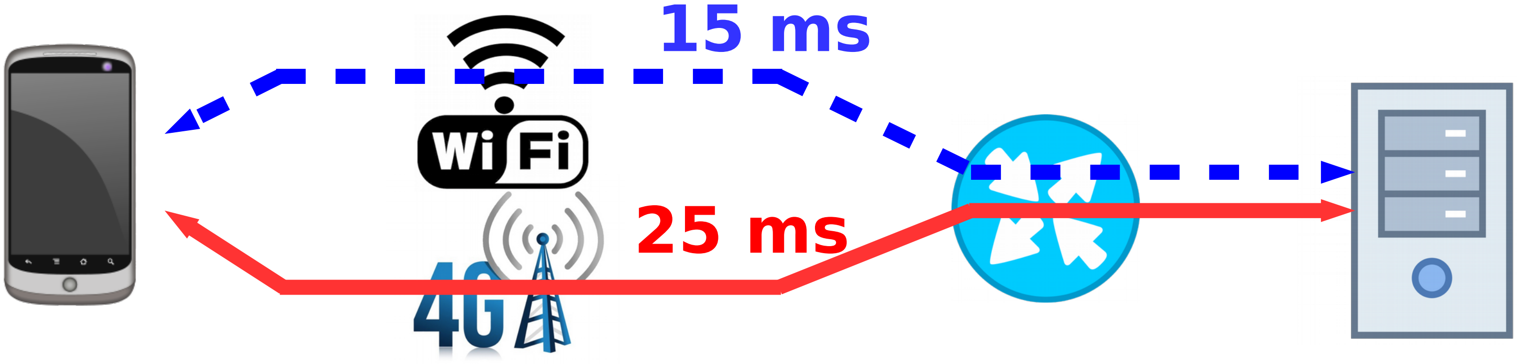

How is it possible?

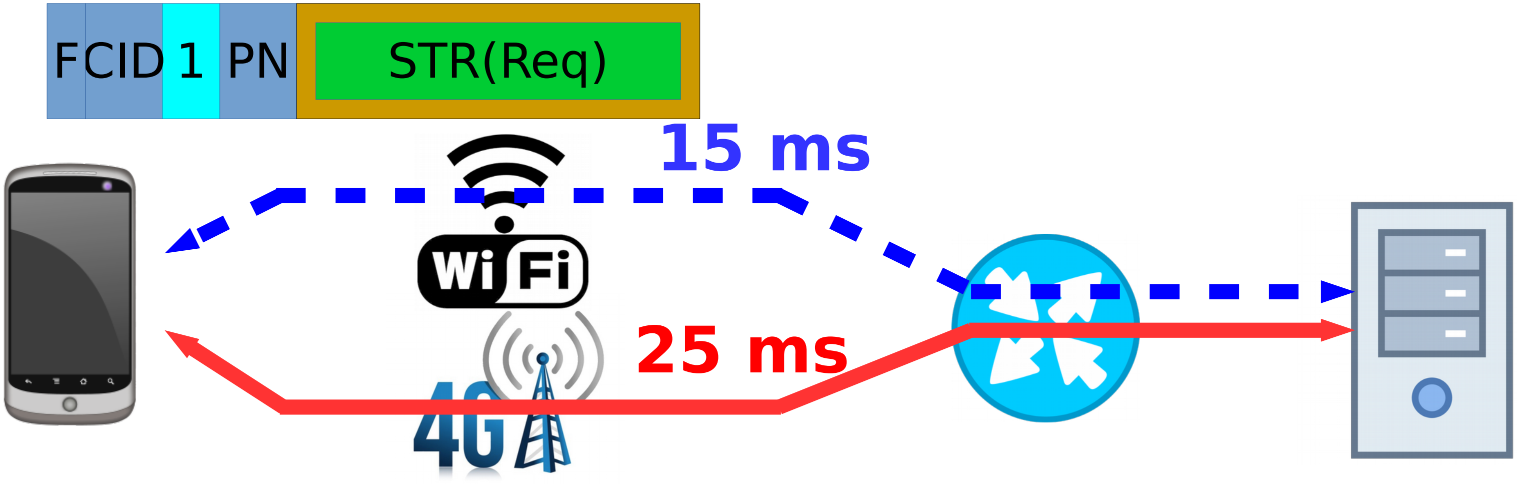

How is it possible?

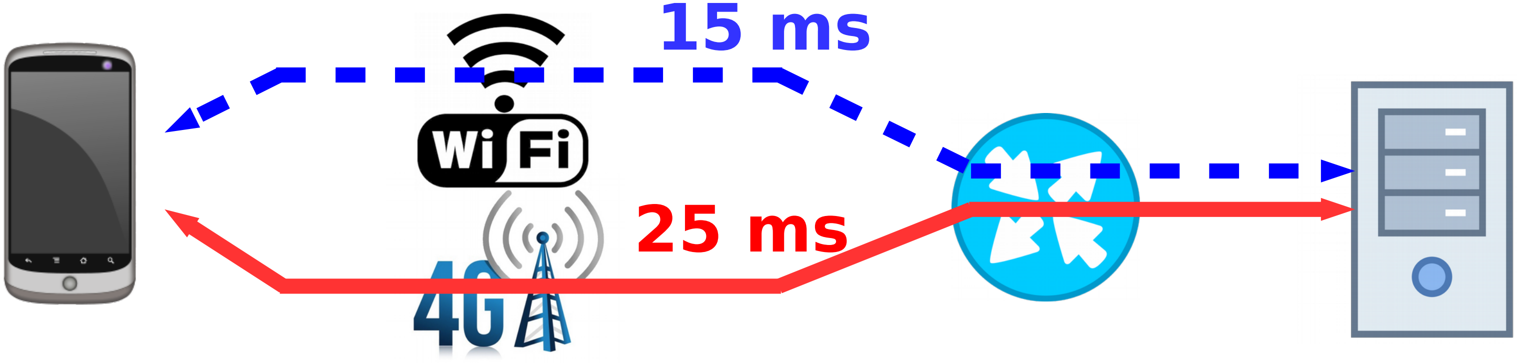

How is it possible?

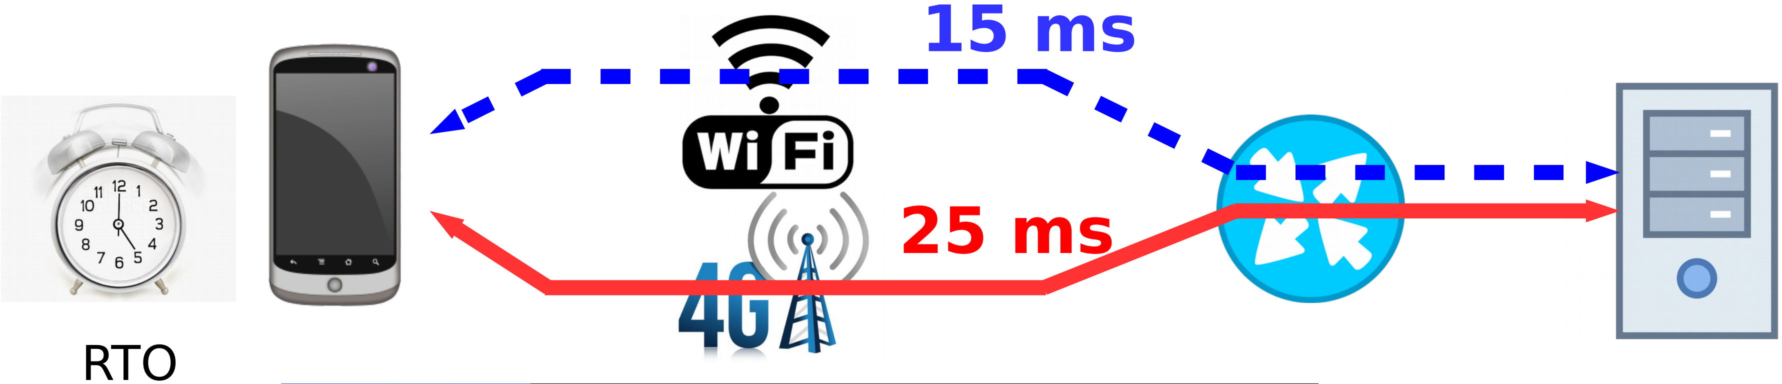

F	CID	2	PN	STR(Req)	PATHS(1 lossy)

How is it possible?

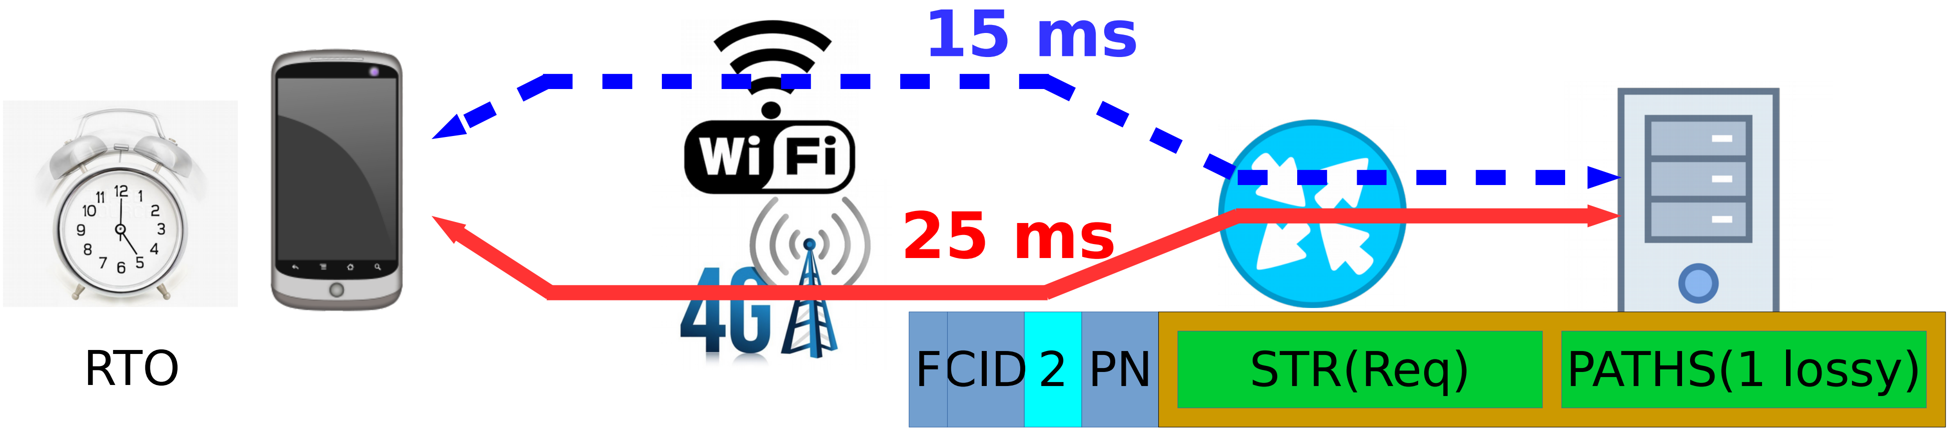

How is it possible?

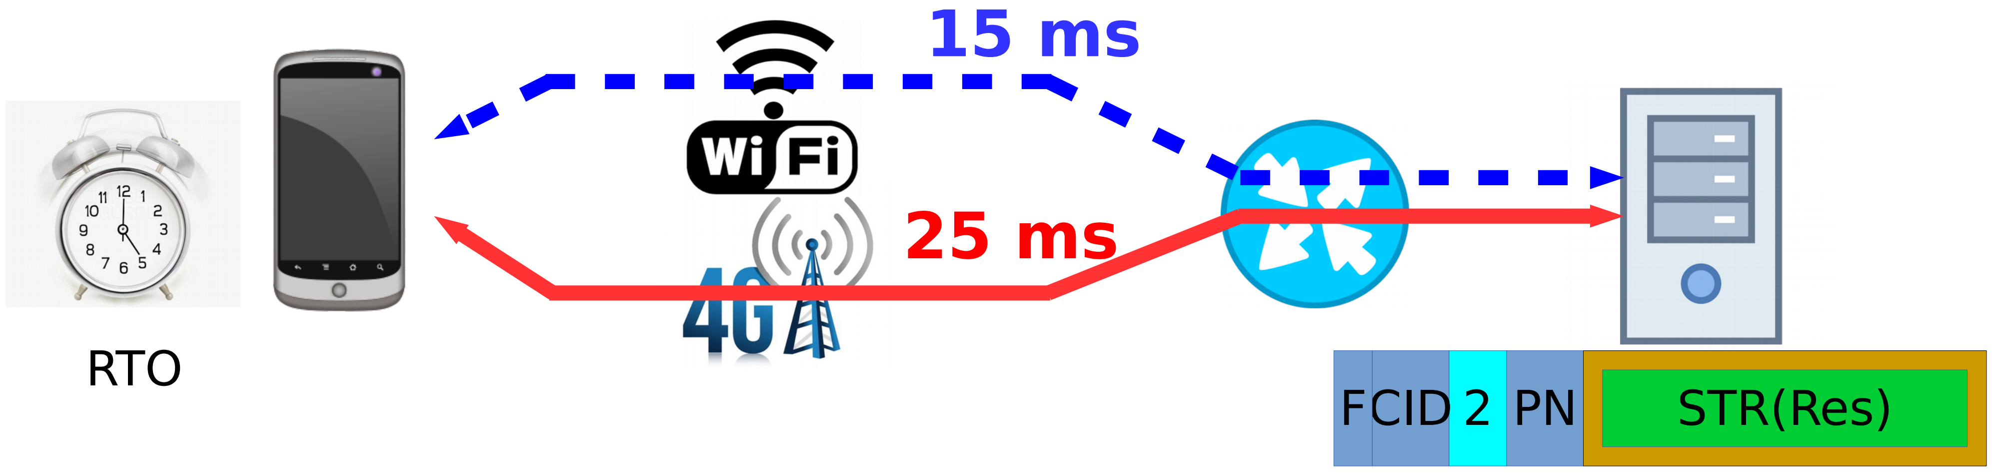

How is it possible?

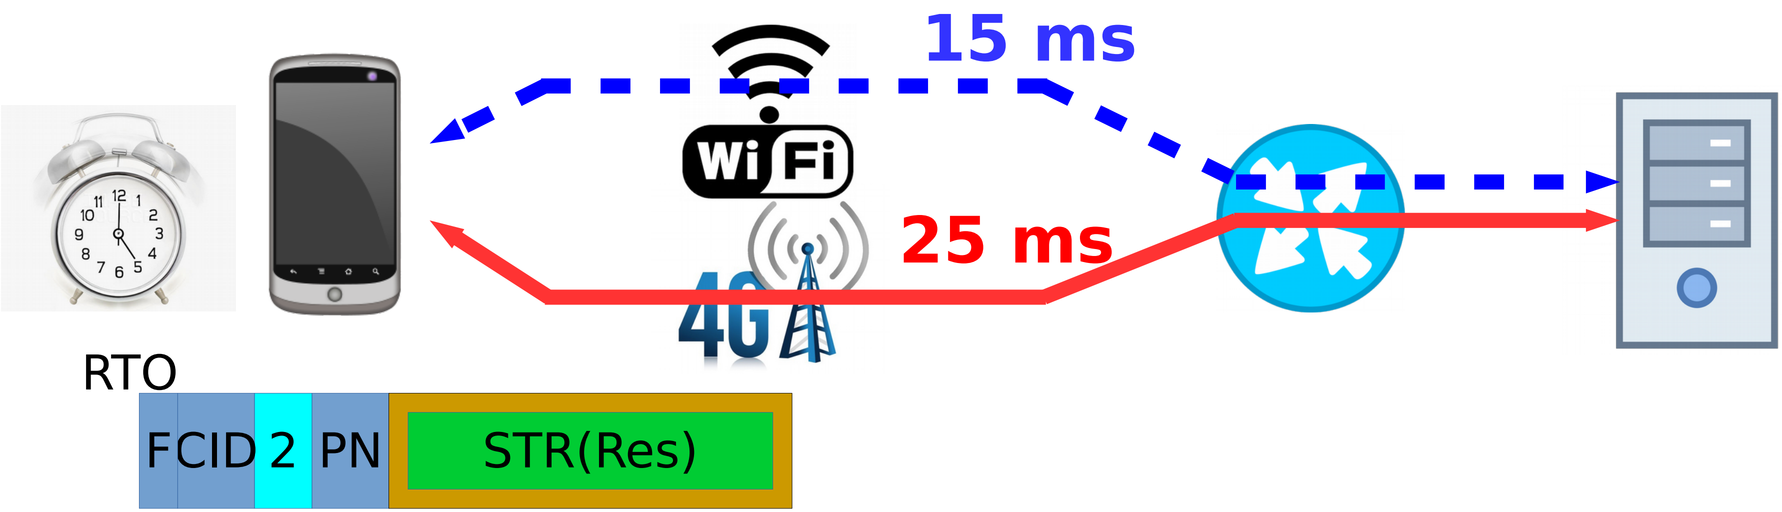

Adapting to IETF QUIC

Issues with (G)QUIC Design

- **Path ID in clear-text public header**
  - Easy to correlate paths :-(
- **IETF QUIC changed a lot in ~2 years**
  - GQUIC very different from current IETF version
  - Source/Destination Connection Idµ
- **Core idea: use CIDs as implicit path ID**

Negotiating Multipath Usage

MPQUIC-ID

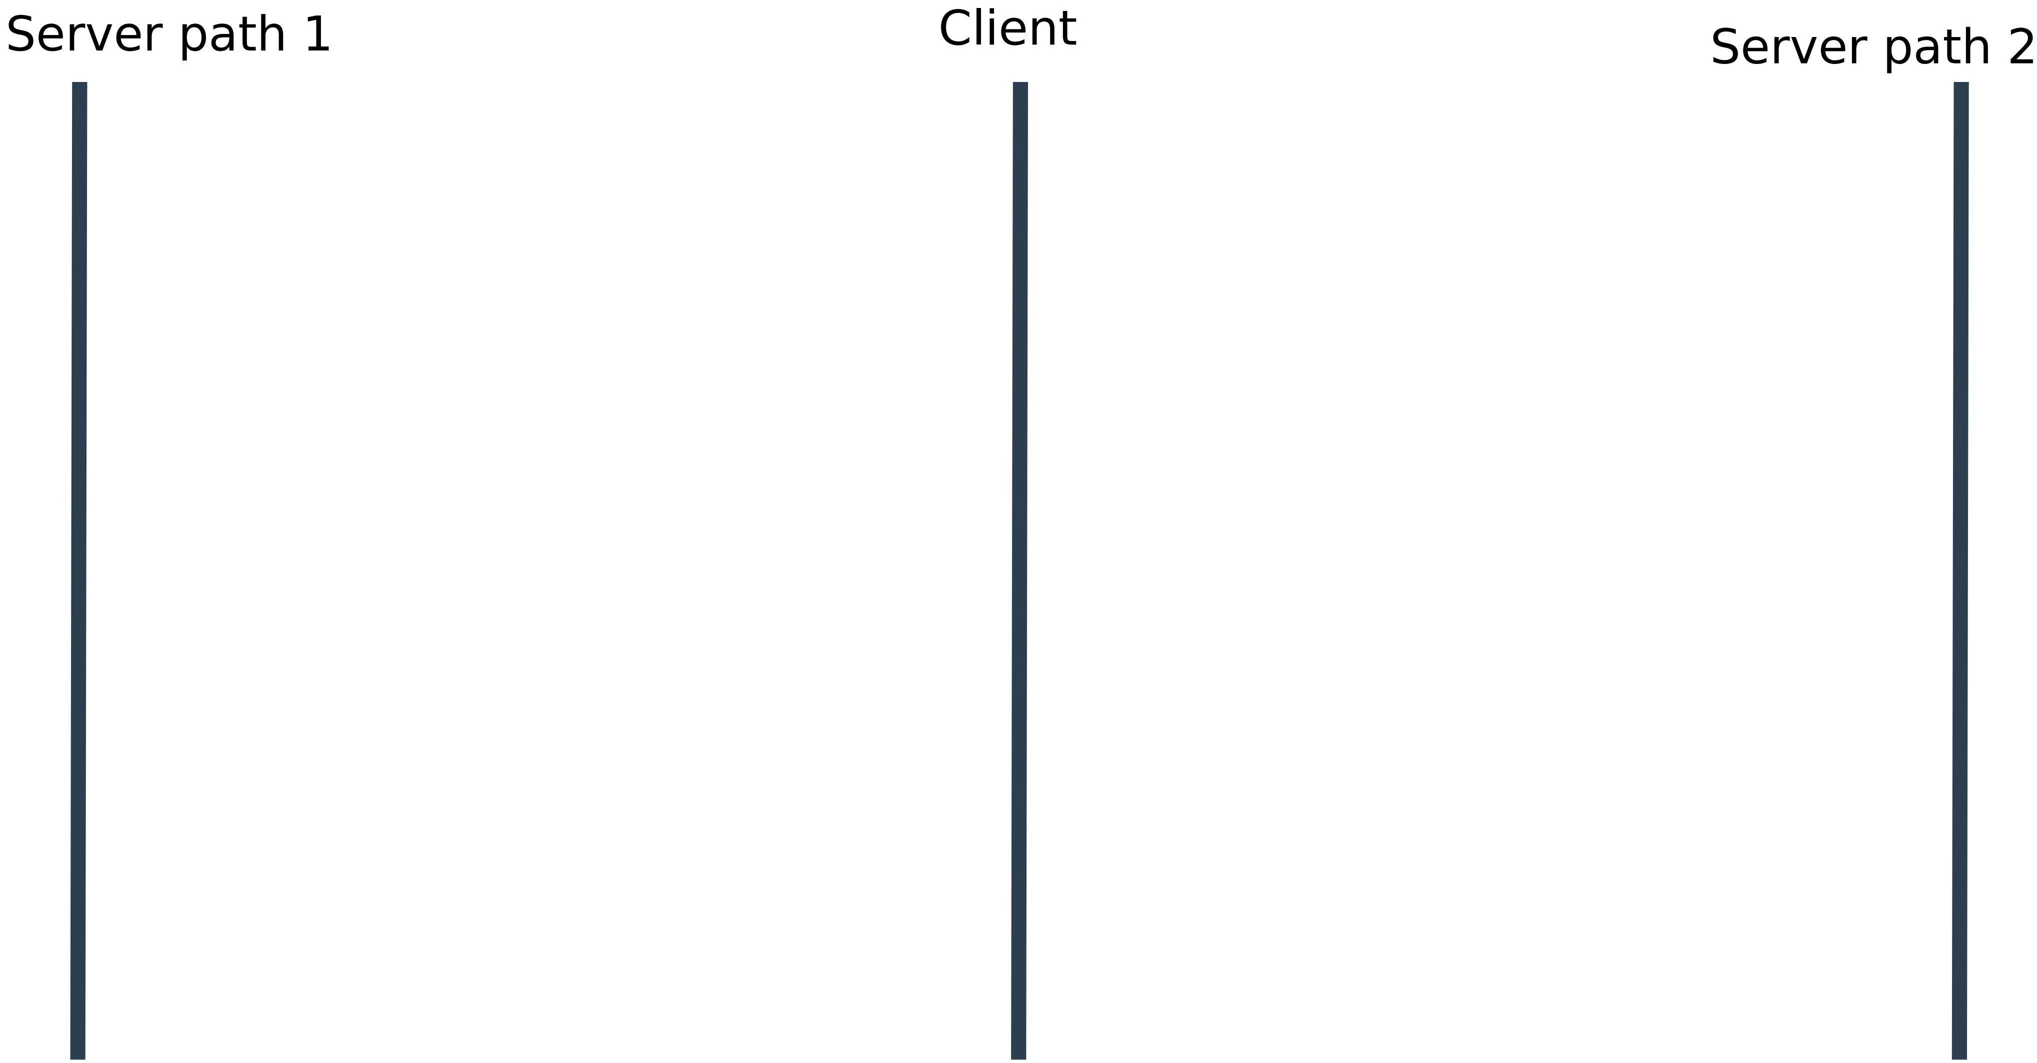

Negotiating Multipath Usage

MPQUIC-ID

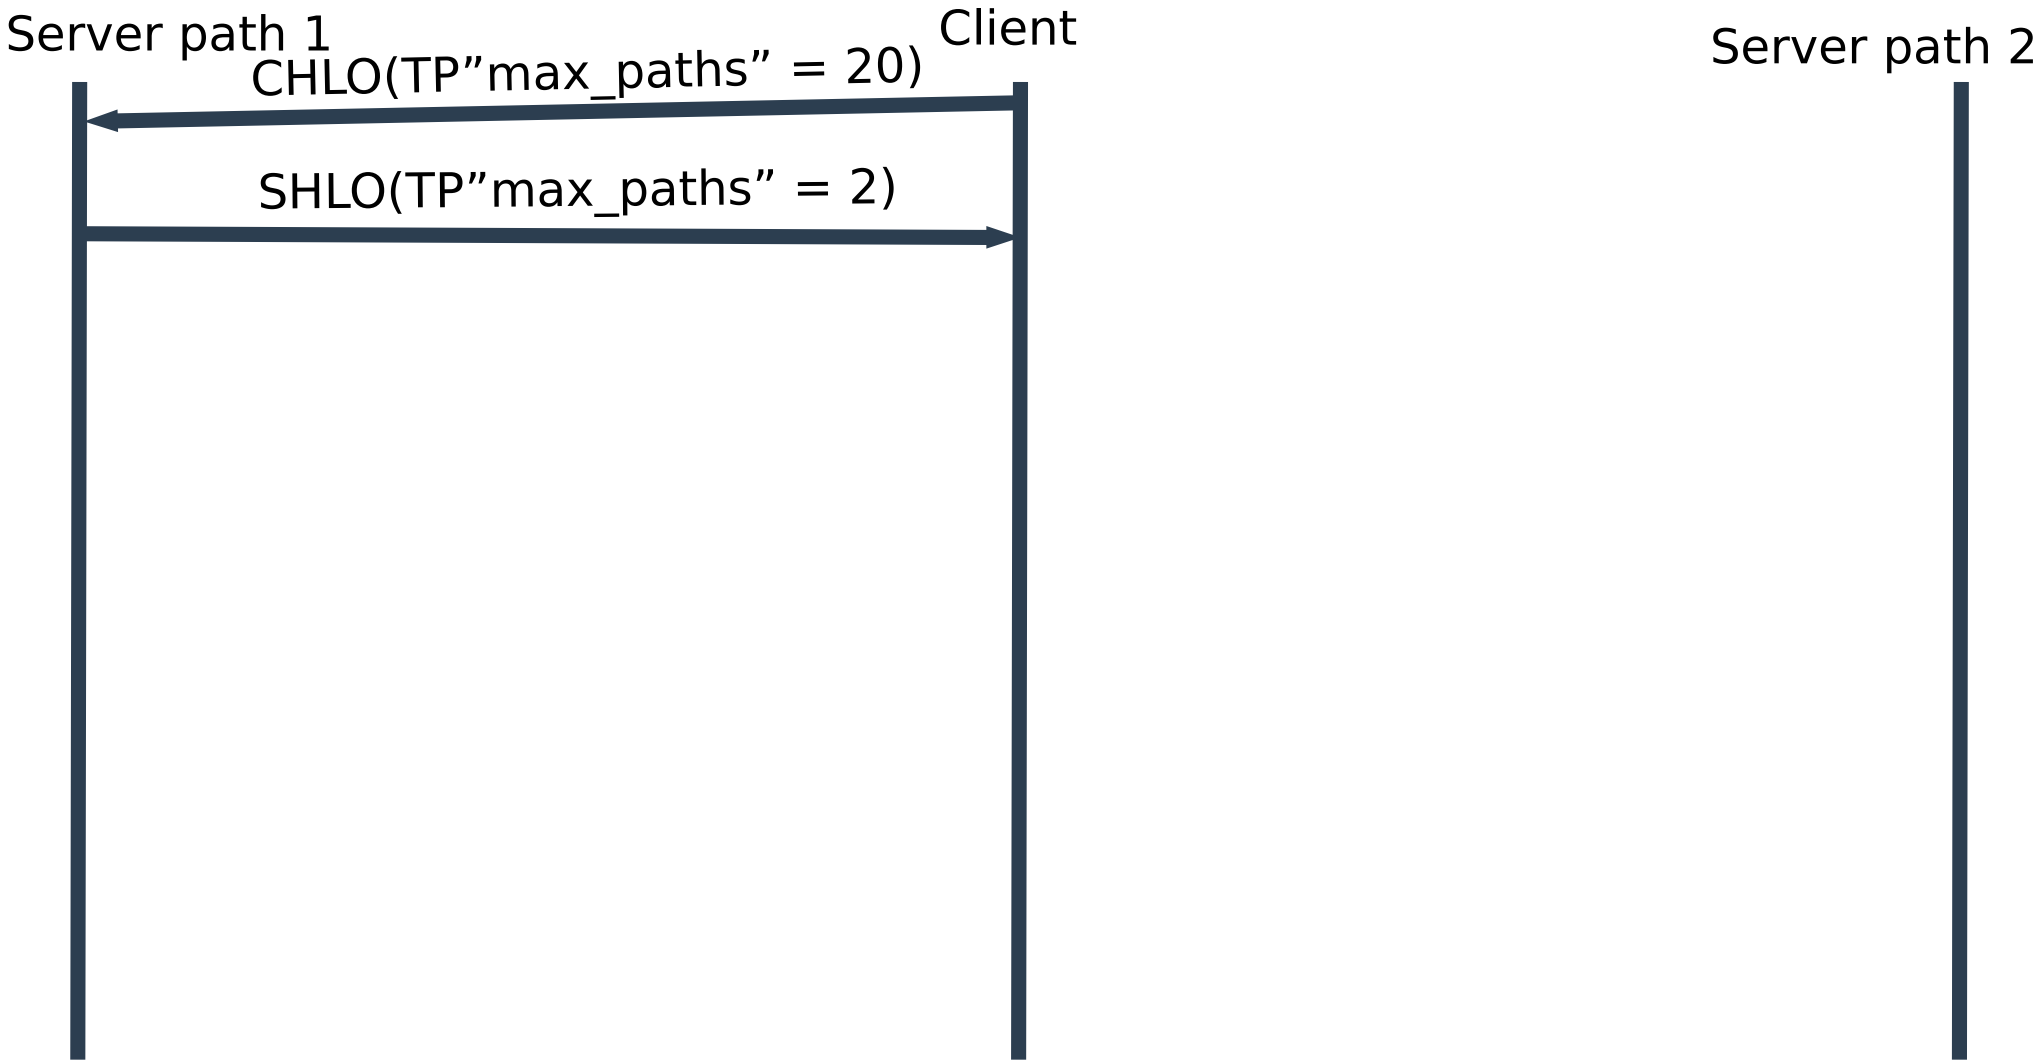

Negotiating Multipath Usage

MPQUIC-ID

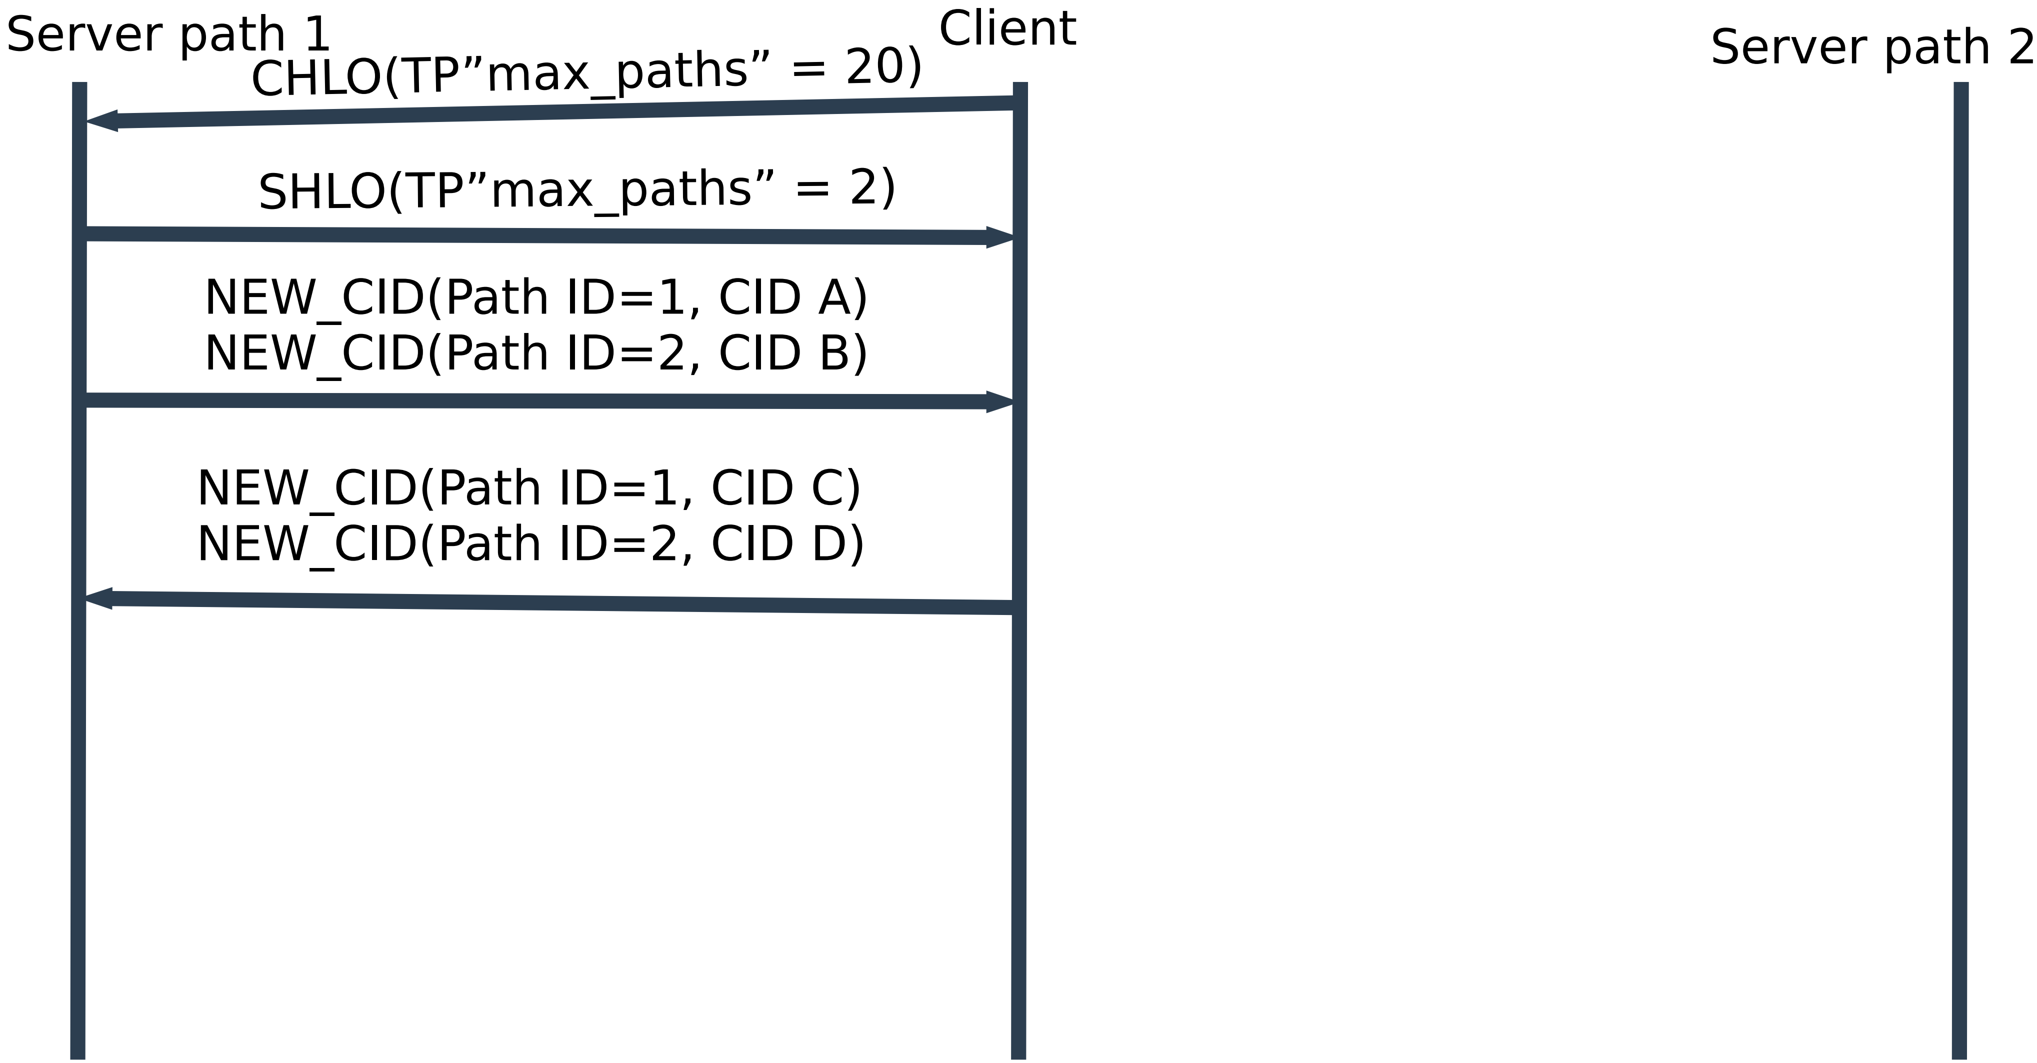

Negotiating Multipath Usage

MPQUIC-ID

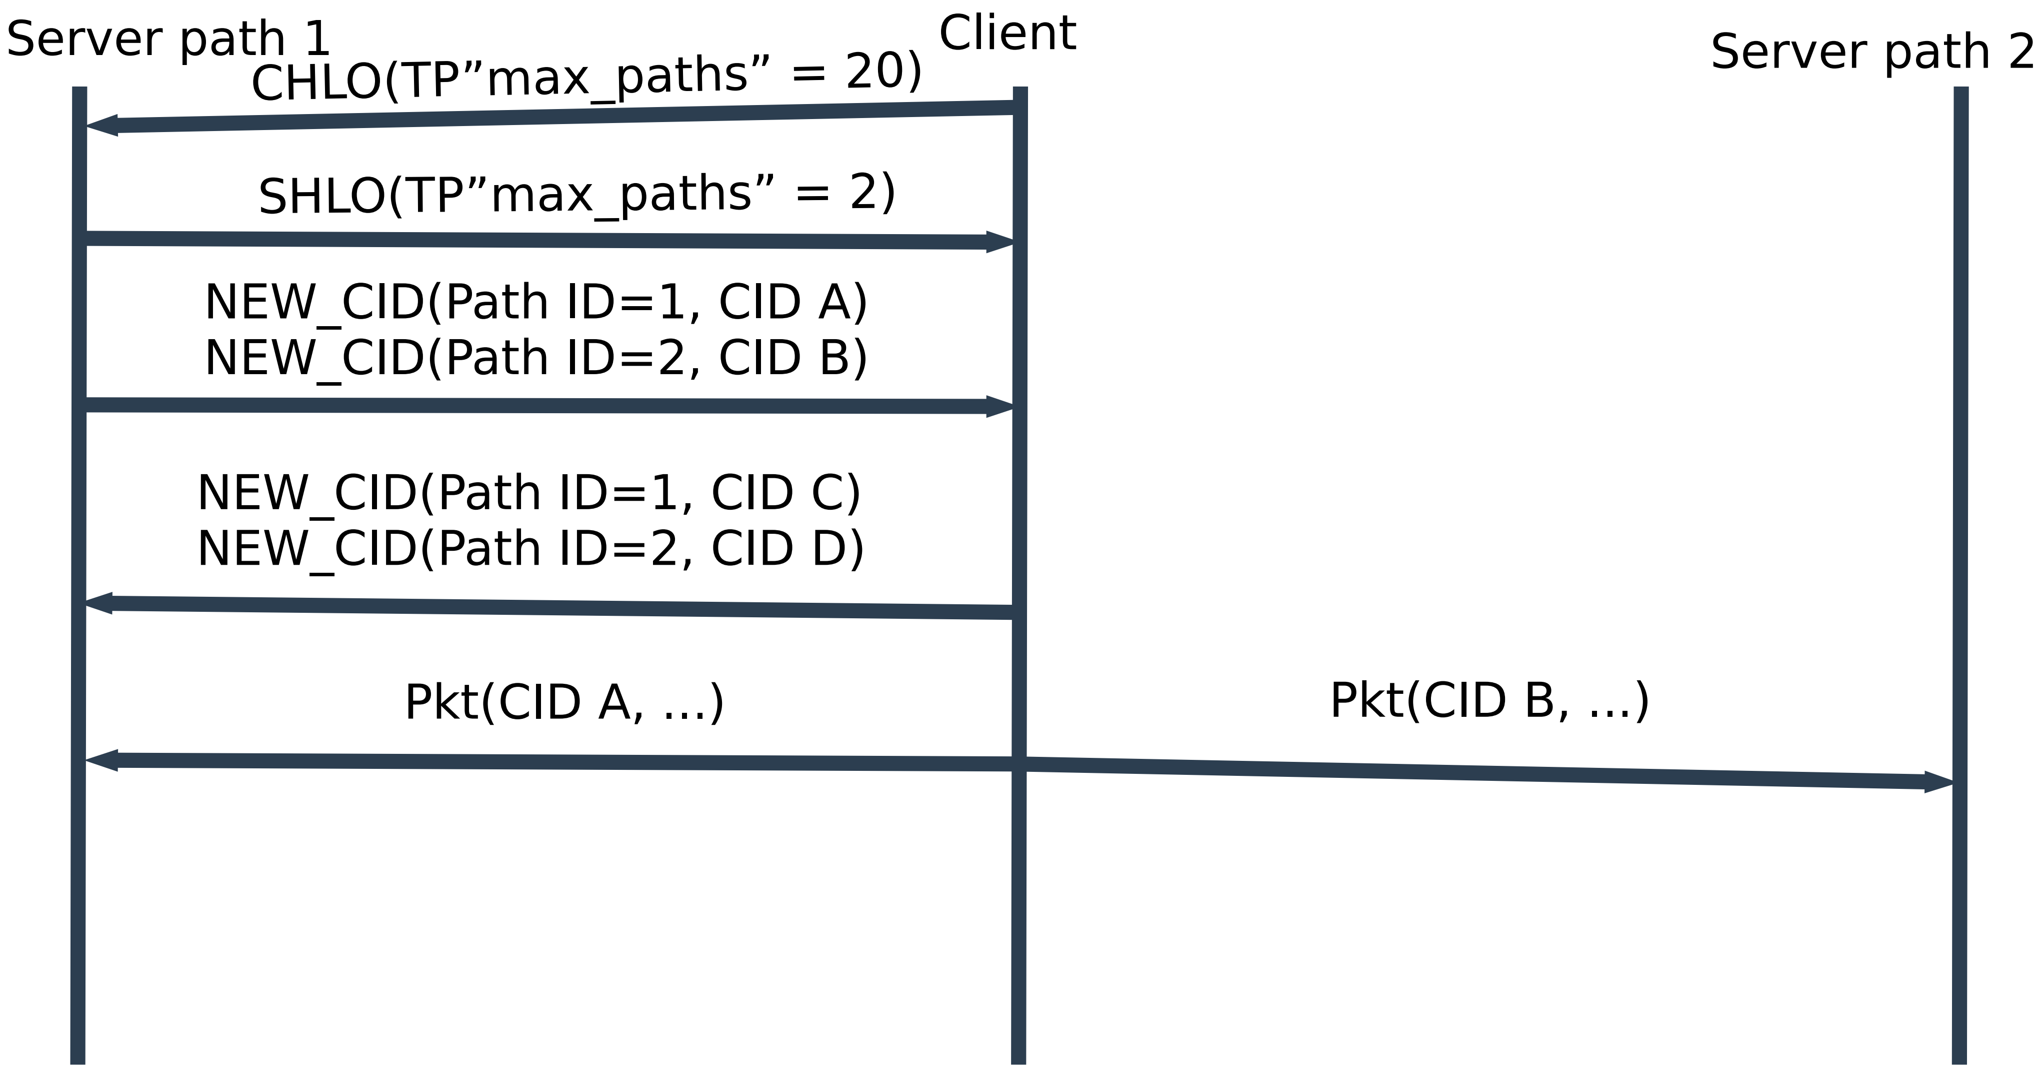

Negotiating Multipath Usage

MPQUIC-ID

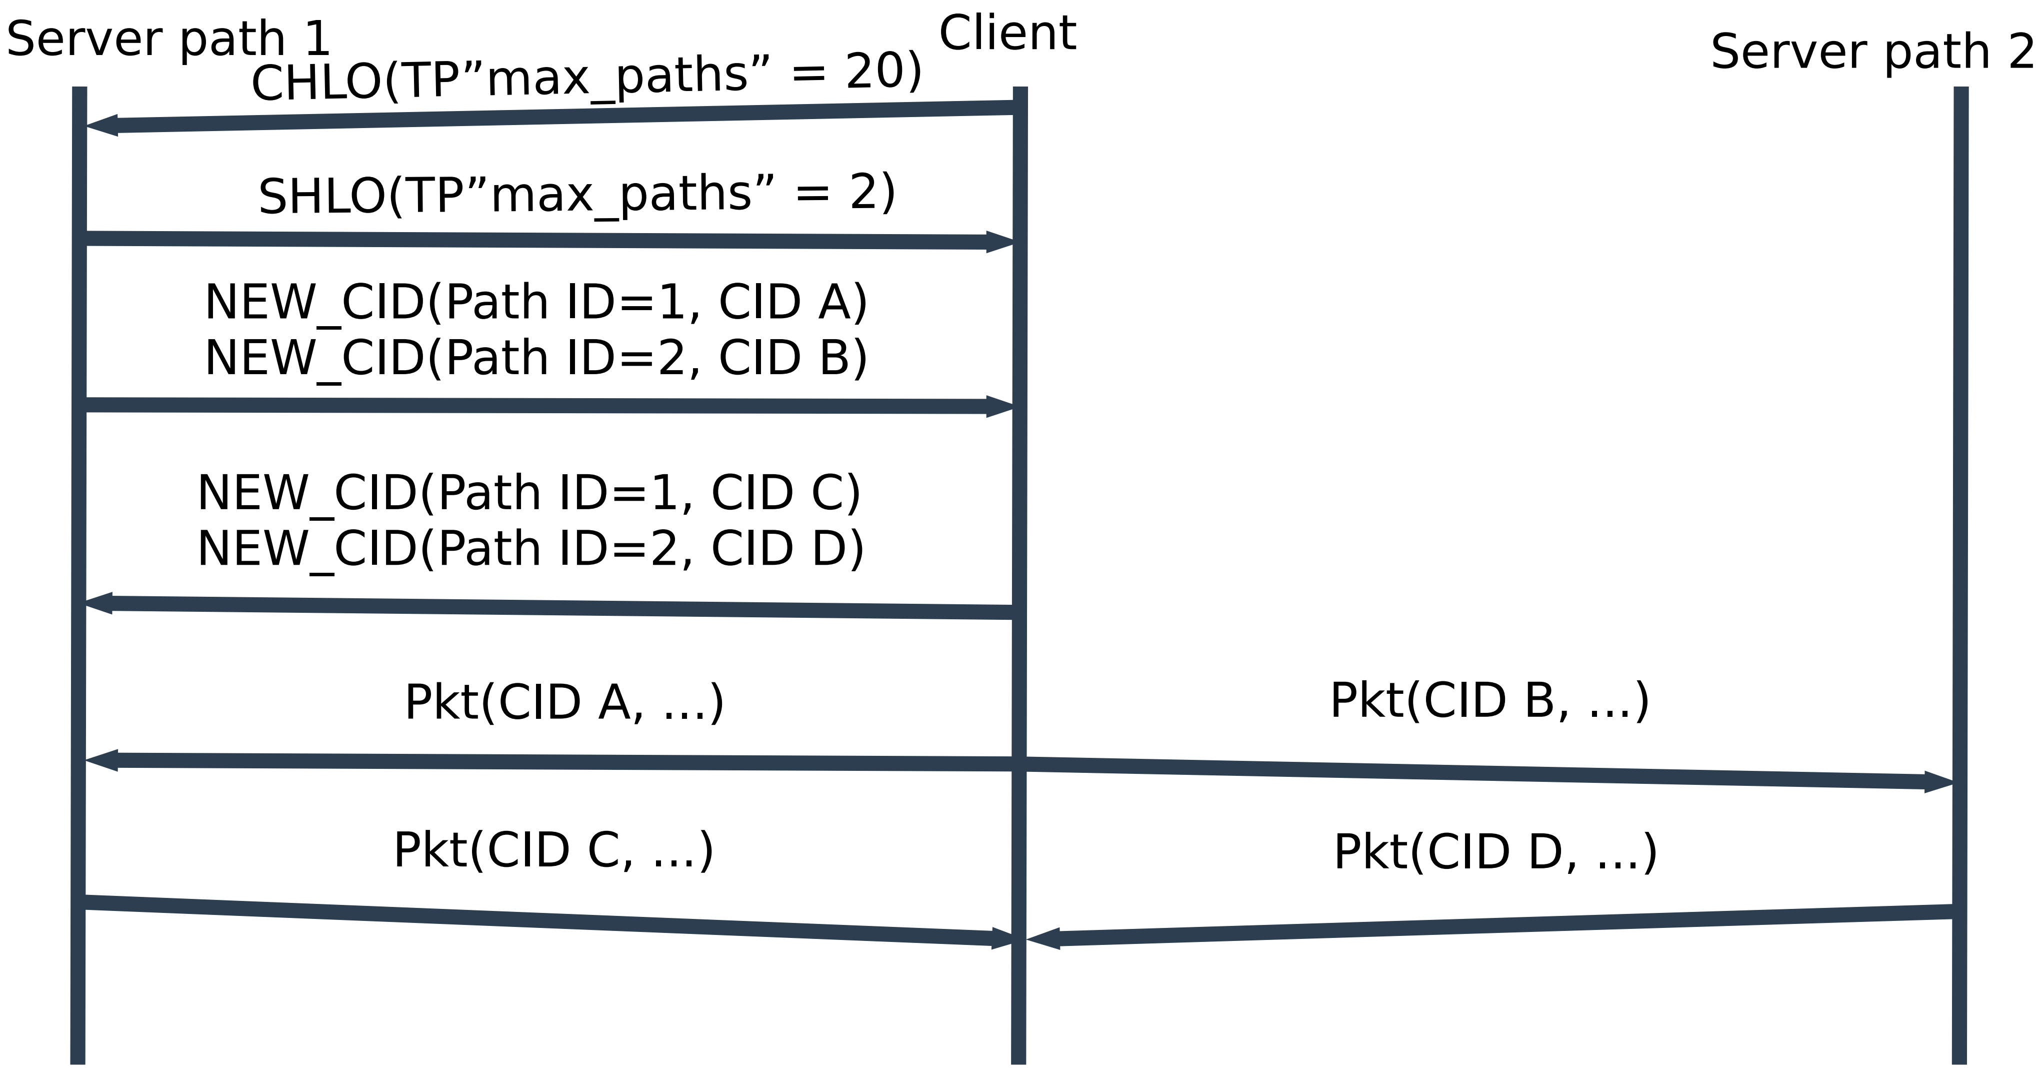

Architecture of Multipath QUIC

MPQUIC-ID

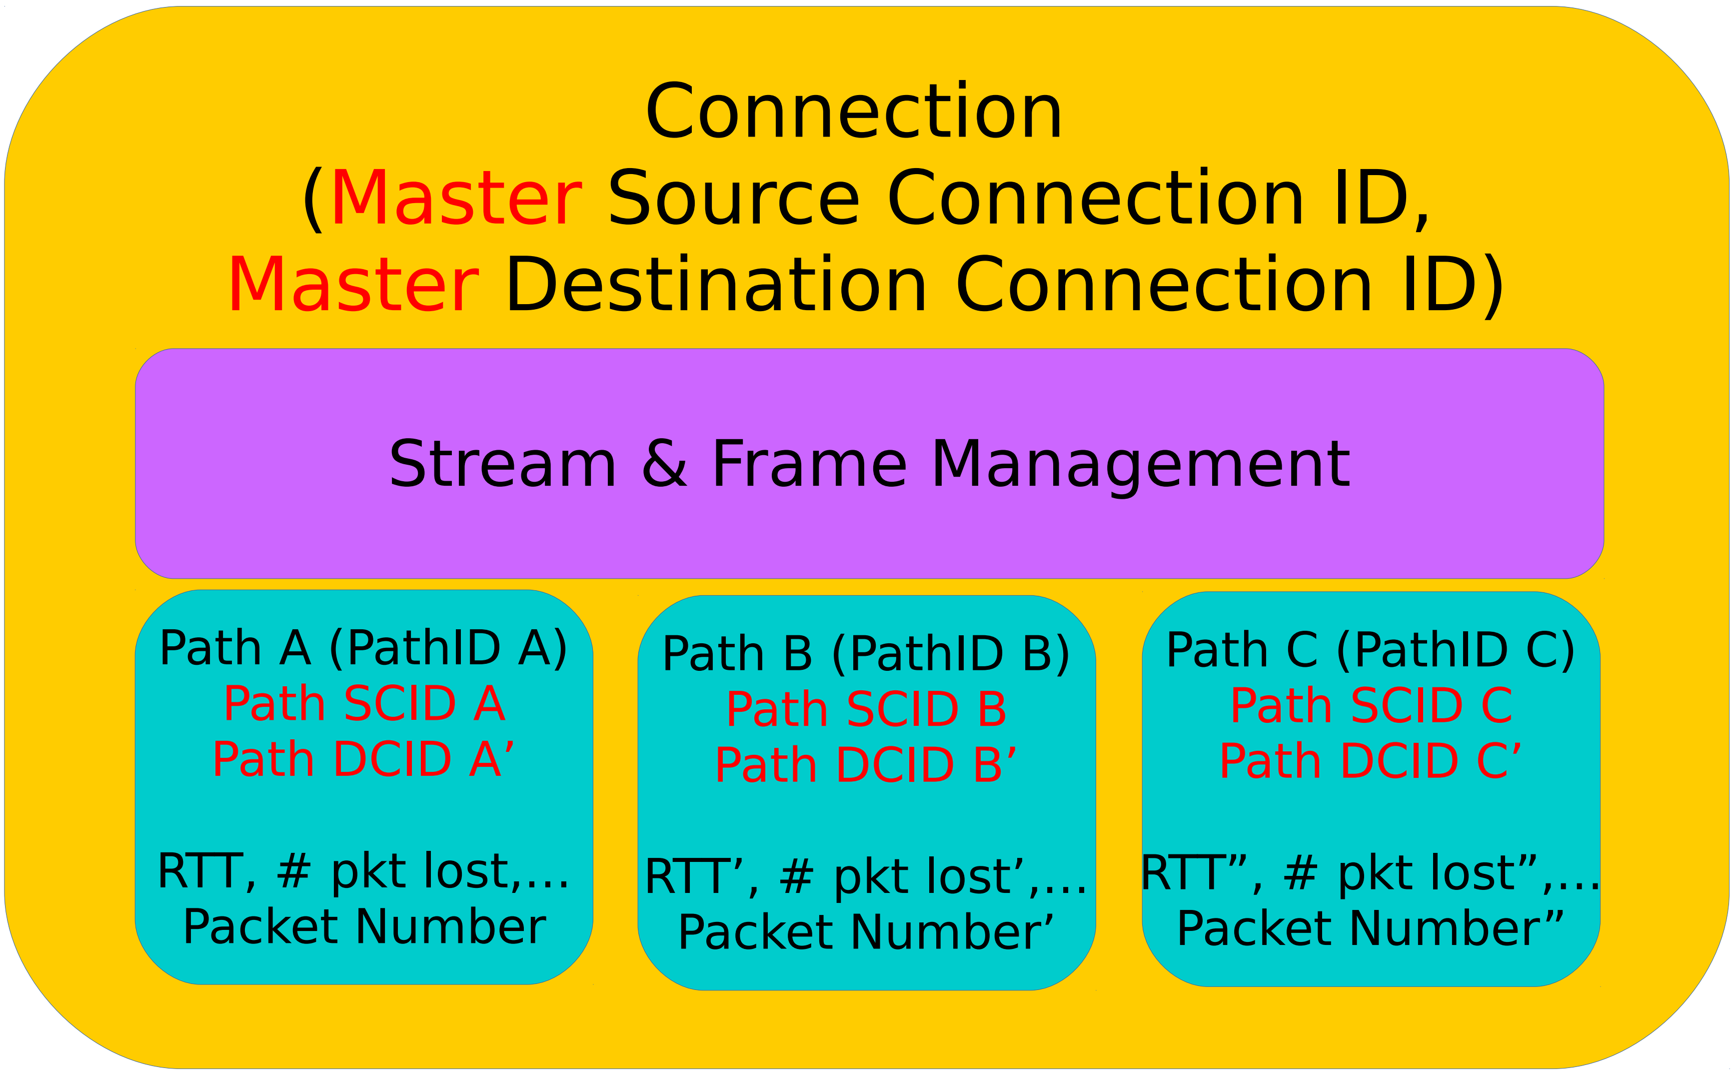

Summary of Changes

- **Transport parameter for MP usage**
- **Wait end of handshake before MP usage**
- **Adding PathID varint in frames**
  - NEW_CONNECTION_ID
  - RETIRE_CONNECTION_ID
  - ACK
- **New frames**
  - ADD_ADDRESS
  - REMOVE_ADDRESS
  - MAX_PATHS
  - PATH_UPDATE

Open Challenges and Opportunities

Multipath Scheduling

- Multiple paths choice, like MPTCP…
- … but with more than 1 stream!

Multipath Scheduling

- **Multiple paths choice, like MPTCP…**
- **… but with more than 1 stream!**

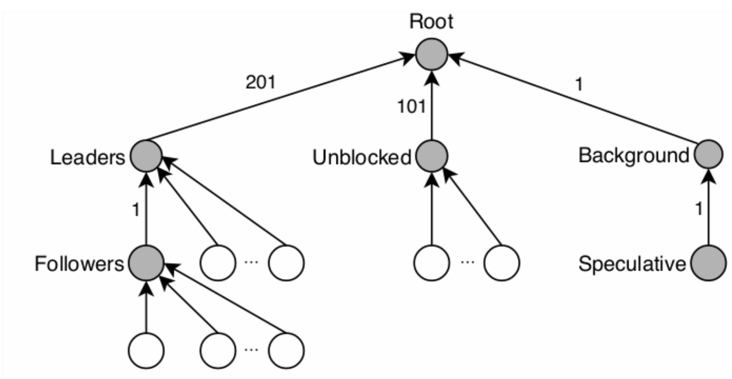

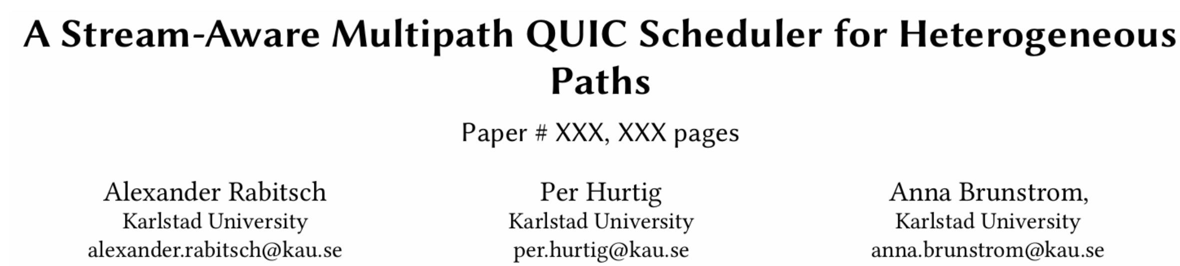

A Stream-Aware Multipath QUIC Scheduler for Heterogeneous Paths

Paper # XXX, XXX pages

Alexander Rabitsch
Karlstad University
alexander.rabitsch@kau.se

Per Hurtig
Karlstad University
per.hurtig@kau.se

Anna Brunstrom,
Karlstad University
anna.brunstrom@kau.se

Congestion Control Scheme

- **How to remain fair but efficient?**
  - And with multipath?

Congestion Control Scheme

- **How to remain fair but efficient?**
  - And with multipath?

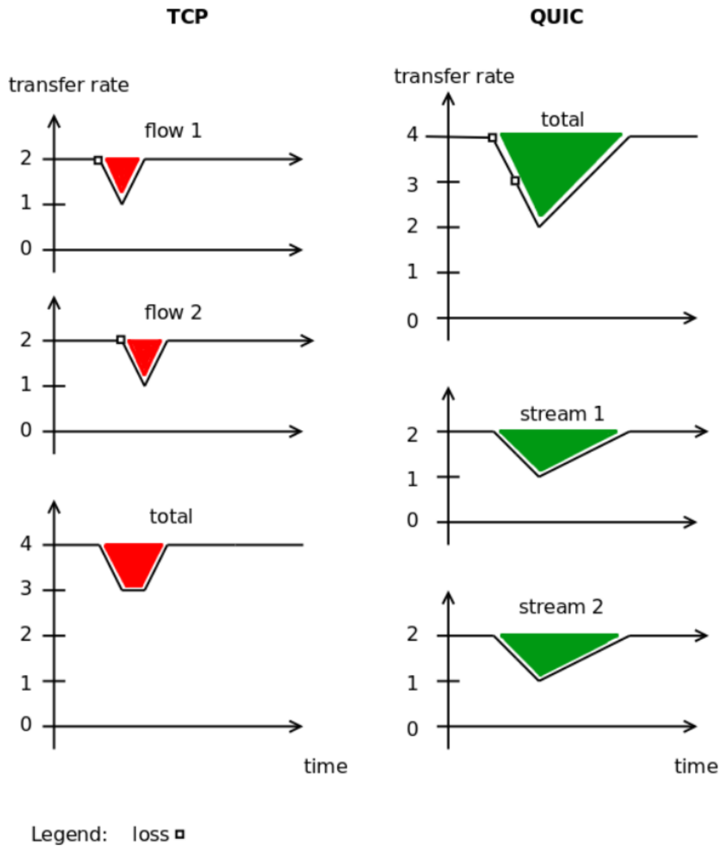

Congestion Control Scheme

- **How to remain fair but efficient?**
  - And with multipath?

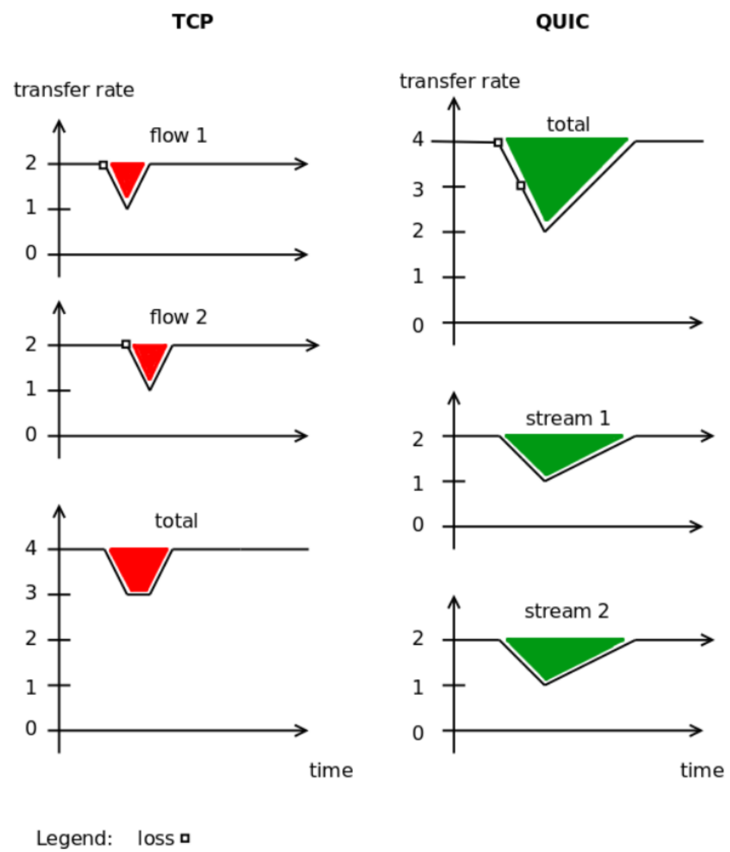

Taking a Long Look at QUIC

An Approach for Rigorous Evaluation of Rapidly Evolving Transport Protocols

Arash Molavi Kakhki
Northeastern University
arash@ccs.neu.edu

Samuel Jero
Purdue University
sjero@purdue.edu

David Choffnes
Northeastern University
choffnes@ccs.neu.edu

Cristina Nita-Rotaru
Northeastern University
c.nitarotaru@neu.edu

Alan Mislove
Northeastern University
amislove@ccs.neu.edu

Scenario	Flow	Avg. throughput (std. dev.)
QUIC vs. TCP	QUIC	2.71 (0.46)
	TCP	1.62 (1.27)
QUIC vs. TCPx2	QUIC	2.8 (1.16)
	TCP 1	0.7 (0.21)
	TCP 2	0.96 (0.3)
QUIC vs. TCPx4	QUIC	2.75 (1.2)
	TCP 1	0.45 (0.14)
	TCP 2	0.36 (0.09)
	TCP 3	0.41 (0.11)
	TCP 4	0.45 (0.13)

Table 4: Average throughput (5 Mbps link, buffer=30 KB, averaged over 10 runs) allocated to QUIC and TCP flows when competing with each other. Despite the fact that both protocols use Cubic congestion control, QUIC consumes nearly twice the bottleneck bandwidth than TCP flows combined, resulting in substantial unfairness.

Handover situations

- **How does Multipath QUIC help under mobility in wireless environment?**
  - Especially with ≠ path priorities

Handover situations

- **How does Multipath QUIC help under mobility in wireless environment?**
  - Especially with ≠ path priorities

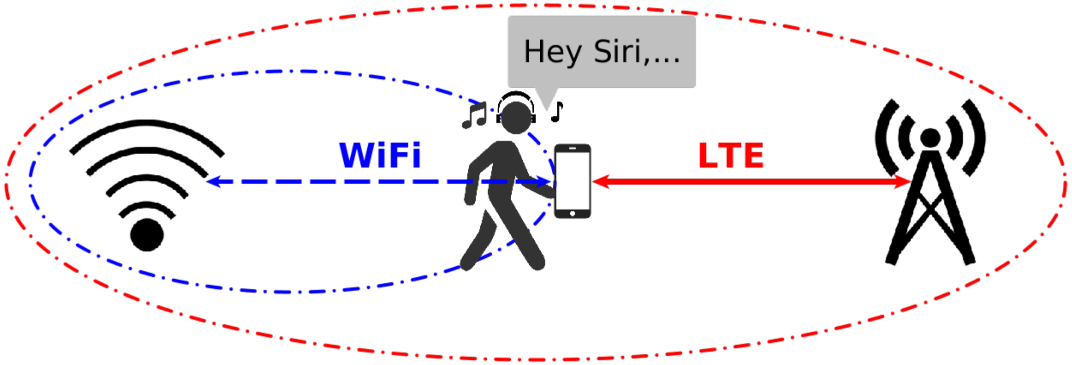

MultipathTester

https://itunes.apple.com/us/app/multipathtester/id1351286809?mt=8

Asymmetric Paths

- So far, assume that paths are symmetric...
- ... what if they are not?

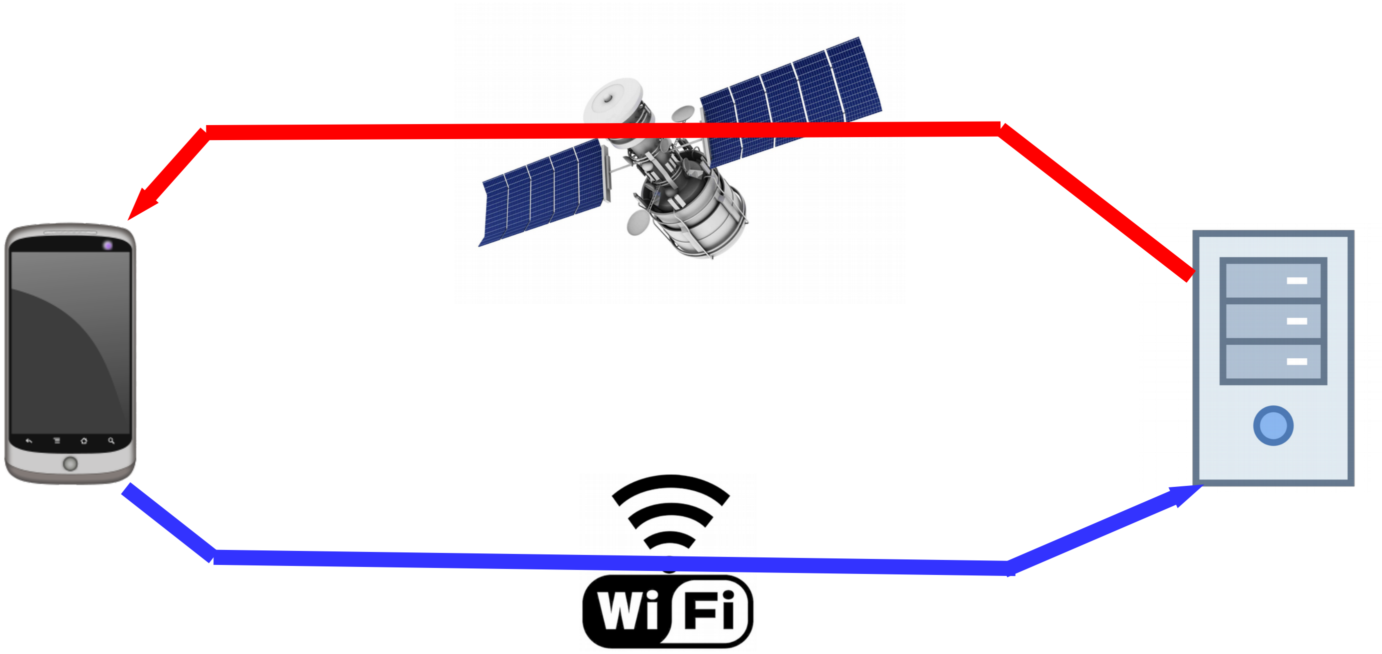

And All The Others…

- **Don’t hesitate to discuss your own challenges :-)**

Thanks for your attention!

References

- **[CoNEXT'17] Quentin De Coninck and Olivier Bonaventure.** ***Multipath QUIC: Design and Evaluation.***
- **[ICC'18] Tobias Viernickel, Alexander Froemmgen, Amr Rizk, Boris Koldehofe and Ralf Steinmetz.** ***Multipath QUIC: A Deployable Multipath Transport Protocol.***
- **[MPQUIC-ID] Quentin De Coninck and Olivier Bonaventure.** ***Multipath Extensions for QUIC (MP-QUIC)*****. Draft-deconinck-quic-multipath-02.**
- **[EPIQ-18] Alexander Rabitsch, Per Hurtig and Anna Brunstrom.** ***A Stream-Aware Multipath QUIC Scheduler for Heterogeneous Paths*****.**
- **[IMC'17] Arash Molavi Kakhki, Samuel Jero, David Choffnes, Cristina Nita-Rotaru and Alan Mislove.** ***Taking a Long Look at QUIC*****.**

Backup slides

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

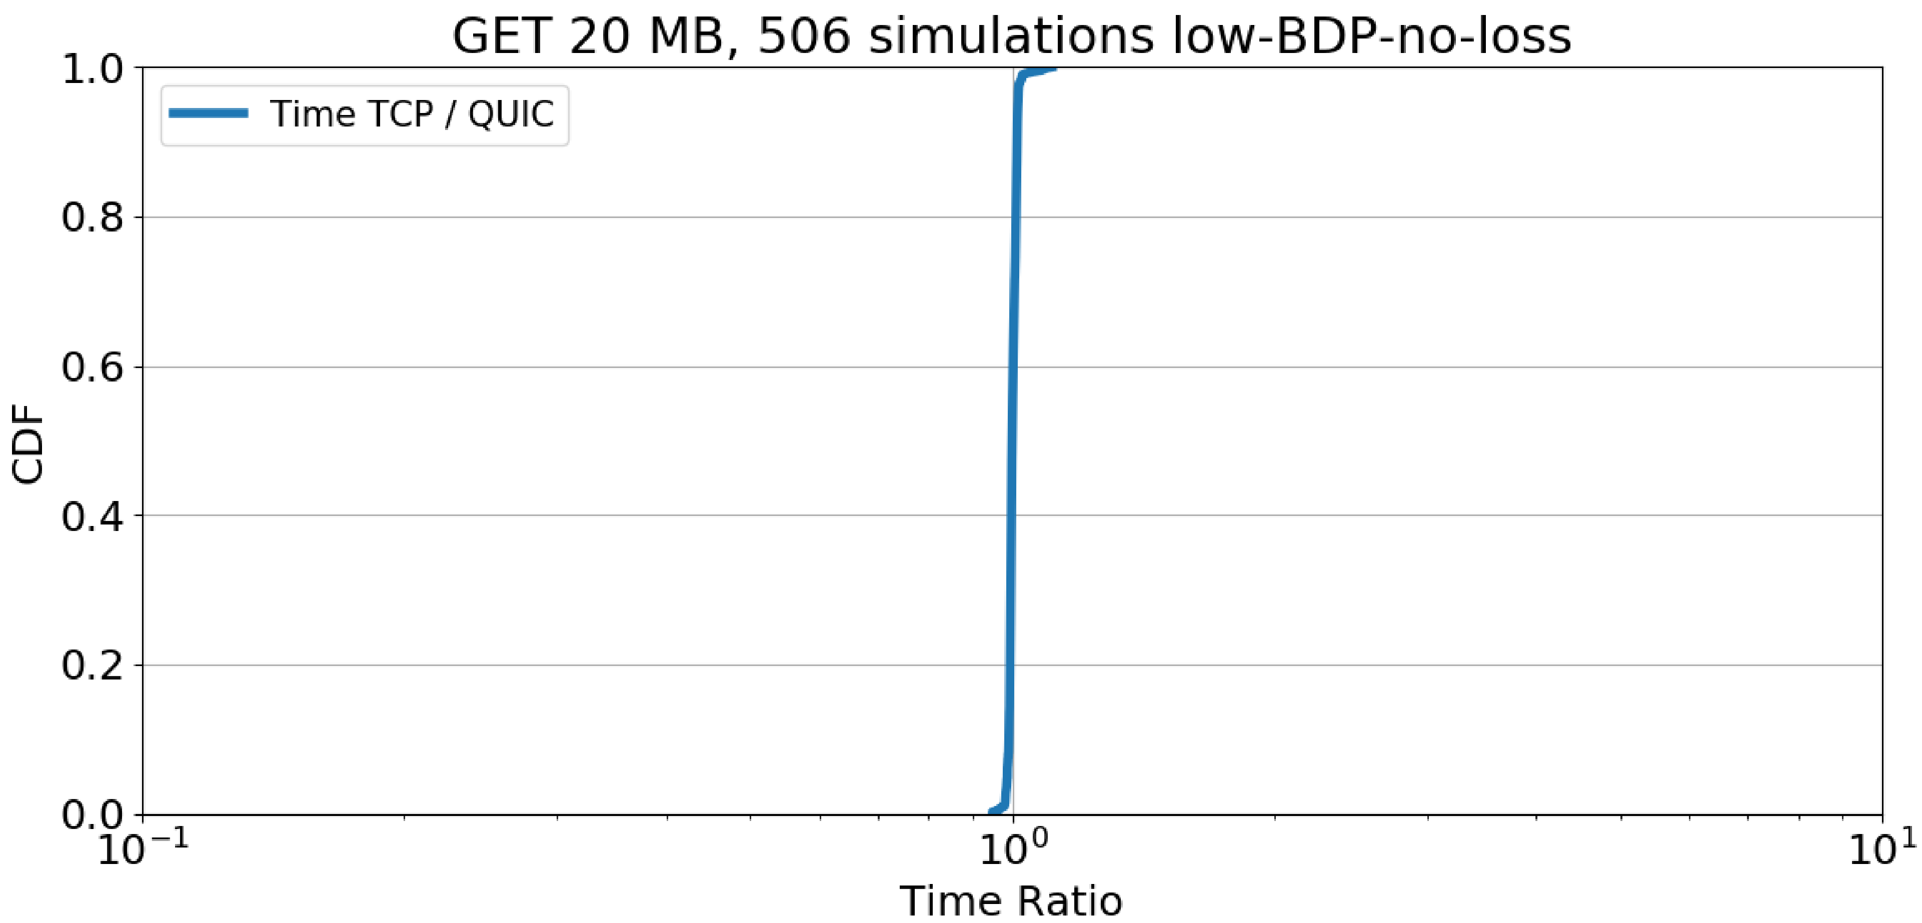

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

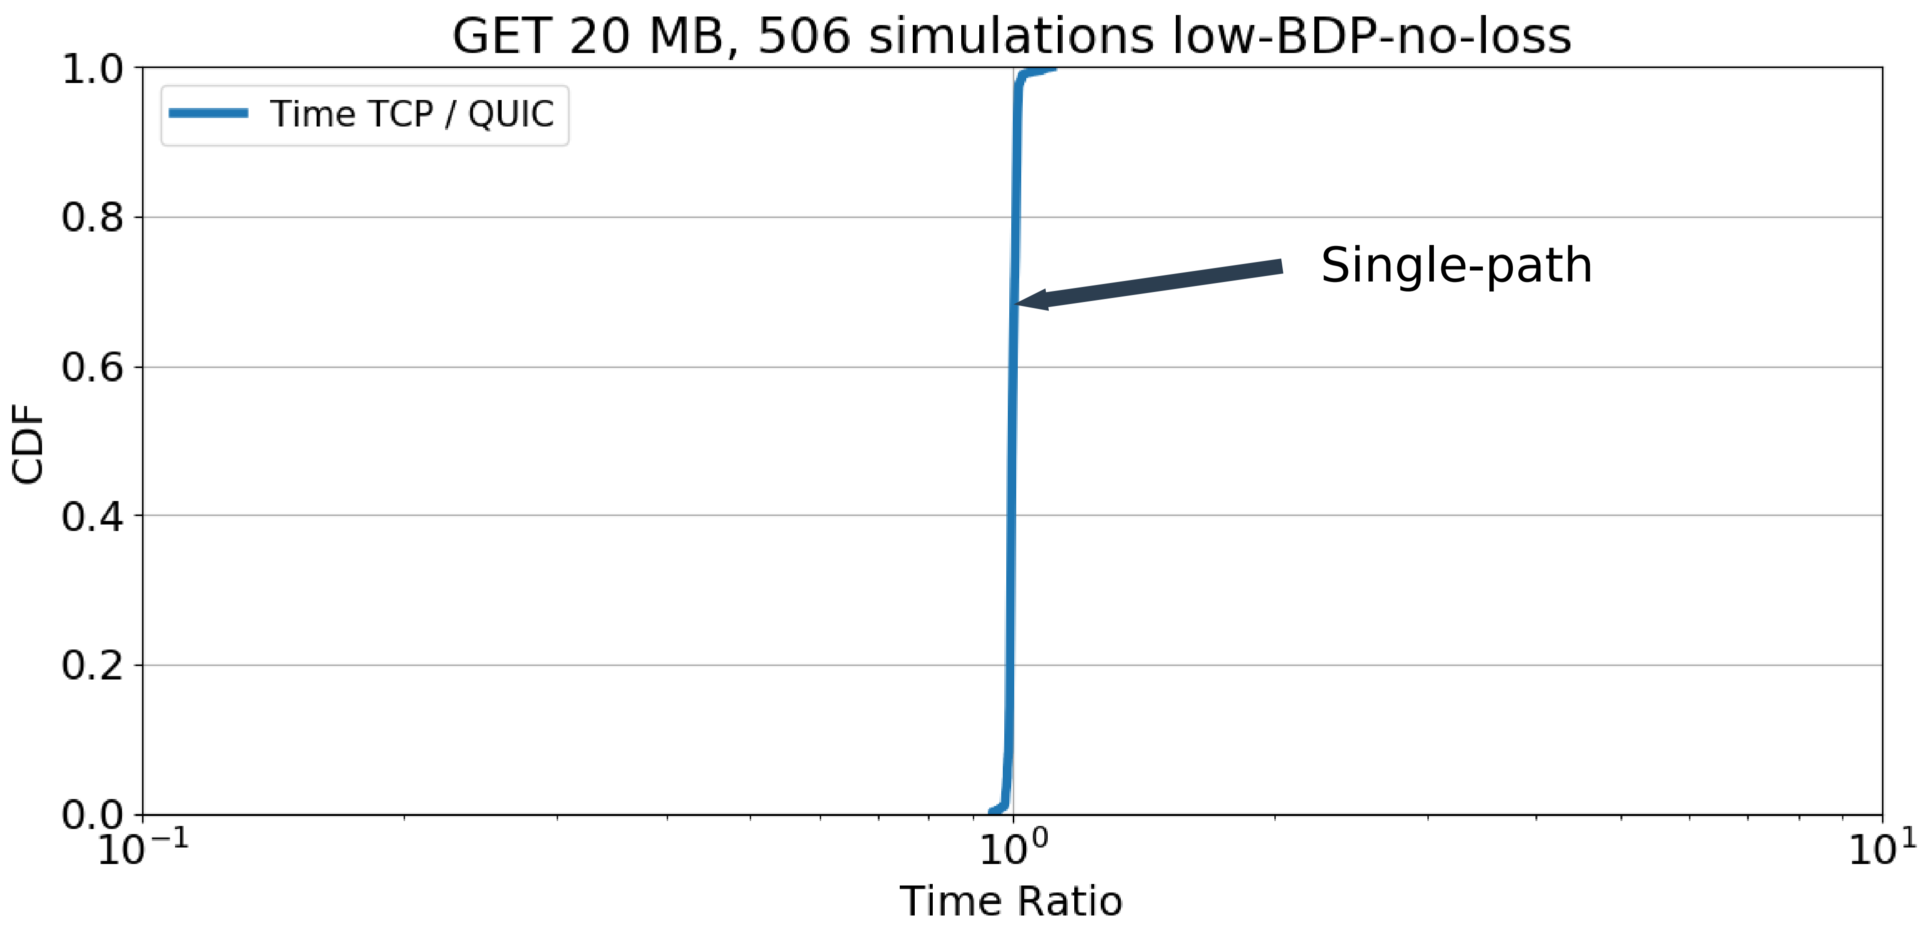

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

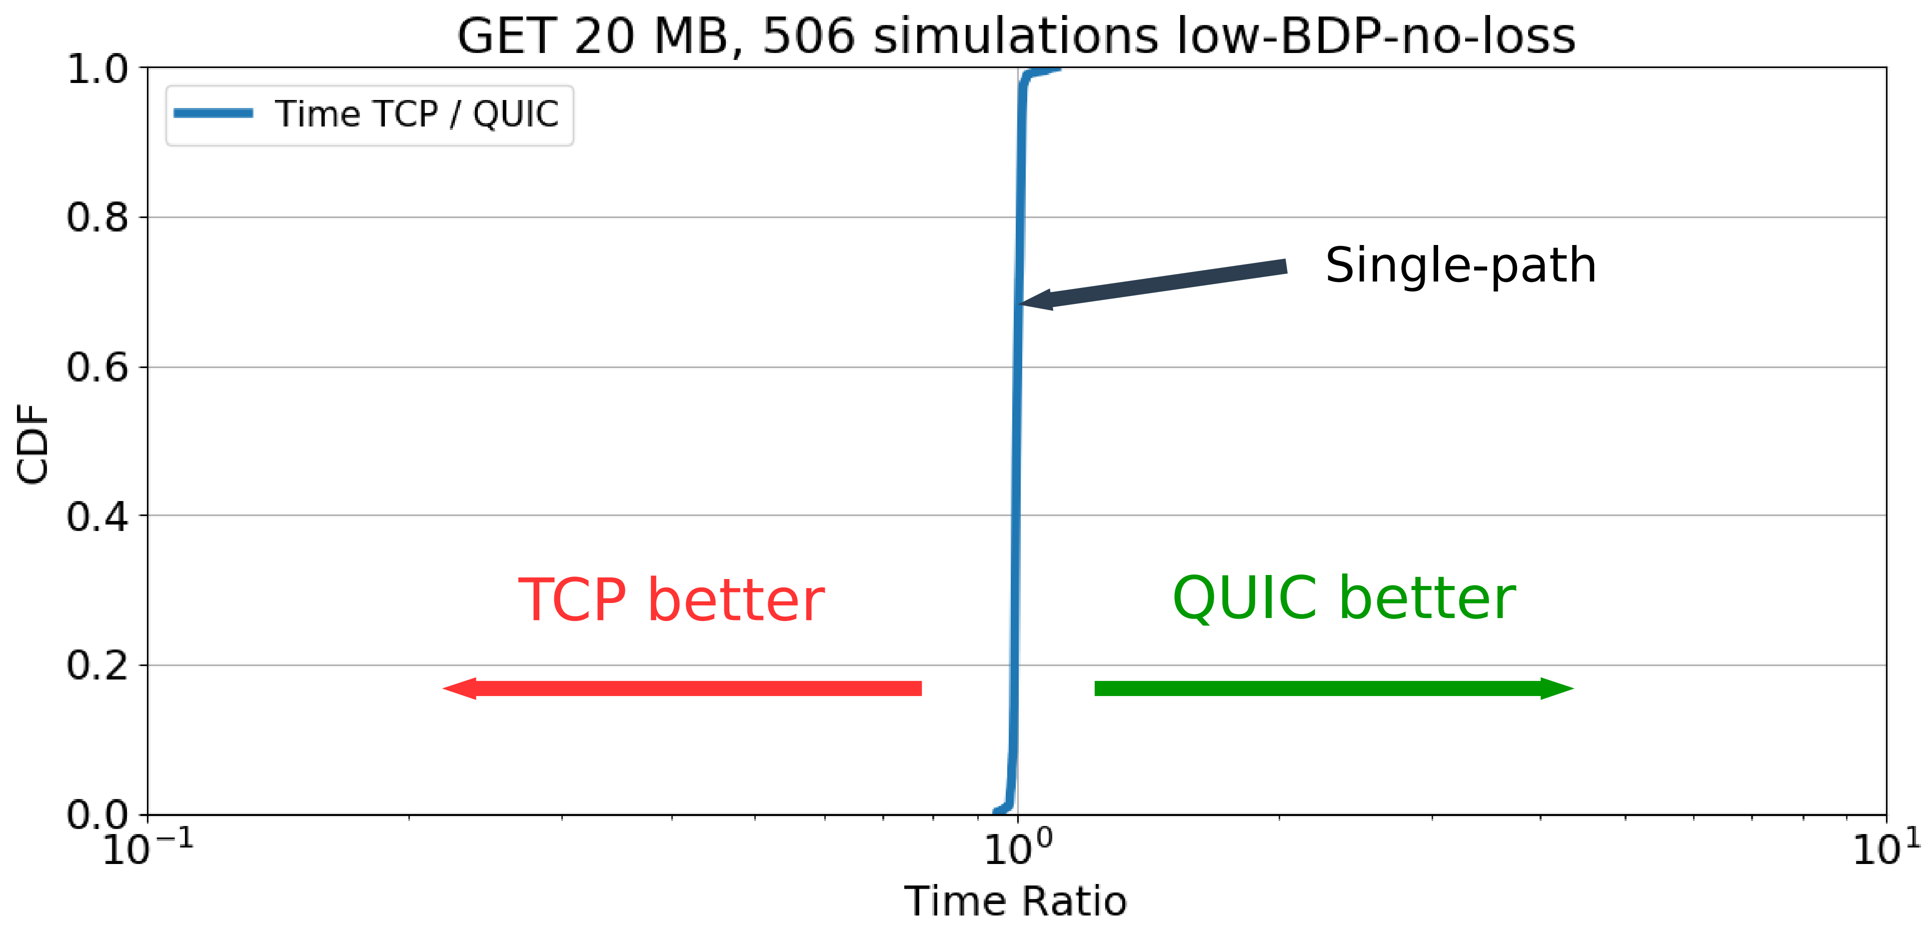

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

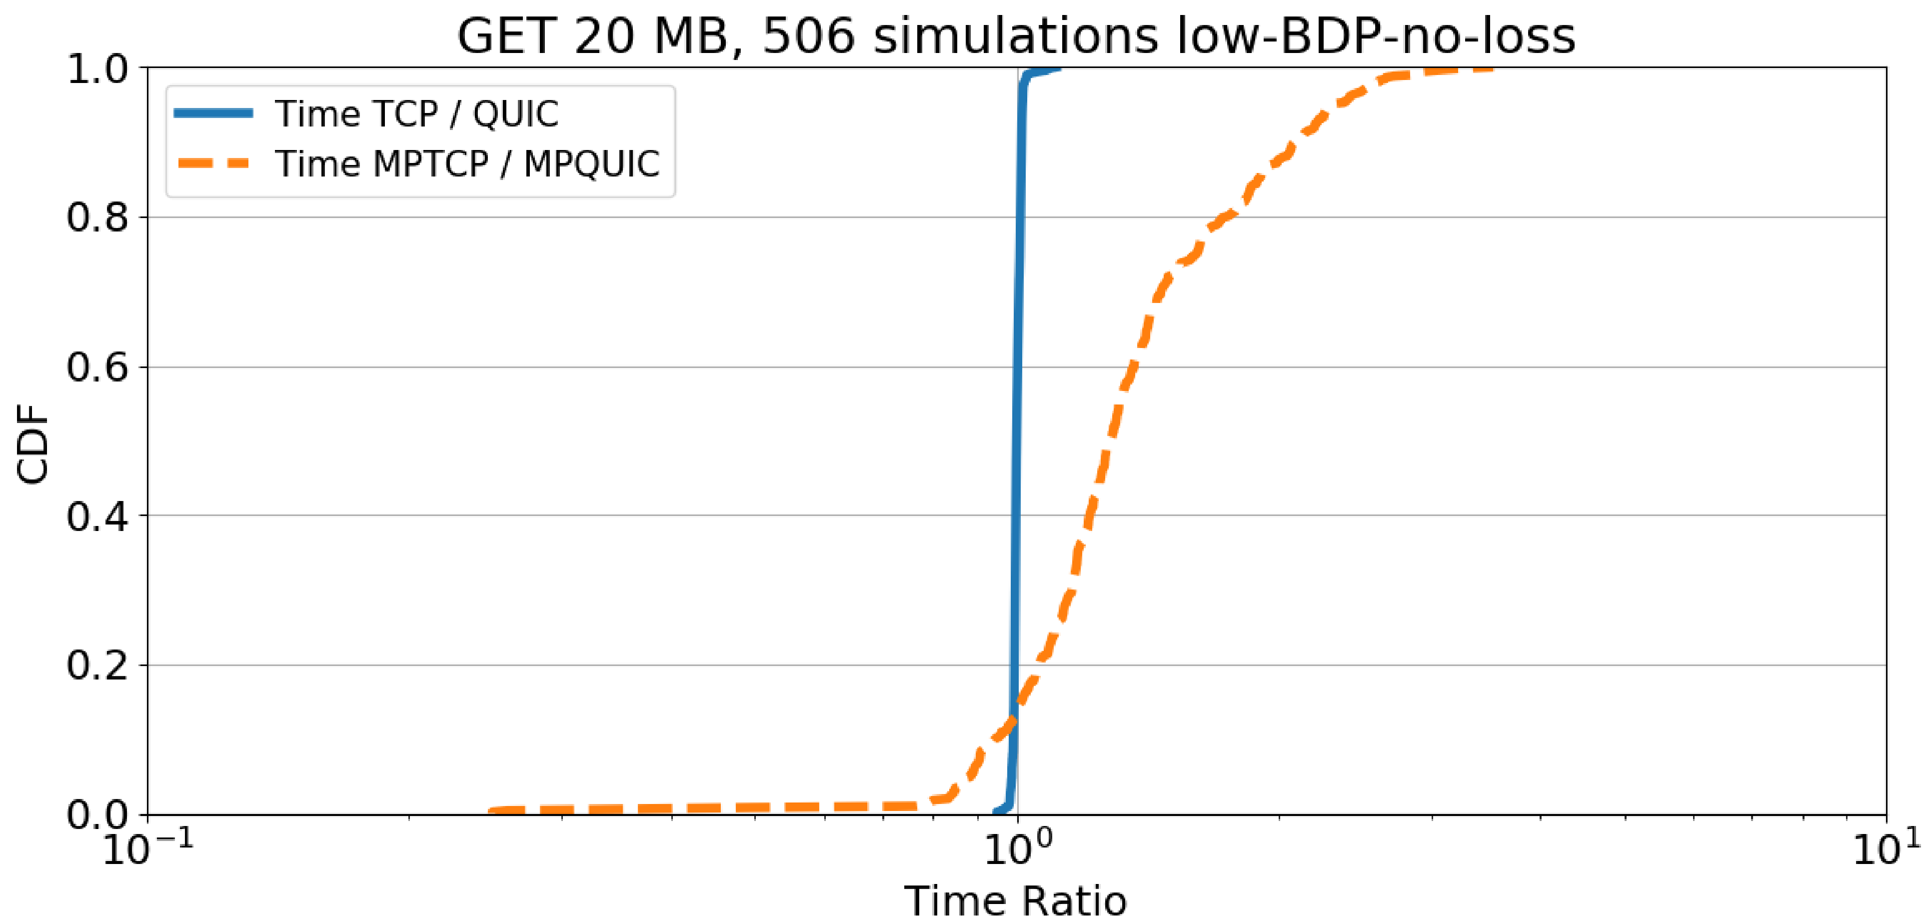

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

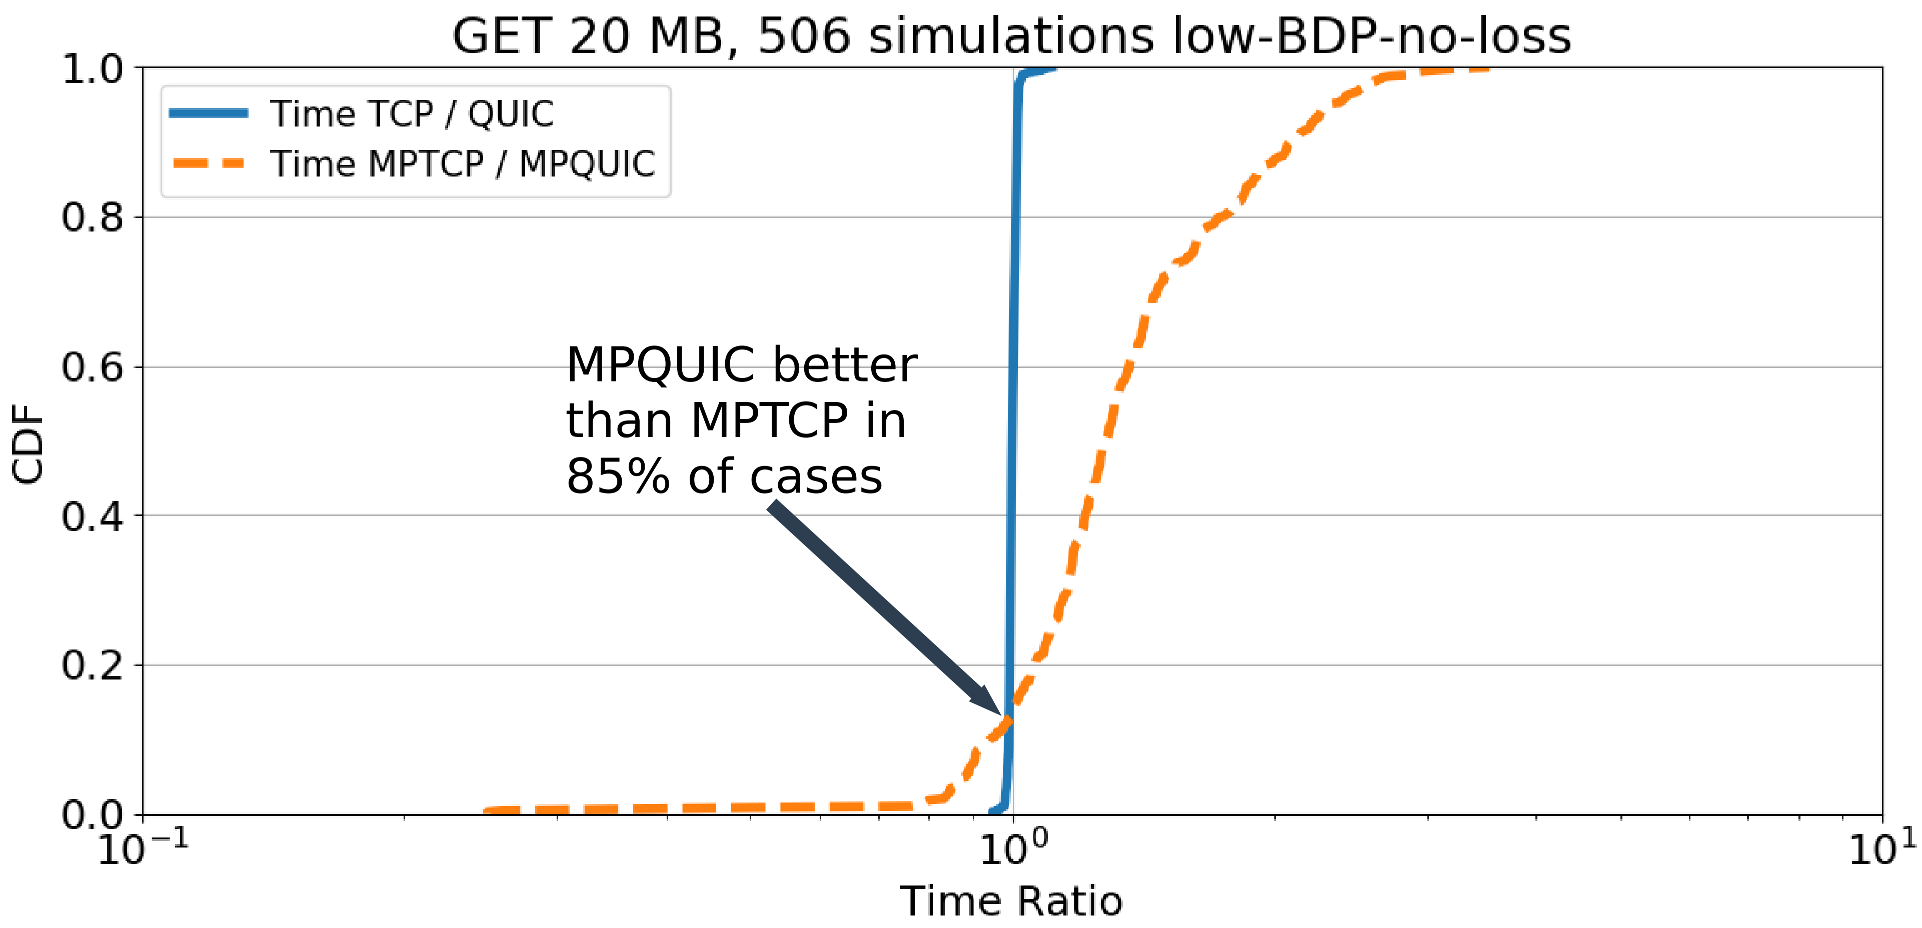

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

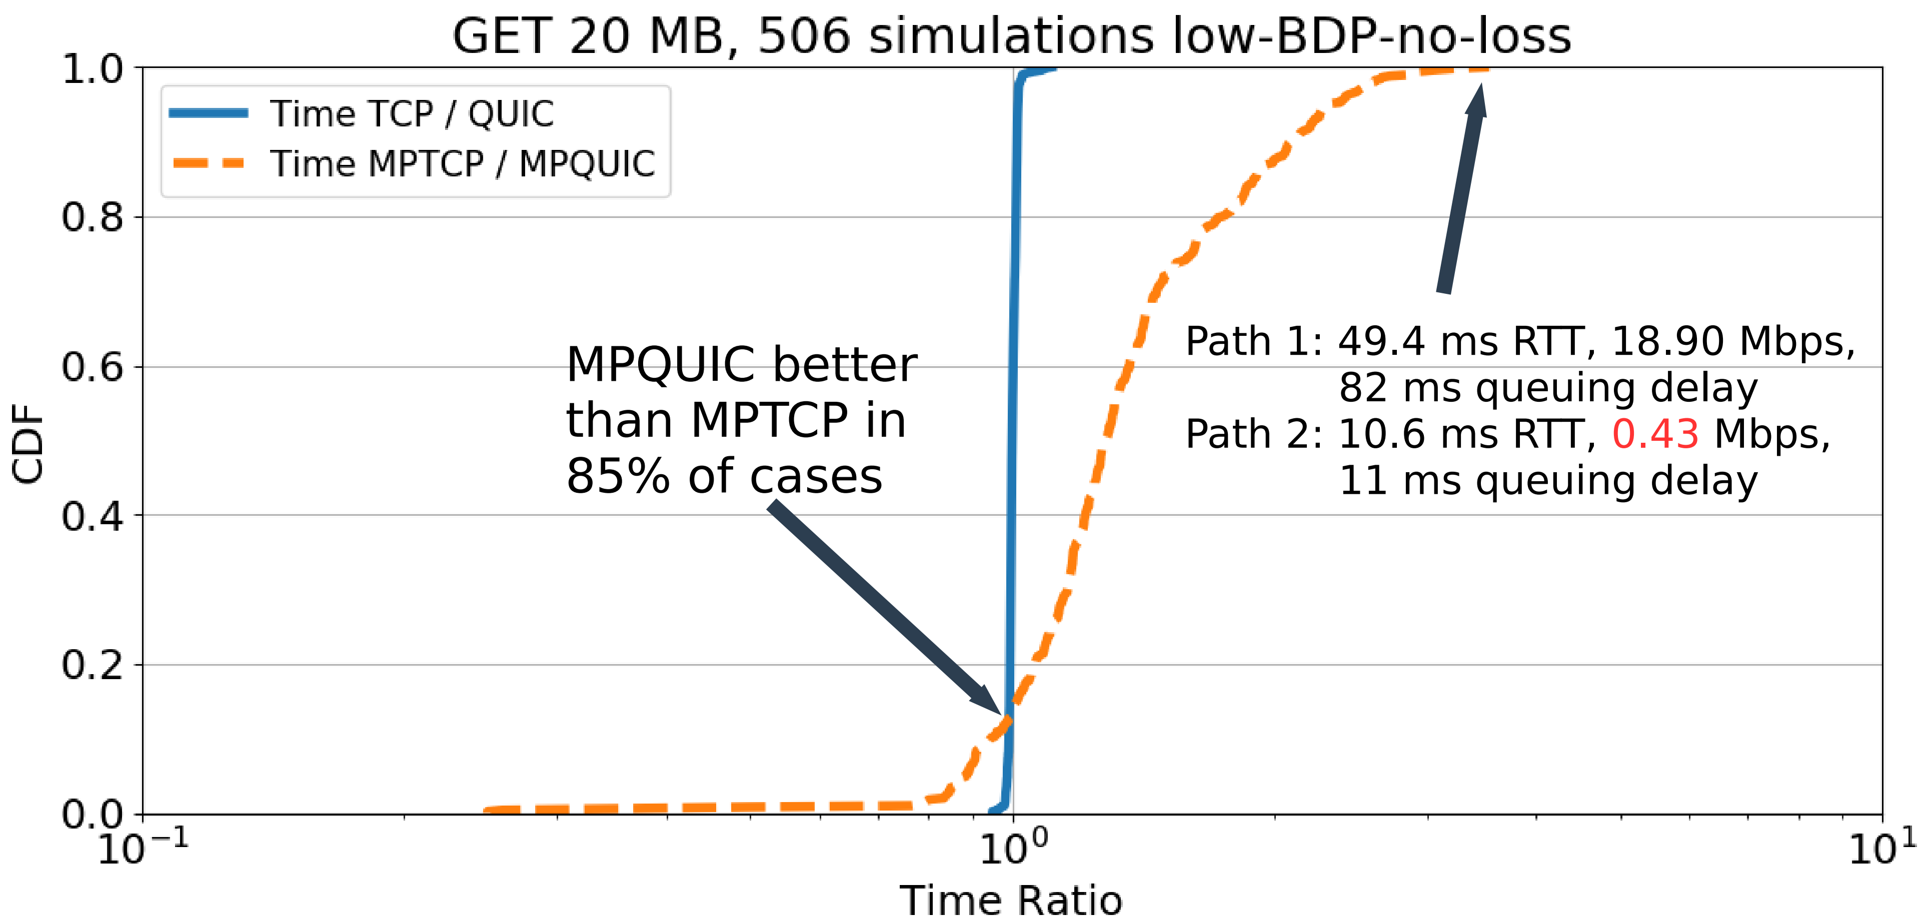

(MP)TCP vs. (MP)QUIC – No Loss

CoNEXT'17

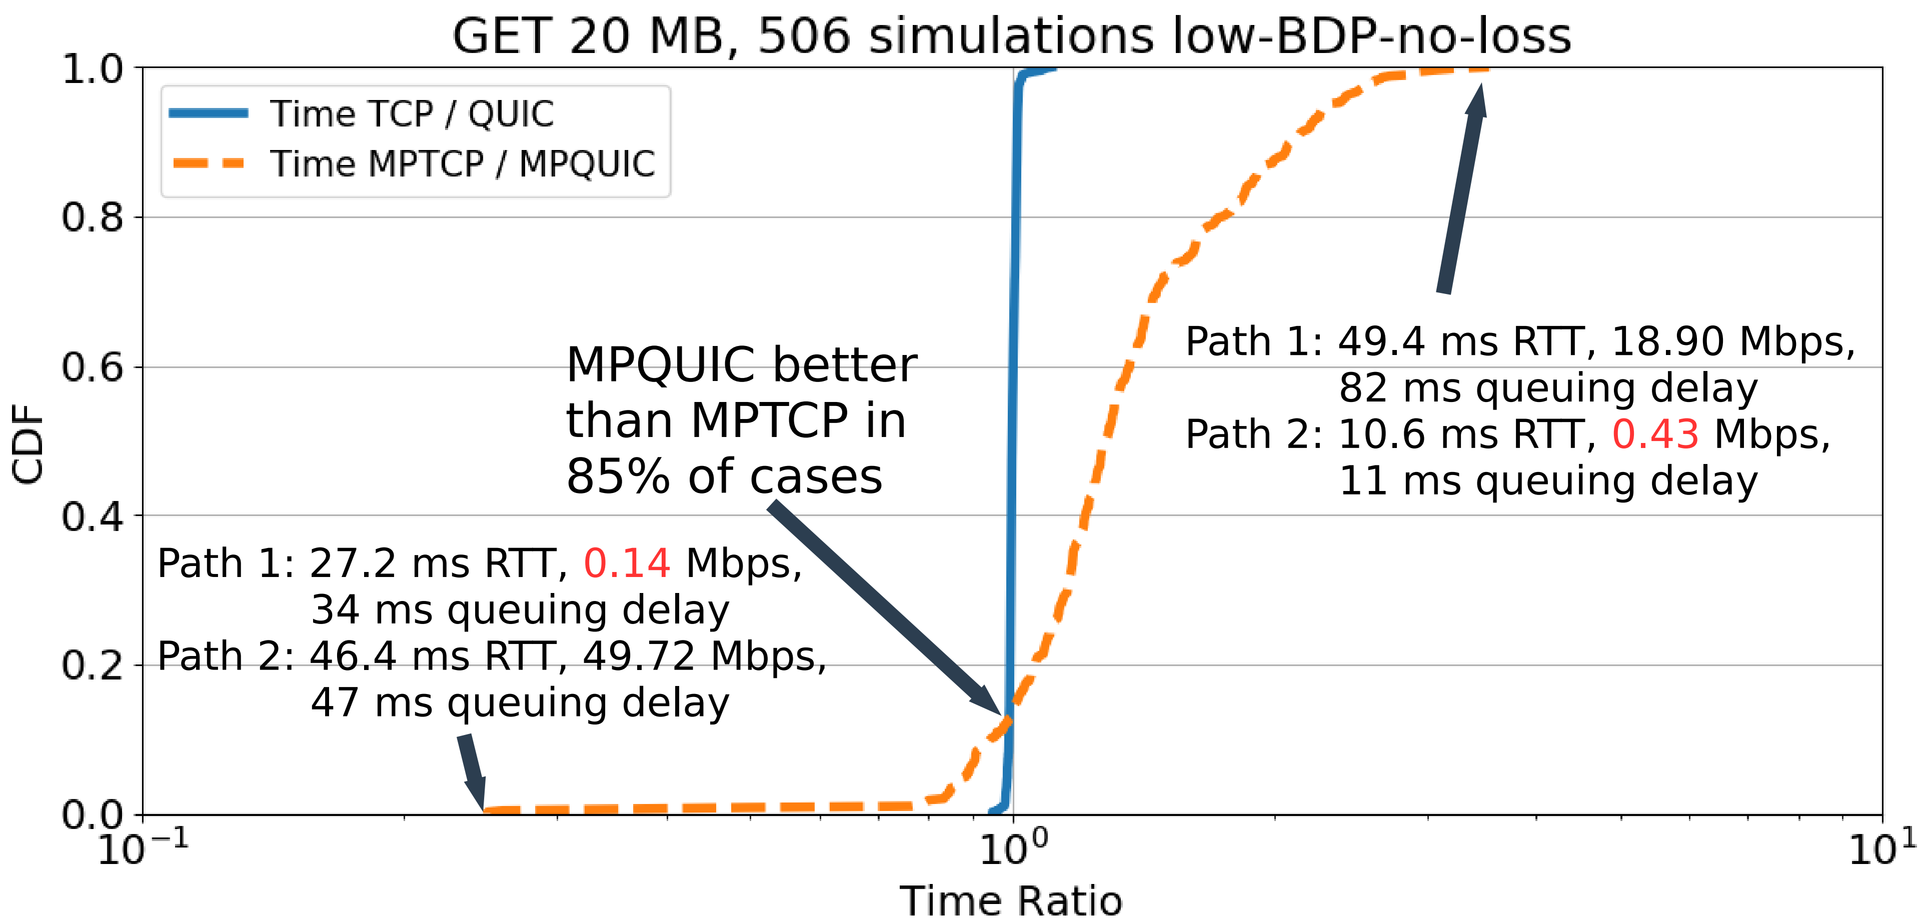

(MP)TCP vs. (MP)QUIC – Losses

CoNEXT'17

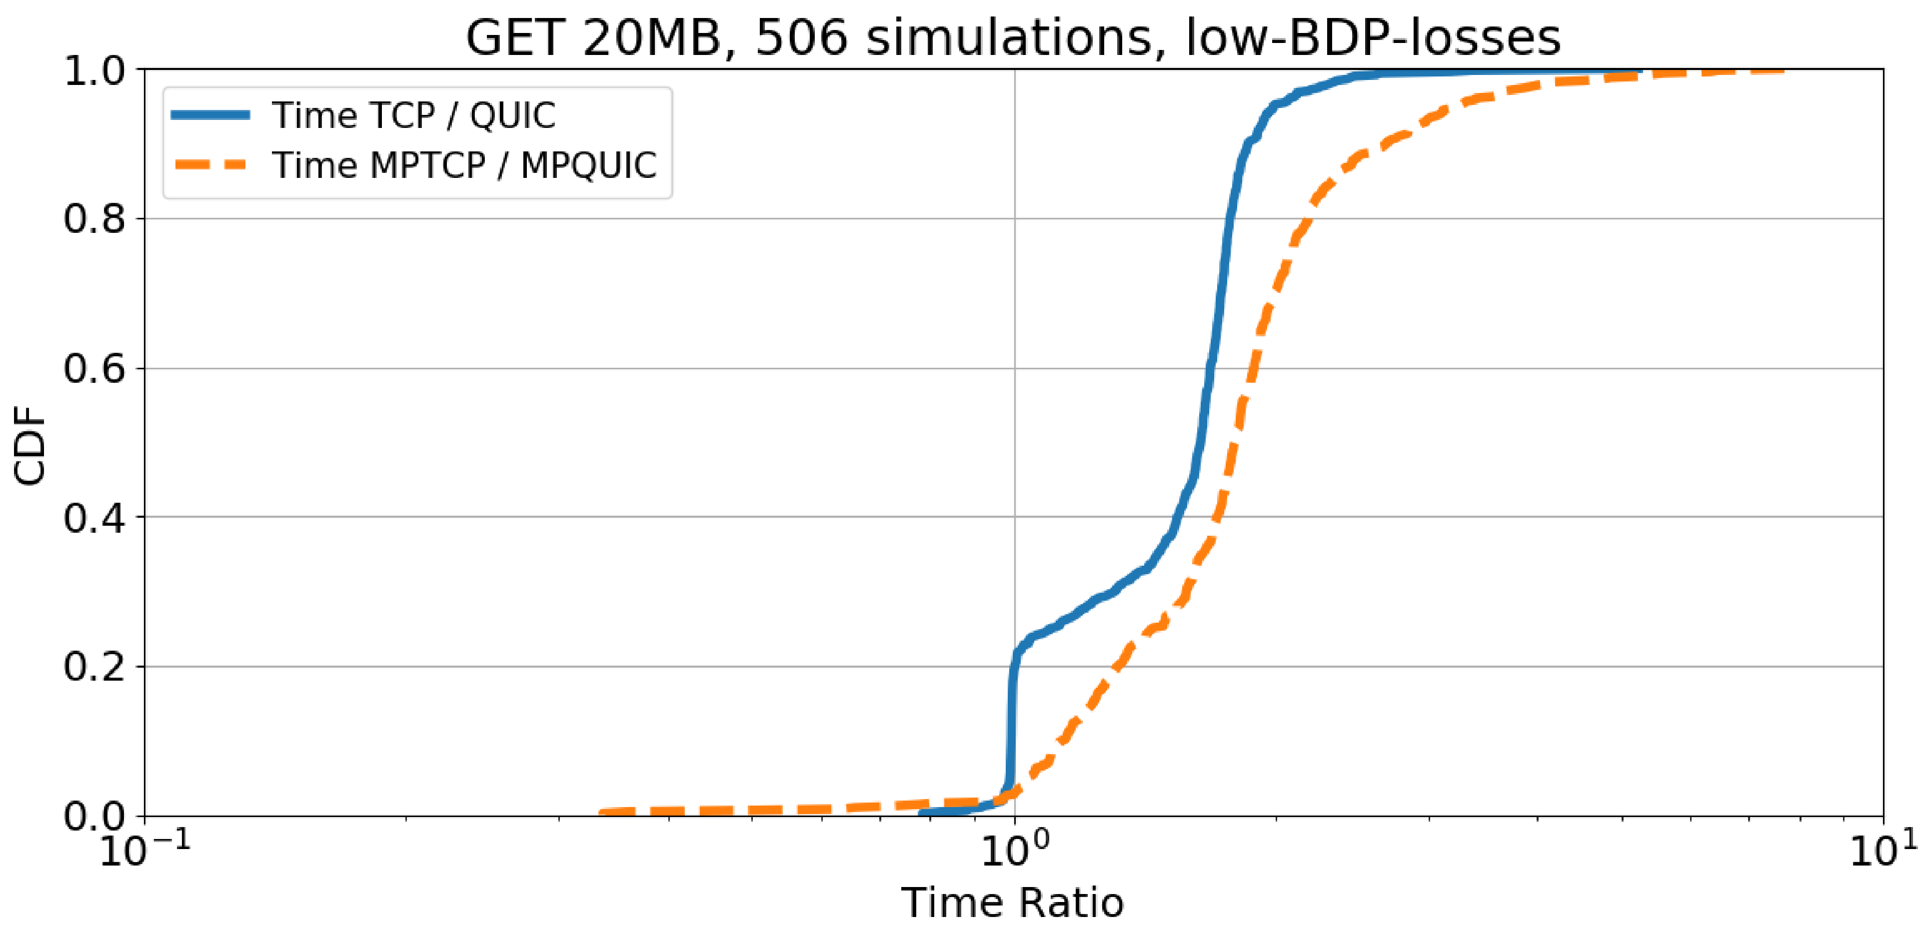

(MP)TCP vs. (MP)QUIC – Losses

CoNEXT'17

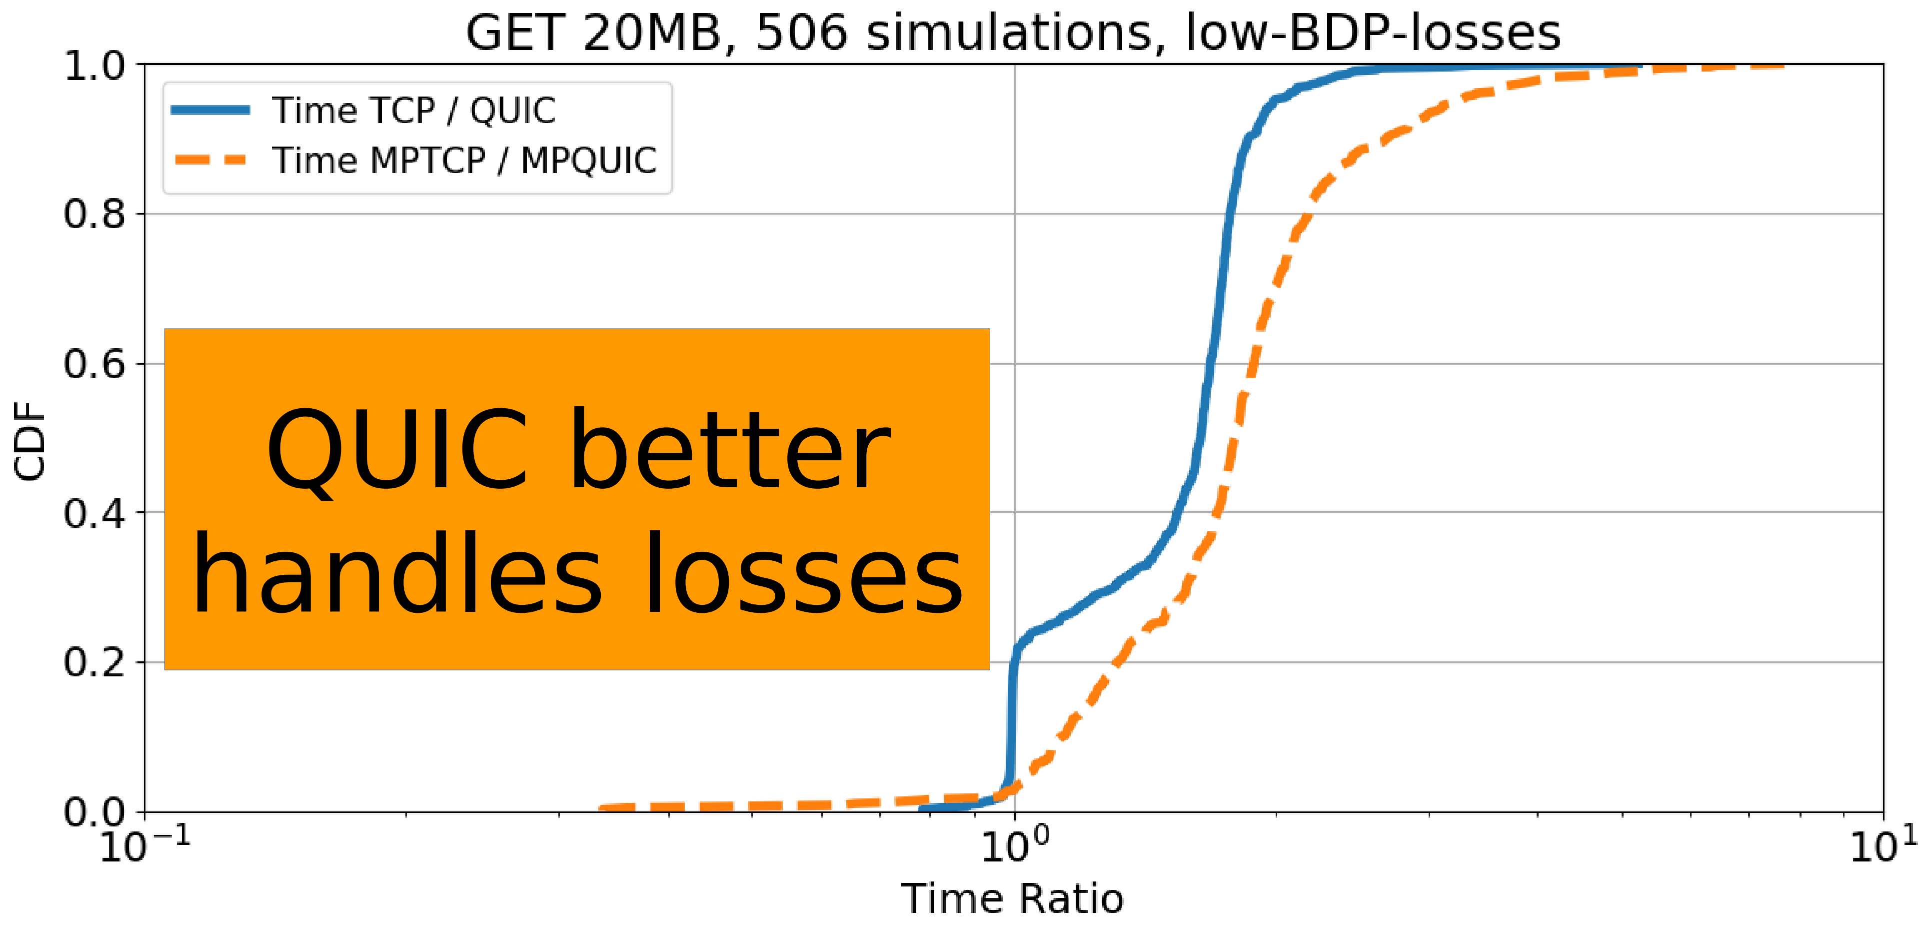

(MP)TCP vs. (MP)QUIC – Losses

CoNEXT'17

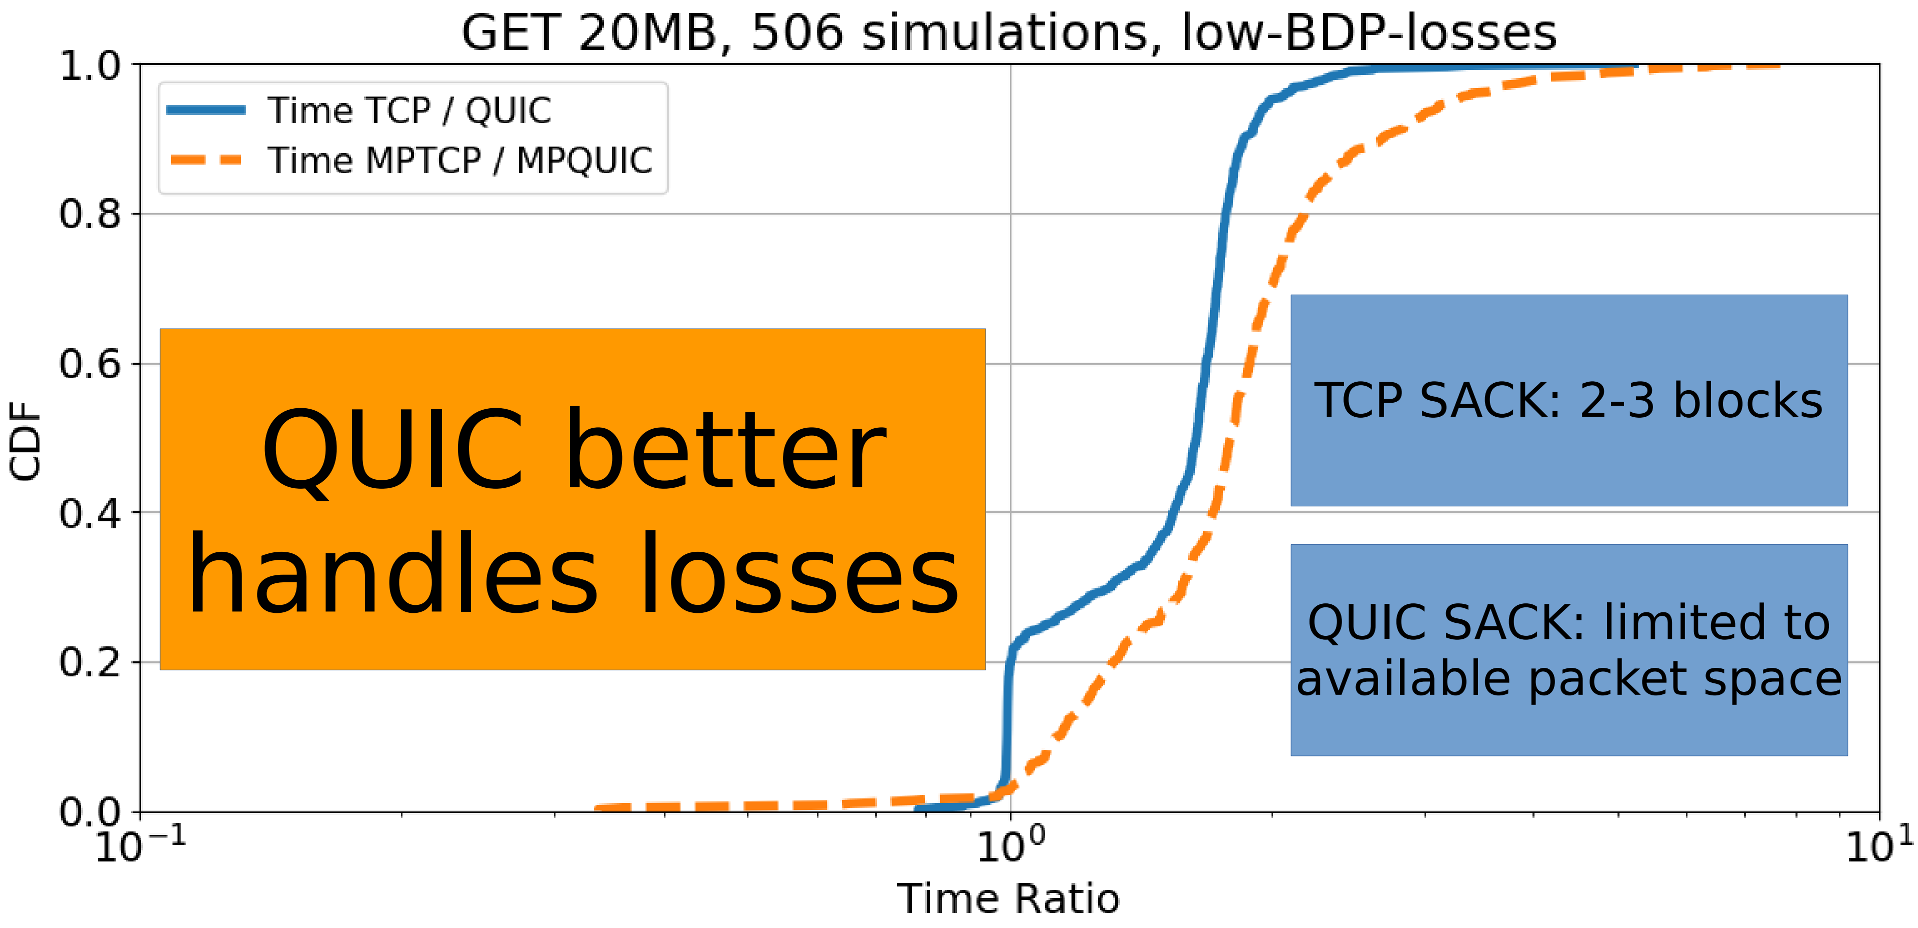

Short Files, Multipath Less Useful...