A Multipath Extension to the QUIC Module for ns-3

Shengjie Shu, Wenjun Yang, Jianping Pan, Lin Cai

Speaker: Shengjie (Shirley) Shu

June 29, 2023

QUIC Protocol

- A recently proposed transport protocol
- Address some limitations of TCP
  - connection establishment latency, head-of-line blocking, packet loss recovery, and mobility and handover support
- More promising for modern applications
- Standardized by IETF and integrated into HTTP/3

QUIC Protocol

- Key features

- ❑ **Connection Establishment:**
  - ❑ Zero round-trip time handshake
  - ❑ Encrypted connection with no additional handshake times
  - ❑ Transport Layer Security (TLS) 1.3 integration
- ❑ **Packet Header and Frame Structure:**
  - ❑ Payloads can contain various frame types
  - ❑ Stream multiplexing within a single connection
- ❑ **Loss Recovery and Error Control:**
  - ❑ Built-in retransmission and congestion control
  - ❑ Forward error correction (FEC) with an ACK frame

Multipath Scenarios

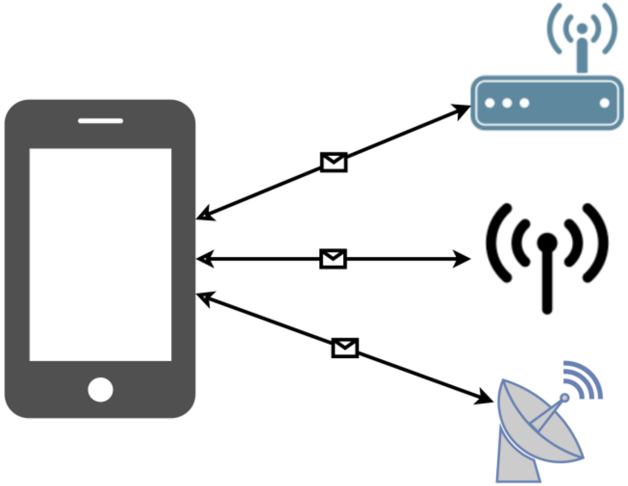

- End devices can connect with multiple network interfaces
- Potential benefits
  - Increase throughput
  - Uninterrupted communication and resilience
  - Load balancing

Multipath QUIC

- Extend the QUIC protocol to leverage multiple network interfaces simultaneously
- Aim to enhance performance, improve throughput, and fortify the protocol against link failures
- MPQUIC is under discussion by IETF
- Current experimental platforms for MPQUIC rely on either real systems or network emulators
- Absence of a hands-on MPQUIC simulation platform

Multipath QUIC

- MPQUIC follows the design logic of MPTCP and inherits the essential feature in QUIC
  - connection establishment, stream multiplexing, and frame structures

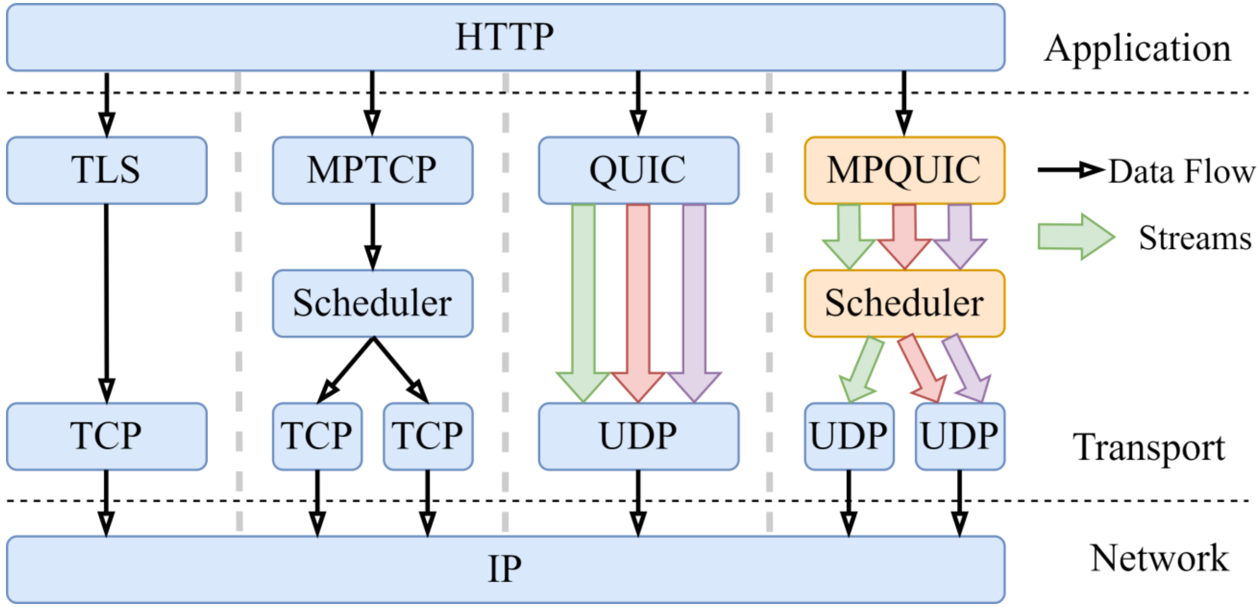

Figure 1: Structure of MPQUIC in comparison with others.

Challenge

- ❑ Address advertisement
- ❑ Path separation
- ❑ Algorithm extension
- ❑ Scheduler design

Challenge

- ❑ Address advertisement
- ❑ Path separation
- ❑ Algorithm extension
- ❑ Scheduler design

Maintain original QUIC transmission features

Our Implementation

- ❑ New classes
- ❑ New functions
- ❑ New variables

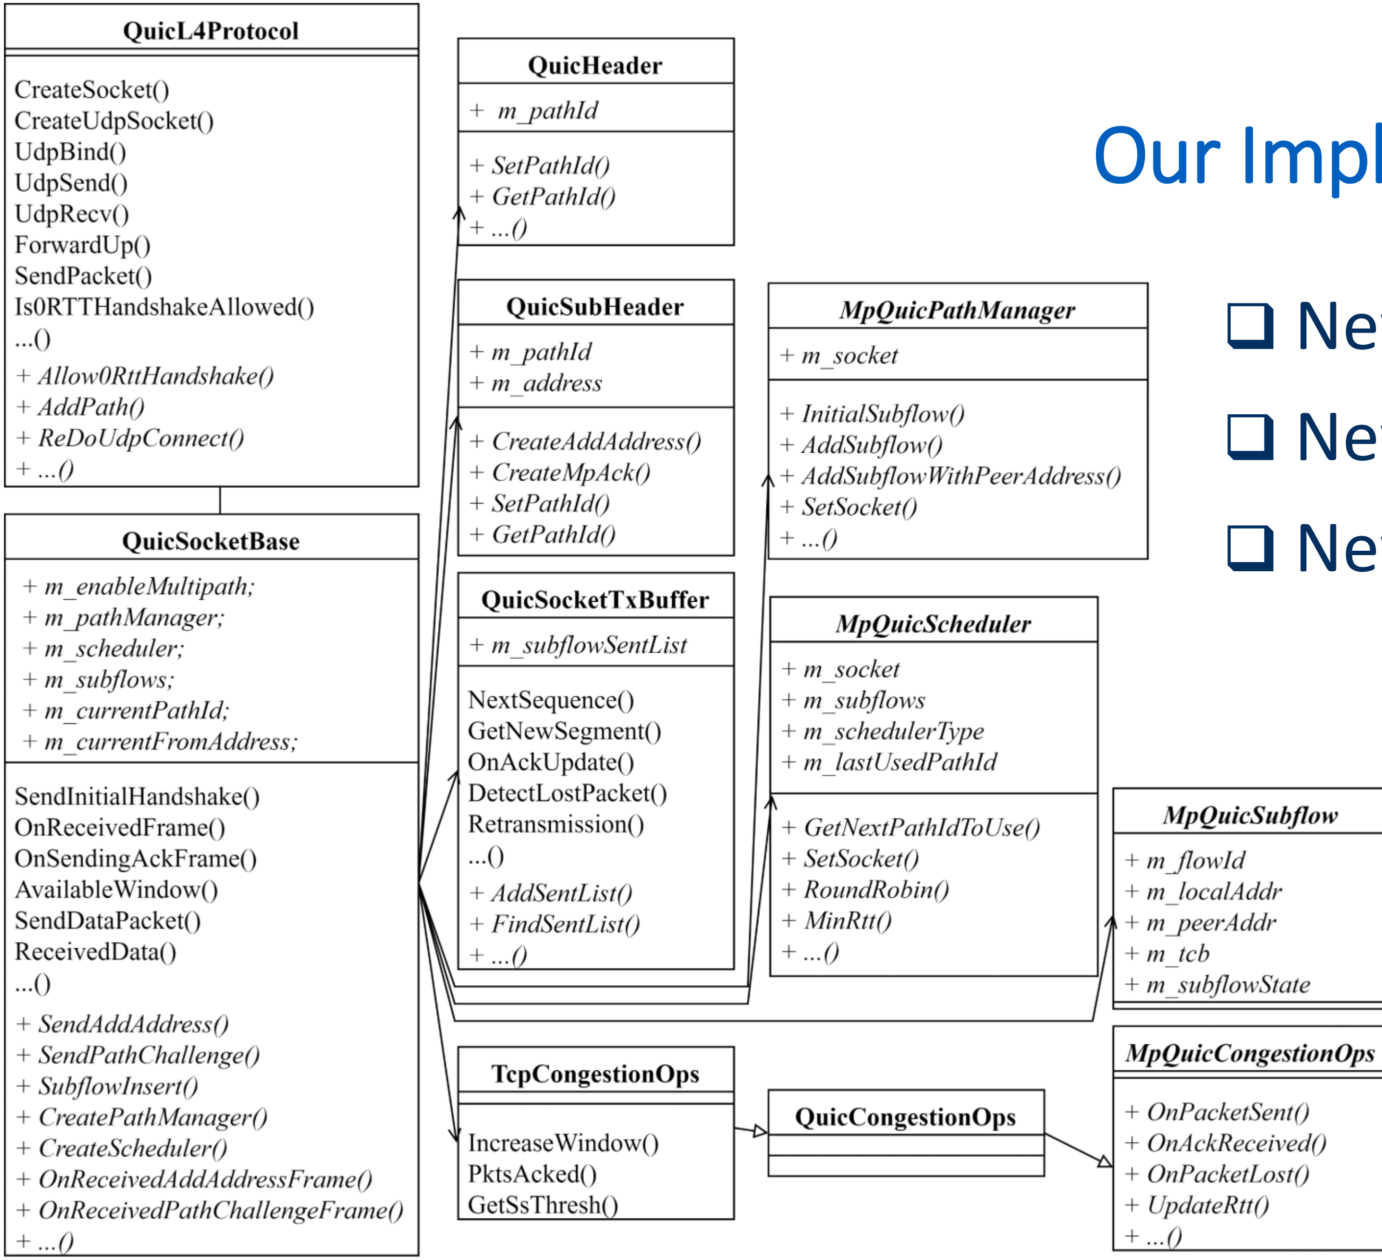

Figure 3: MPQUIC UML diagram (new classes, functions, and variables shown in italics).

Our Implementation

- ❑ New classes
- ❑ New functions
- ❑ New variables

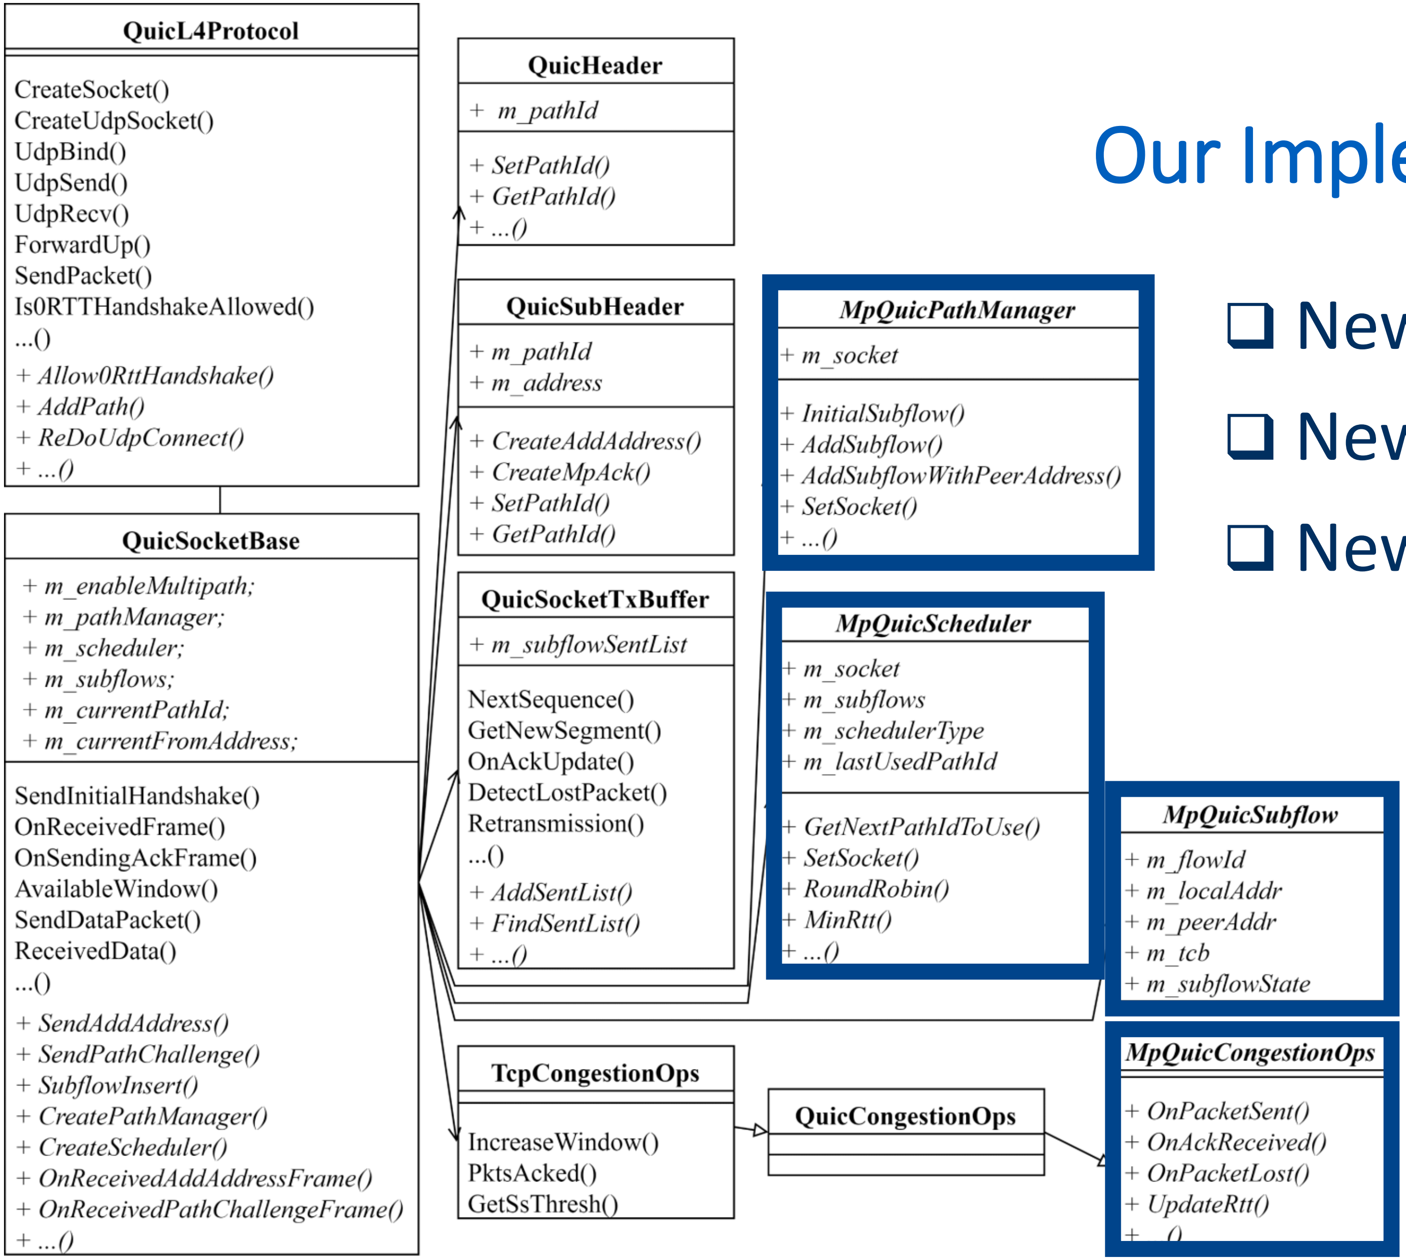

Figure 3: MPQUIC UML diagram (new classes, functions, and variables shown in italics).

Our Implementation

- New classes
- New functions
- New variables

Subflow

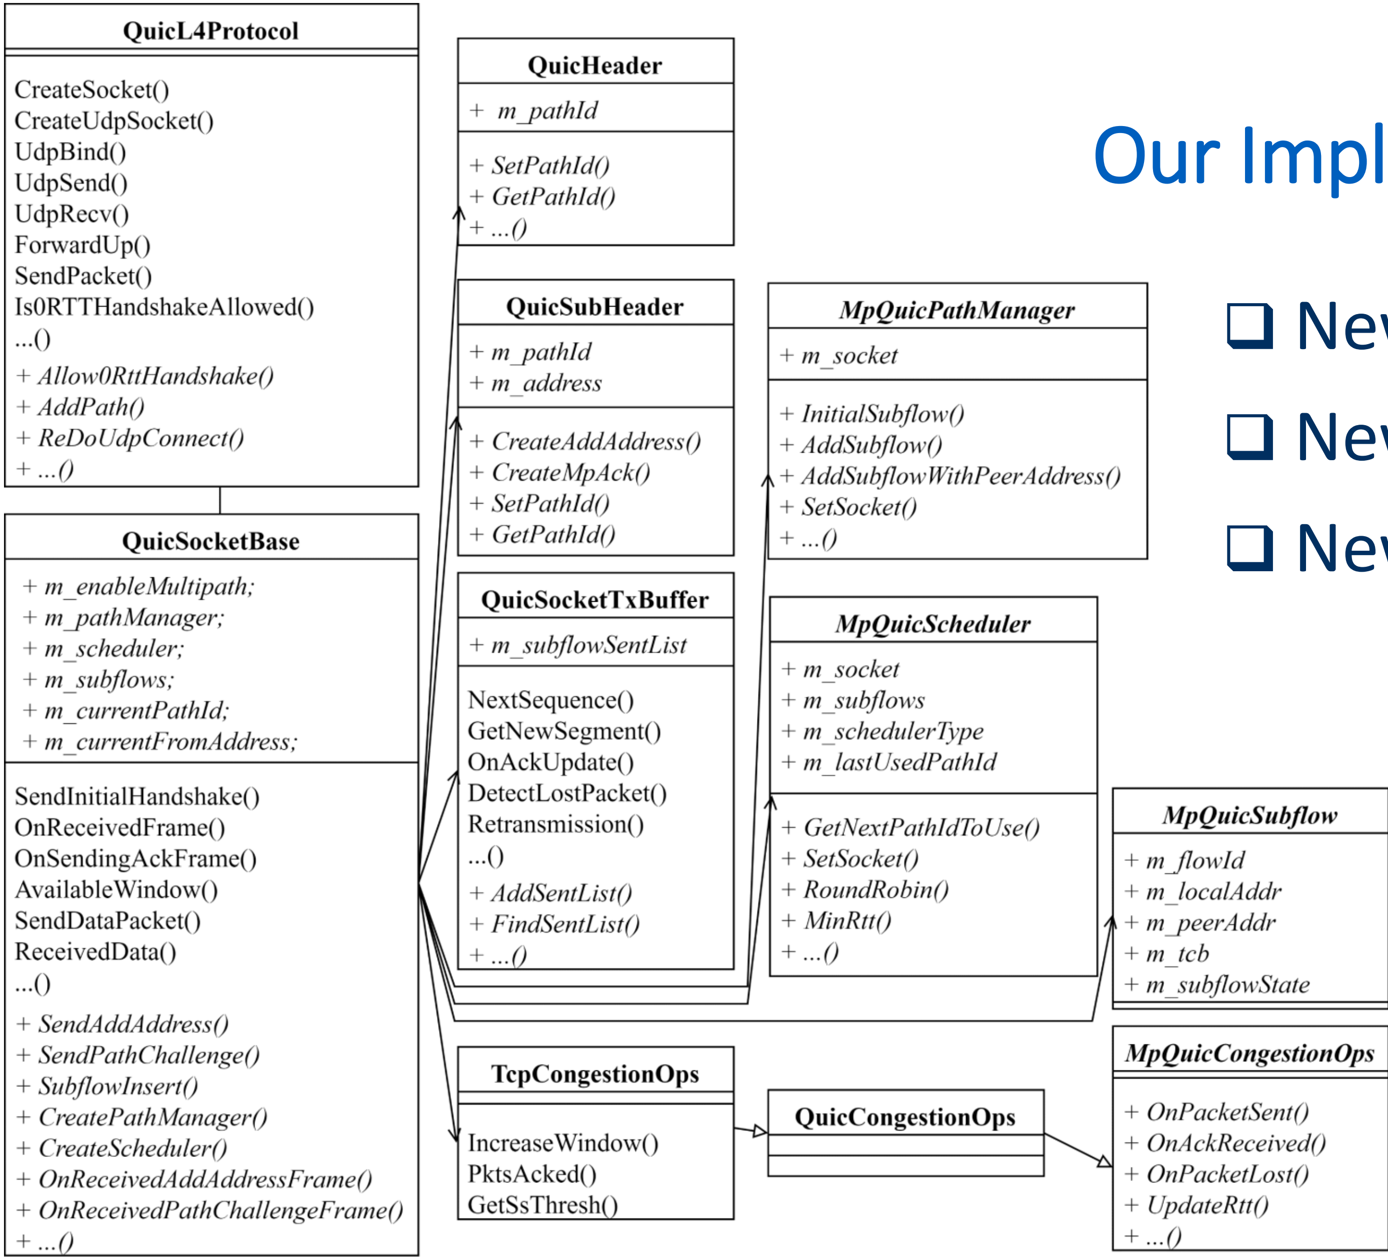

Figure 3: MPQUIC UML diagram (new classes, functions, and variables shown in italics).

Packet Header and New Frames

MPQUIC uses the frame structure to create additional sorts of frames for carrying multipath information.

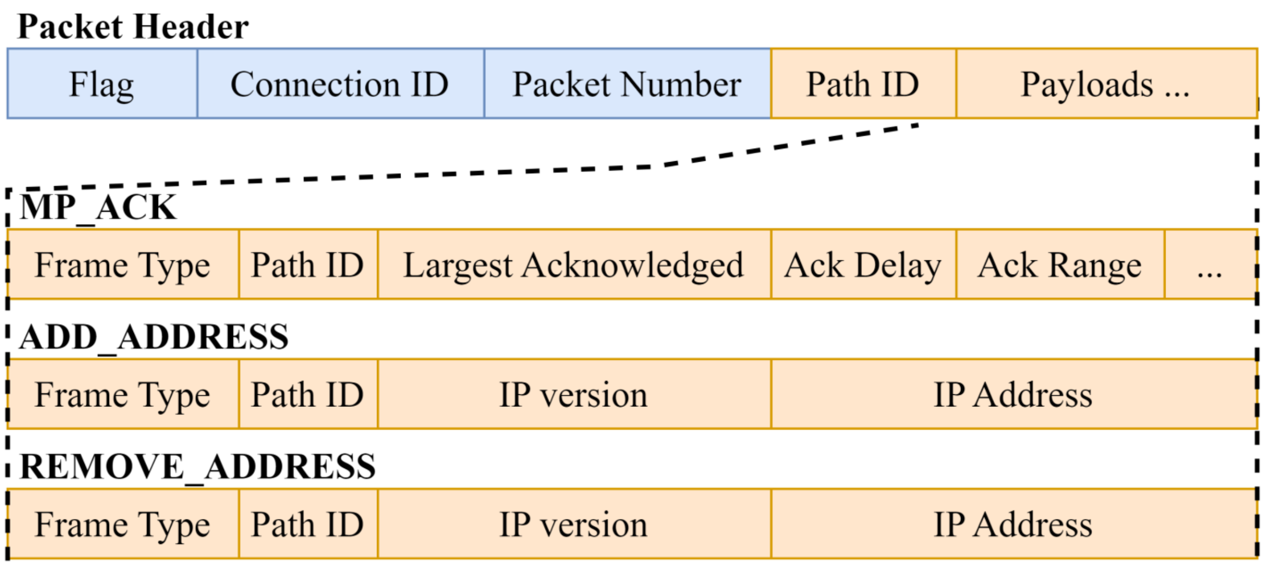

Figure 2: MPQUIC Header and New Frames.

Path Identification

- m_pathId

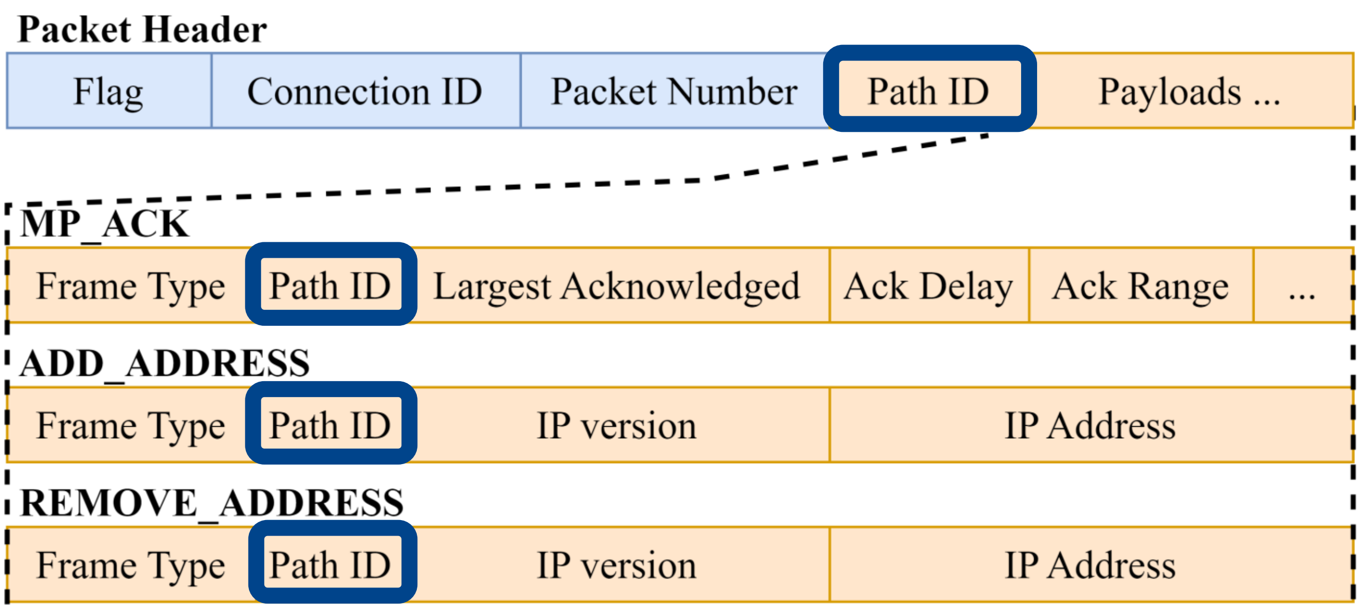

Figure 2: MPQUIC Header and New Frames.

Path Management

- ❑ New class:
  - ❑ MpQuicPathManager
  - ❑ MpQuicSubflow

m_enableMultipath=True

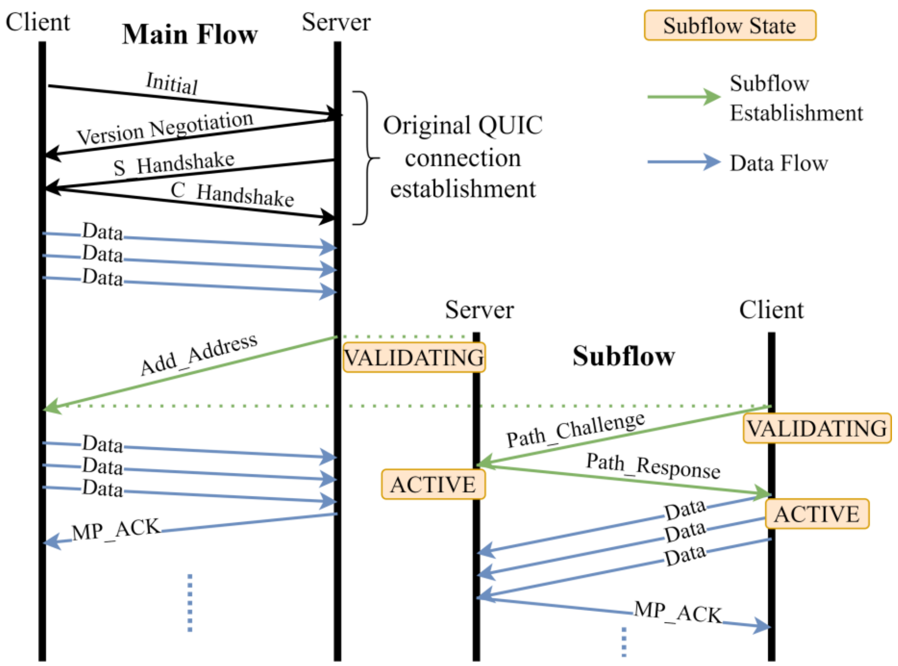

Figure 4: Procedures for subflow establishment.

Subflow State Machine

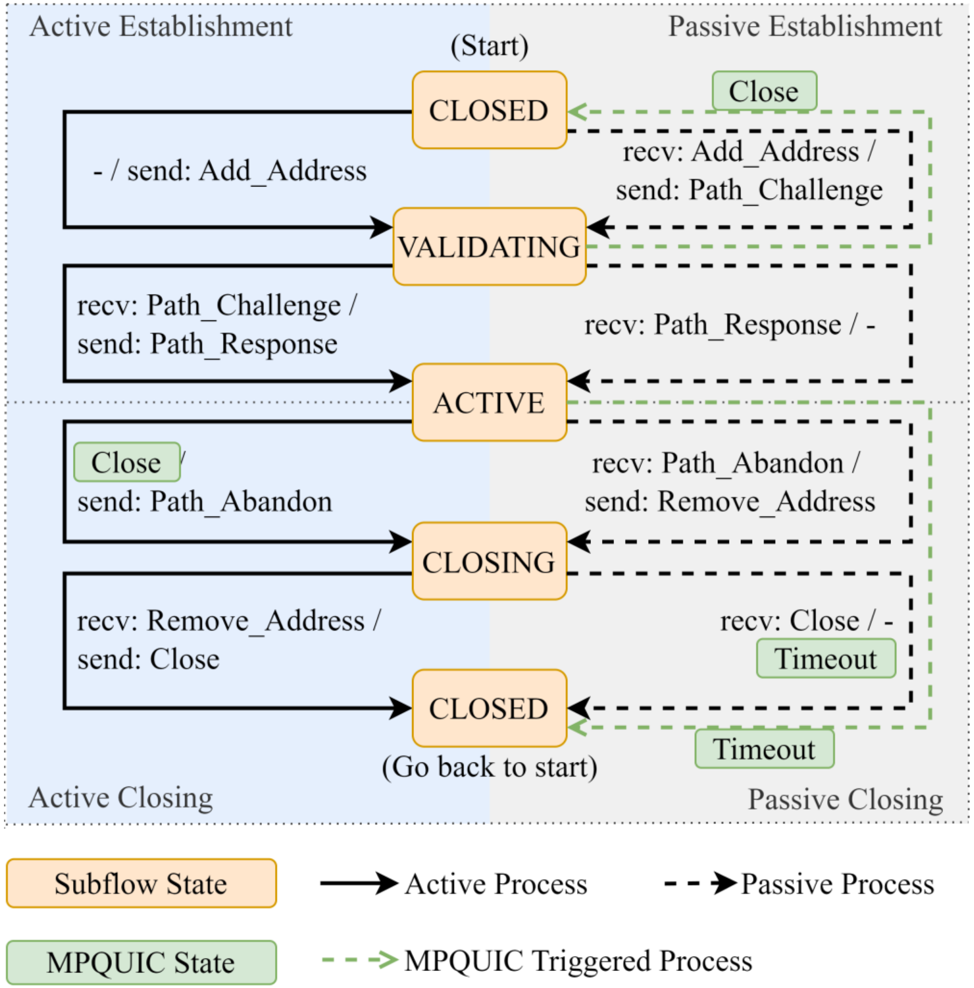

Figure 5: State machine of a subflow.

Packet Scheduling

- MpQuicScheduler: m_schedulerType

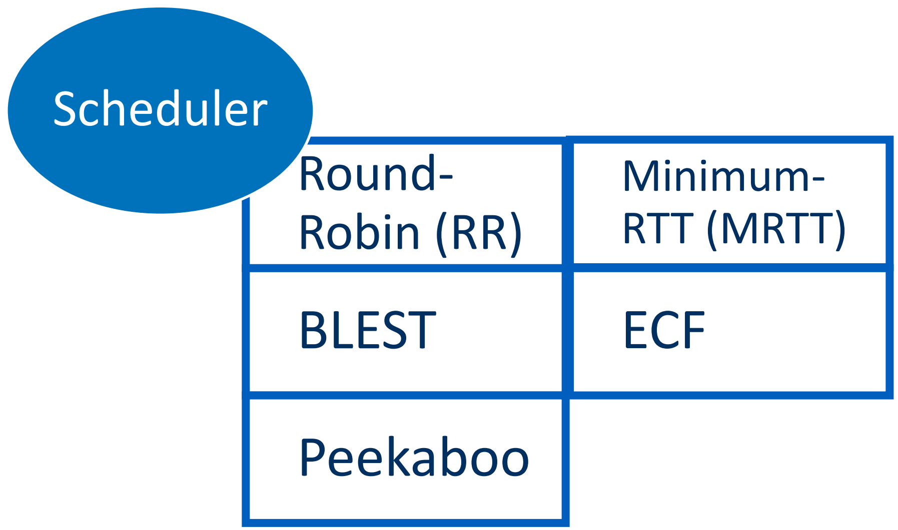

Congestion Control

- MpQuicCongestionOps
- QuicSocketBase::m_ccType

Current Status

2021

Implemented the fundamental transmission elements .

Used in some internal research projects .

2022

MPQUIC - 1.0

Fixed several issues.

Lightning talk in WNS3 2022.

Incorporated several scheduling .

2023

MPQUIC - 1.2

Evaluated the implementation with more complicated scenarios.

Future work:

Align with subsequent IETF drafts.

Compare MPQUIC implementations in emulation scenarios.

Evaluation: Scalability

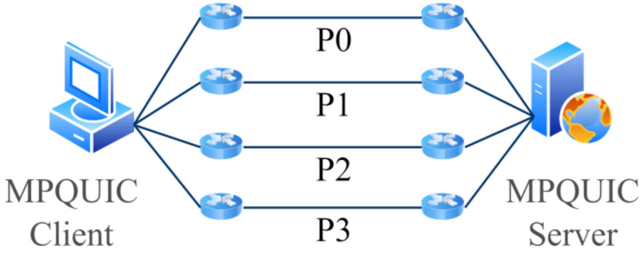

(a) Four paths

Bandwidth	One-way Delay	Loss Rate	Data Size	Repeat
5-5.5 Mbps	50-55 ms	0-0.08%	5 MB	50

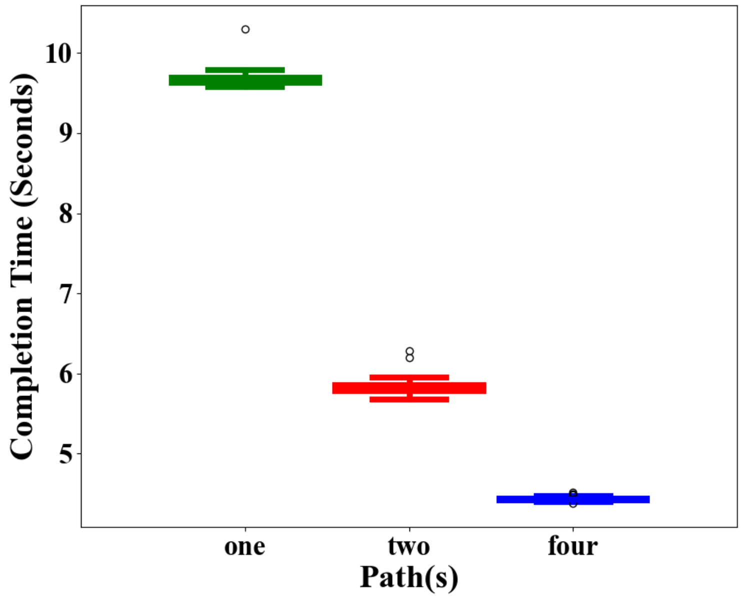

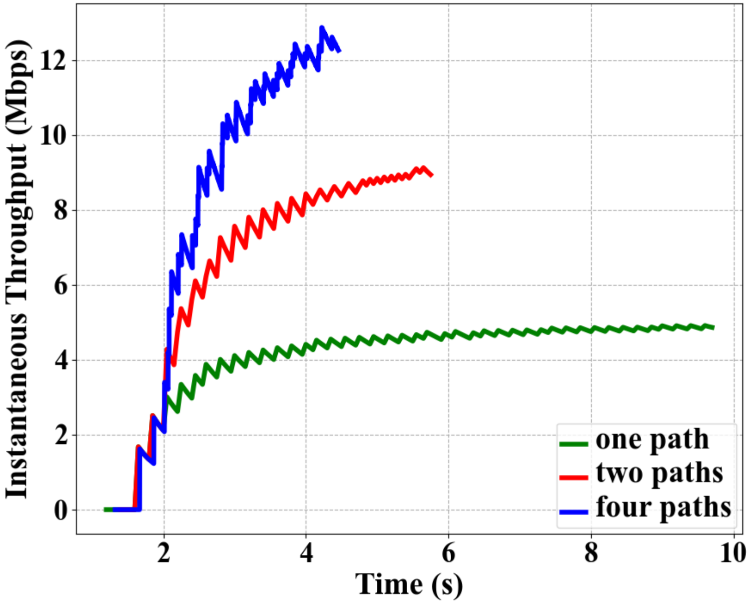

Figure 7: Completion time and instantaneous throughput comparison for one, two, and four paths.

Evaluation: Congestion Control

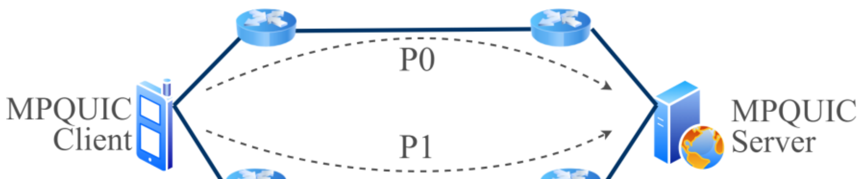

(b) Two paths

Bandwidth	One-way Delay	Loss Rate	Data Size	Duration
10-11 Mbps	10-11 ms	0-0.08%	Unlimited	50 s

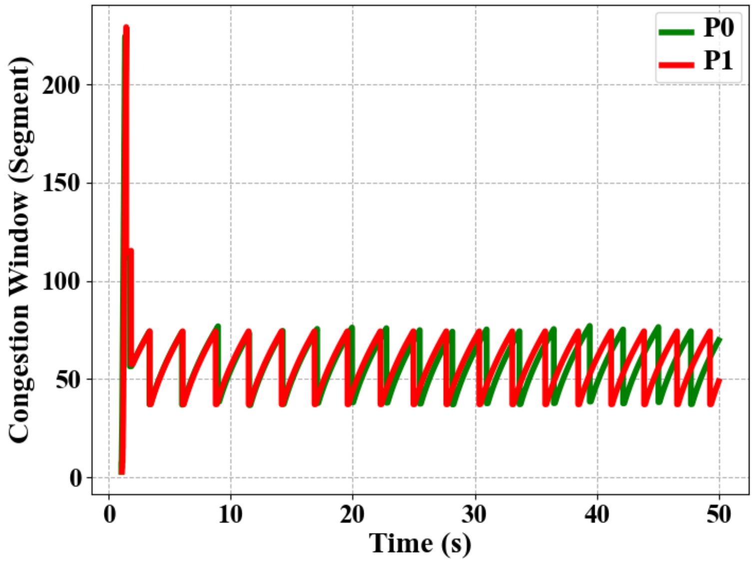

(a) NewReno

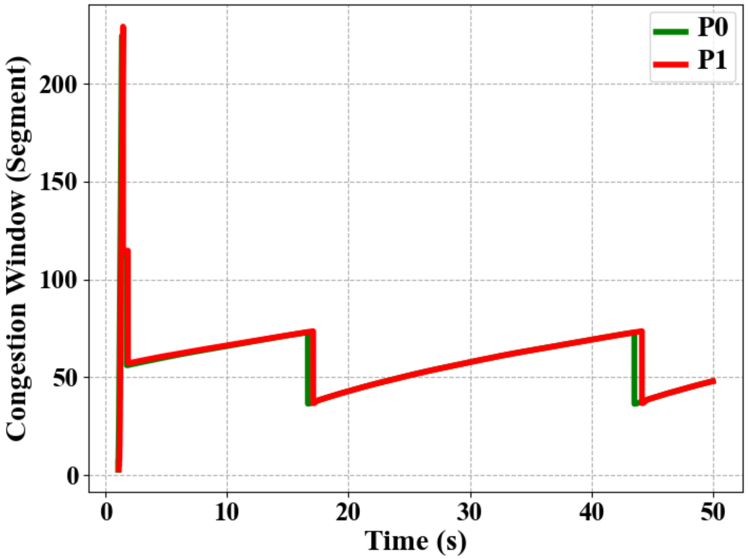

(b) OLIA

Figure 8: Congestion window comparison for NewReno and OLIA

Evaluation: Schedulers

- Dominating Scenario

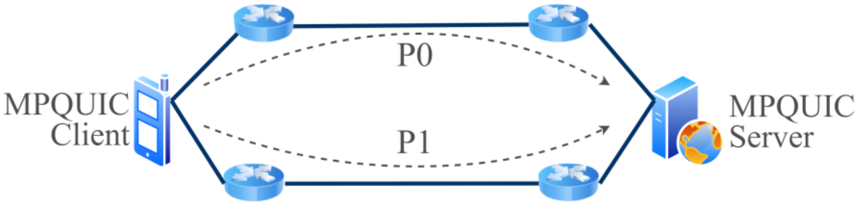

(b) Two paths

Path	Bandwidth	One-way Delay	Loss Rate	Data Size	Repeat
P0	5-5.5 Mbps	50-55 ms	0-0.08%	5 MB	50
P1	10-11 Mbps	10-11 ms	0-0.08%	5 MB	50

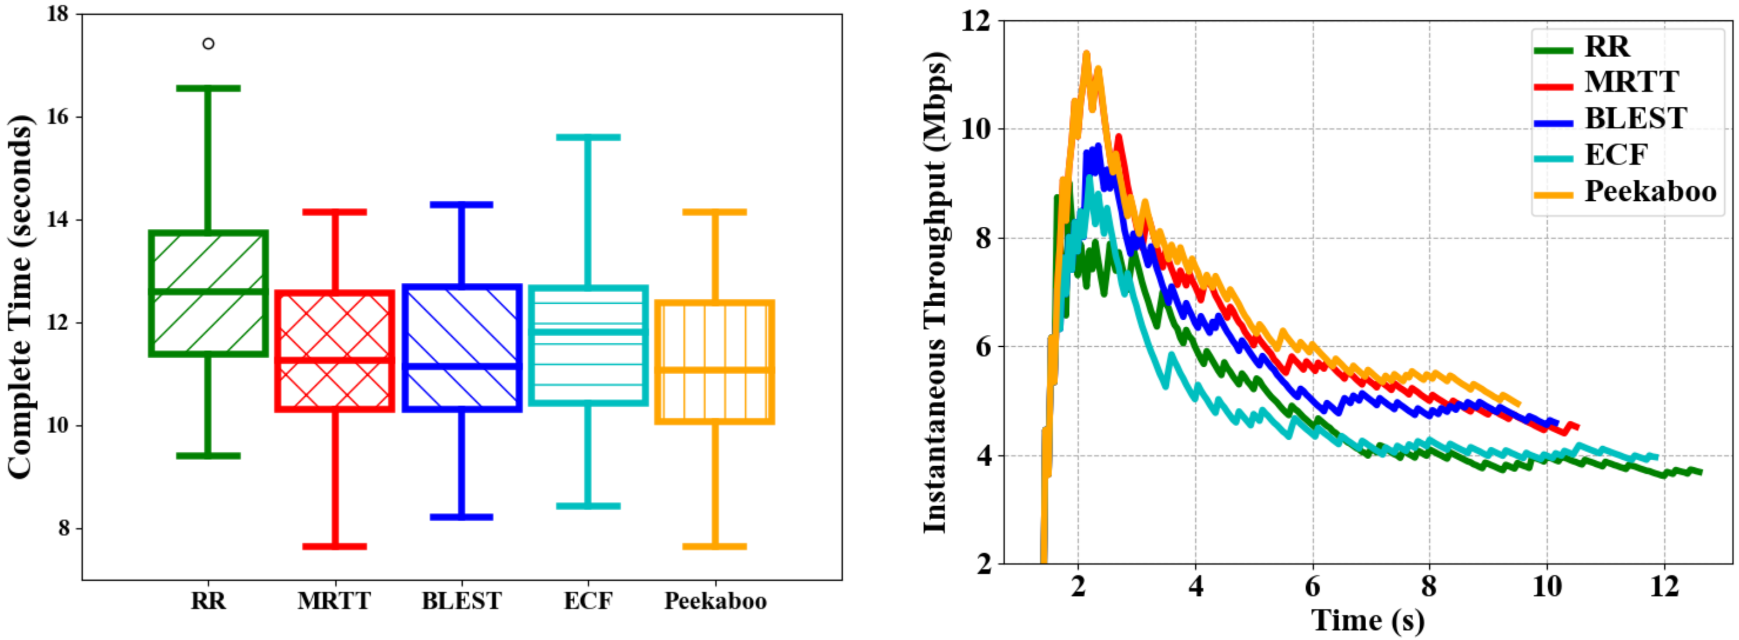

Figure 9: Completion time and instantaneous throughput comparison in the dominating scenario.

Evaluation: Schedulers

- Dominating Scenario

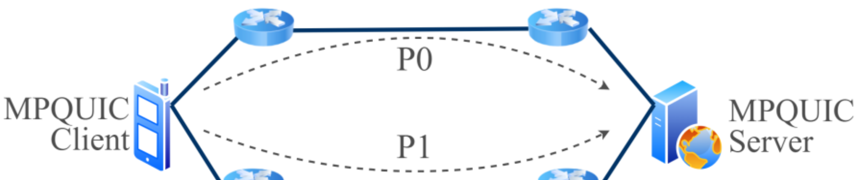

(b) Two paths

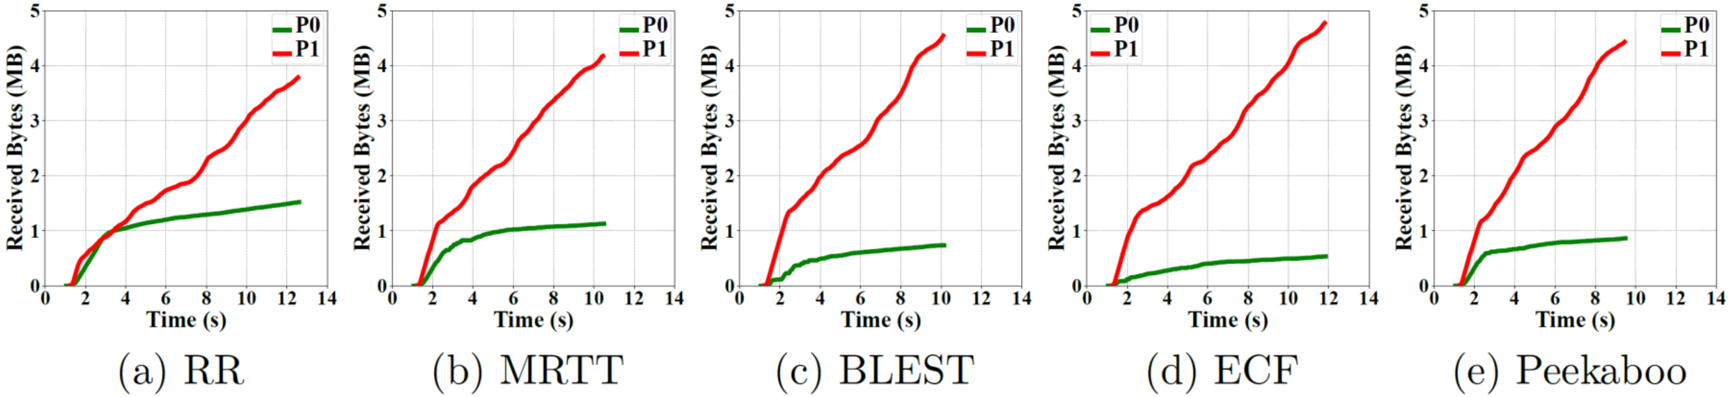

Figure 11: Received bytes of two paths in the dominating scenario

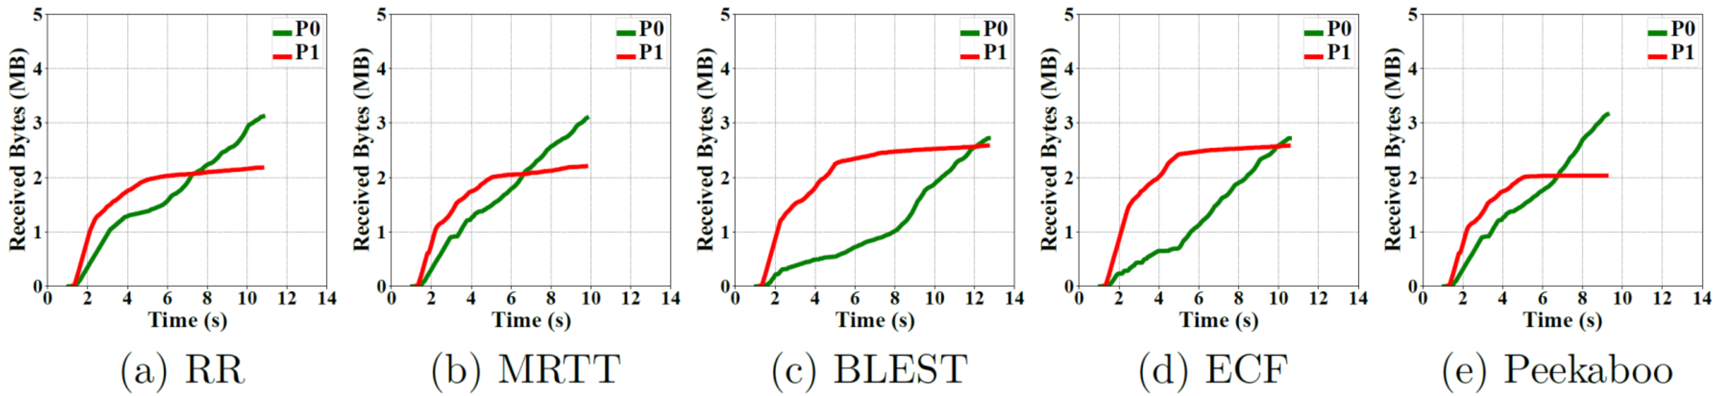

Figure 12: Received bytes of two paths under the dominating scenario with swapped setting after 5 seconds

Evaluation: Schedulers

- Competing Scenario

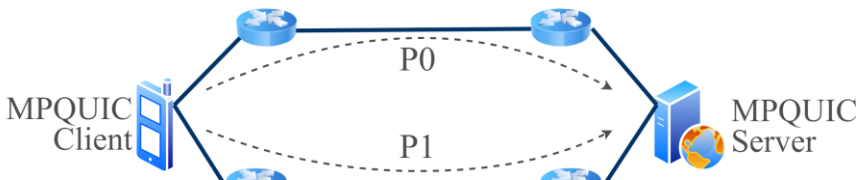

(b) Two paths

Path	Bandwidth	One-way Delay	Loss Rate	Data Size	Repeat
P0	5-5.5 Mbps	10-11 ms	0-0.01%	5 MB	50
P1	10-11 Mbps	50-55 ms	0-0.01%	5 MB	50

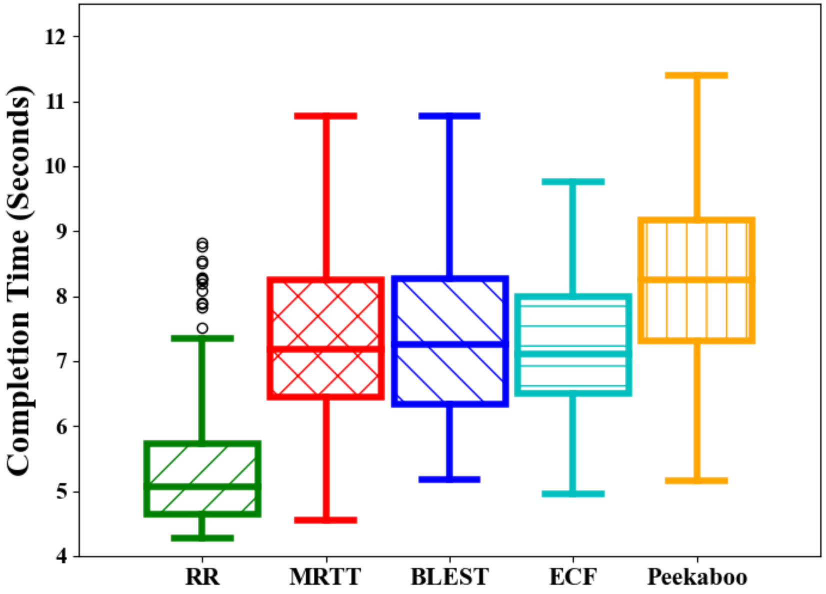

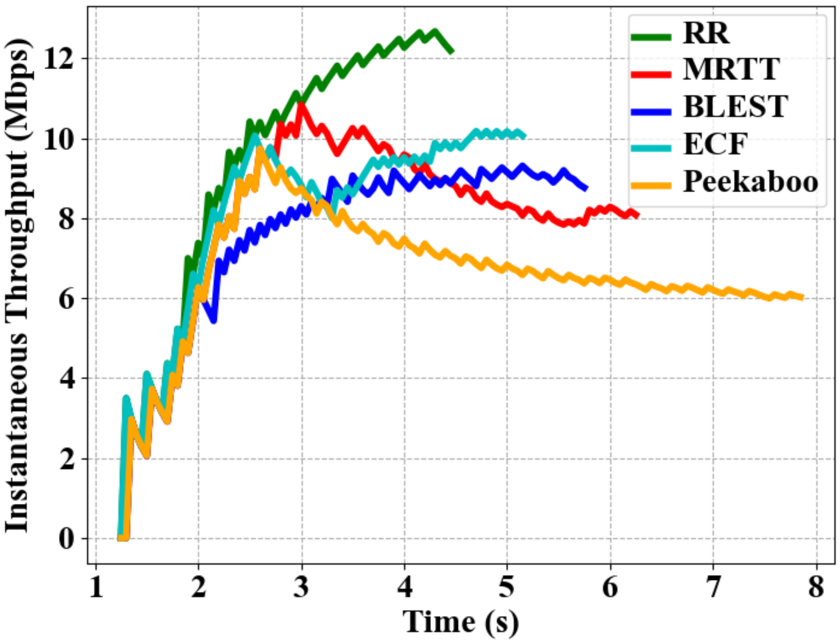

Figure 10: Completion time and instantaneous throughput comparison in the competing scenario.

Evaluation: Schedulers

- Competing Scenario

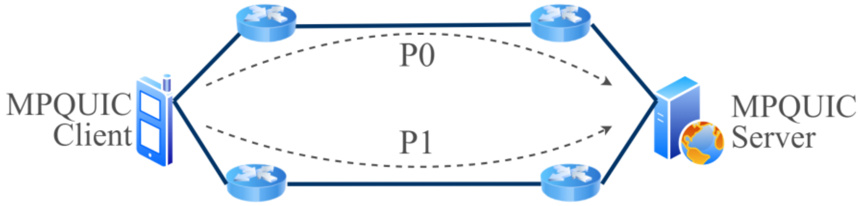

(b) Two paths

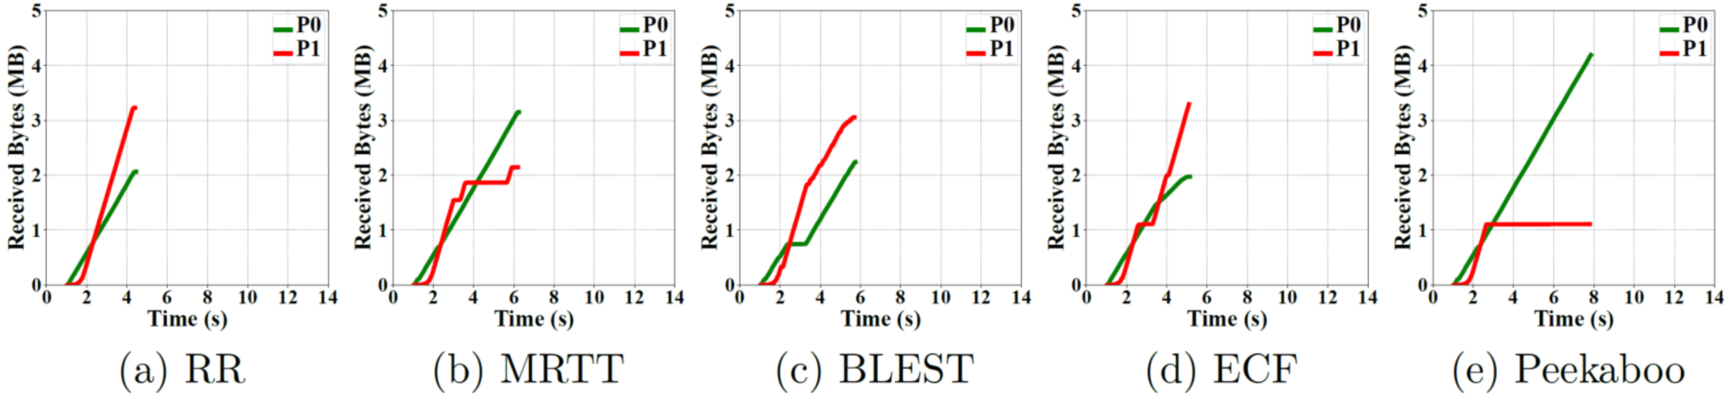

Figure 13: Received bytes of two paths in the competing scenario

Conclusions

- Provided a stable simulation platform of MPQUIC in ns-3
- Overcame the challenges of multipath transmission features
  - address advertisement, path separation, and congestion control and scheduling algorithms
- Evaluated its correctness, scalability, and flexibility with a set of experimentations

Future work:

- Align with the future IETF draft
- Compare MPQUIC implementations in emulation scenarios
- Investigate better scheduling and congestion control techniques

Thank you!