

# LAMP5G: A Link-Aware Multipath QUIC Scheduler for Heterogeneous and Disruption-Prone 5G Cellular Networks

Hongmiao Yu\*, Yuanhao Chang\*, Jorge Perez Ortiz\*, Quentin De Coninck<sup>†</sup>, K. K. Ramakrishnan\*

\* University of California Riverside, Riverside, USA

<sup>†</sup> University of Mons, Mons, Belgium

**Abstract**—Multipath transport has emerged as a promising solution to address key challenges in 5G cellular networks, including improved performance with limited RAN resources, resilience against link failures, and operator heterogeneity. While emerging 5G architectures such as NR dual connectivity (NR-DC) and multi-operator deployments potentially enable simultaneous access to multiple paths, efficiently utilizing these paths remains challenging due to dynamic link conditions, frequent handovers, and environmental impairments.

In this paper, we present *LAMP5G*, a link-aware scheduler built on Multipath QUIC (MPQUIC) for 5G cellular networks. *LAMP5G* employs a dynamic path-ranking mechanism designed for heterogeneous path environments. The ranking algorithm leverages both end-to-end transport-layer metrics—such as throughput, round-trip time (RTT), loss rate, and congestion window size—and local link-layer statistics to estimate available path bandwidth. To mitigate disruptions caused by handovers, *LAMP5G* implements a make-before-break strategy through explicit path management, enabling seamless traffic migration between RANs.

We evaluate *LAMP5G* on a 5G testbed. Experimental results show that *LAMP5G* reduces completion time by 57% compared to the dual-QUIC baseline under heterogeneous path conditions. Furthermore, during mobility scenarios, *LAMP5G* attains approximately 60% higher average throughput compared to traditional handover.

## I. INTRODUCTION

Multipath transport is a promising approach to overcome several challenges faced by applications communicating over 5G cellular networks. First is the need to provide resilience: access links—especially those using high-frequency technologies such as mmWave—are highly susceptible to environmental impairments. Multipath transport allows traffic to be rerouted over alternative paths when one path becomes obstructed. Second, we need to overcome the capacity limitations of the Radio Access Network (RAN): limited RAN resources are typically shared among many mobile devices (user equipment (UEs)), and performance can degrade significantly as the number of active UEs increases. Using multipath transport, a UE can utilize multiple RAN resources simultaneously, thus improving overall throughput. Third, operator heterogeneity: measurement studies [1] show that different cellular operators exhibit substantially different performance at the same geographic location. Multipath transport enables a UE to exploit resources from multiple operators, improving robustness and performance. Although multipath transport is not yet deployed in all

current 5G UEs and servers, several existing systems support its deployment. The architecture of the cellular infrastructure also supports multipath capability with NR Dual Connectivity (NR-DC), as defined in 3GPP TS 37.340 [2]. It allows a UE to simultaneously connect to two RANs that are anchored to the same 5G core network and to transmit data over both RANs concurrently. A further capability being explored is support for resource sharing among multiple operators [3], [4]. With the widespread deployment of technologies such as O-RAN, together with UE support for dual-SIM and eSIM capabilities, a UE can potentially leverage network resources from multiple operators. In this scenario, the UE connects to multiple RANs that are attached to different cores operated by different cellular operators. Given the performance, resilience, and flexibility benefits offered by multipath transport, together with the feasibility of these deployment architectures, a key remaining challenge is how to schedule packets to be able to utilize these multiple heterogeneous paths efficiently.

A key challenge for packet scheduling in multipath 5G cellular networks is highly dynamic link quality. Depending on the UE's location, different RANs may provide different coverage and performance, and the use of heterogeneous access technologies (*e.g.*, sub-6 GHz, mmWave, or WiFi) further amplifies path diversity. The transport scheduler must therefore exploit this heterogeneity effectively. A second challenge is disrupted transmission during handover. Since current 5G inter-gNB handovers are typically break-before-make, the UE disconnects from the source gNB before connecting to the target gNB, causing packets to be buffered at the source gNB or UPF [5]. Packet drops during this process may also reduce the transport-layer congestion window, further degrading performance [6]. In multipath transport, the scheduler can mitigate such disruption by redirecting traffic to other available paths. A third challenge is environmental impairment: high-frequency links such as mmWave are vulnerable to blockage and rapid environmental dynamics, which can sharply degrade link quality or cause temporary outages. Thus, a 5G multipath transport scheduler must adapt to dynamic link conditions while efficiently distributing traffic across available paths.

Another key design choice is the transport protocol. Compared with MPTCP [7], MPQUIC [8], [9] offers finer-grained scheduling and control by organizing packets into frames, which carry data and control functions. Unlike MPTCP's

monolithic byte stream, MPQUIC multiplexes independent streams and control frames within the same packet, reducing head-of-line (HoL) blocking across streams through SPDY-style transport-layer streams [10]. MPQUIC also decouples frames from network paths, allowing traffic from the same connection to be scheduled across different paths based on path conditions and packet requirements. This flexibility improves heterogeneous-path utilization. Finally, MPQUIC provides built-in confidentiality and authenticity, avoiding additional transport-layer security mechanisms.

In this paper, we present *LAMP5G*, a link-aware MPQUIC scheduler for 5G cellular networks. *LAMP5G* combines transport-layer metrics (e.g., throughput, RTT, packet loss, and congestion window) with local link-layer statistics to schedule packets across heterogeneous paths and adapt to dynamic conditions, including handovers and environmental changes. Our main contributions are:

**Link-aware Scheduling.** Link conditions strongly affect cellular-path bandwidth. To schedule traffic more effectively, *LAMP5G* combines end-to-end transport metrics with local link-layer status (e.g., link-quality indicators) to estimate each path’s bandwidth, assuming the wireless link is the bottleneck. It then ranks paths by the estimated bandwidth and schedules packets sequentially on the highest-ranked path, subject to congestion-window constraints.

**Environmental Impairments.** A throughput predictor [11] leverages the local view of link-layer status to anticipate short-term throughput fluctuations caused by environmental impairments. *LAMP5G* integrates these predictions into its scheduler by dynamically lowering the ranking scores of impaired paths, proactively shifting traffic toward more reliable paths. *LAMP5G* achieves 57% faster completion compared to the dual-QUIC baseline under a heterogeneous cellular environment.

**Seamless Handover.** *LAMP5G*’s multipath transport mitigates handover-induced data disruption by explicitly managing paths in a make-before-break manner. Once the target RAN satisfies required conditions (e.g., sufficient signal strength), the UE establishes a new path through it while retaining the source-RAN path until it is no longer suitable for efficient transmission. Thus, even if the source connection is disrupted or terminated, data can continue over the newly established target path. During handover, *LAMP5G* maintains seamless, high-throughput transmission, achieving 60% higher average throughput than traditional handover.

## II. BACKGROUND AND RELATED WORK

### A. Disruptions in 5G NR networks

Although 5G NR enables higher peak data rates, its radio links can be more disruption-prone, especially for high-frequency technologies such as mmWave, which are sensitive to blockage, mobility, and environmental dynamics. Handovers further disrupt 5G data transmission, causing latency spikes, packet reordering, and temporary throughput degradation.

### 1) Disruptions in High-Frequency Wireless Channels

Although high-frequency wireless technologies offer higher bandwidth and lower latency, they are more vulnerable to environmental impairments such as penetration loss, rain attenuation, and blockage. Measurement studies [12], [13], [14], [15], [16] show that high-frequency networks, especially mmWave, are highly sensitive to environmental conditions, causing throughput fluctuations and link instability.

Lower-layer solutions, including intelligent reflecting surfaces [17] and advanced beam management [15], [18], improve mmWave resilience but mainly strengthen individual links. In contrast, multipath transport improves reliability at the transport layer by shifting traffic to alternative paths when high-frequency links degrade. Prior work [19], [20], [11] has shown its effectiveness in mitigating such disruptions. Building on this, we study how multipath transport can improve reliability and performance in disruption-prone 5G cellular networks.

### 2) Disruption of Handover

Another major disruption source in 5G NR is handover. 3GPP handover follows a break-before-make model: in inter-gNB handover, such as N2-based handover, the UE releases source-gNB radio resources and synchronizes with the target gNB before user-plane transmission is restored. During this interval, no active air interface is available, so packets arriving at the source gNB must be forwarded to the target gNB. 3GPP defines two forwarding modes: Indirect Data Forwarding, where packets are tunneled through the core network, typically via the UPF, and Direct Data Forwarding, where packets are sent over the Xn interface when available.

Several measurement studies [6], [21] show that handovers can significantly degrade end-to-end performance through three main effects. First, buffer overflow at the source gNB, forwarding delays, and packet drops can cause packet loss or retransmission timeouts (RTOs), triggering transport-layer retransmissions. Second, such losses or RTOs can shrink the congestion window, causing temporary throughput degradation and slowing recovery. Third, after the UE attaches to the target gNB, packets may arrive from both the forwarding path, especially under indirect forwarding, and the new target-gNB path, causing out-of-order delivery and further degrading transport performance [6].

Prior work mitigates handover impact through proactive handover detection and transport-parameter configuration [6] or improved data-transmission procedures [5]. Multipath transport can further reduce handover-induced degradation: when one path undergoes handover, other active paths can continue carrying traffic, allowing ongoing transmissions to proceed with minimal interruption.

### B. Architecture for Multipath 5G Networks

Emerging 5G architectures can support multipath transport. A representative example is NR-DC, standardized in 3GPP TS 37.340 [2], which allows a UE to connect simultaneously to two NR base stations: a Master Node (MN) and a Secondary Node (SN). The MN handles control-plane signaling and connection management, while the SN provides additional

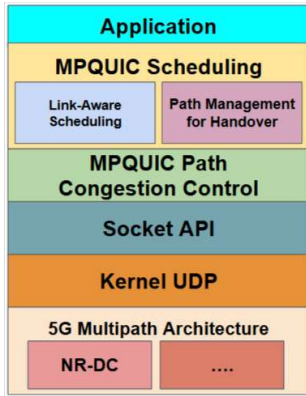


Fig. 1: Design Overview of *LAMP5G*

user-plane resources. Through PDCP-layer integration, traffic can be split across the MN and SN, enabling the UE to aggregate radio resources and improve throughput, reliability, and resilience. Prior work [22], [23] has explored running transport-layer multipath protocols over NR-DC by treating MN and SN traffic as separate transport paths. However, the use of NR-DC to enable more seamless handover remains largely unexplored.

### III. OVERVIEW

Figure 1 illustrates the architecture of *LAMP5G*. Implemented on top of MPQUIC, *LAMP5G* integrates link-aware scheduling with dynamic path management to support efficient multipath utilization and seamless handover in 5G cellular.

#### A. Link-Aware Scheduling

##### 1) Bandwidth based Path Ranking Scheduler

We design a link-aware scheduler to address path heterogeneity and disruptions in 5G cellular networks. At the core of our design is a path-ranking mechanism based on predicted loss-free end-to-end bandwidth. Ideally, data packets in a heterogeneous environment should be transmitted over the “better” path. To quantify the “goodness” of a path, we adopt the concept of *Estimated Loss-Free End-to-End Bandwidth* from [24]. This metric is computed based on the available congestion window, loss rate, and RTT, capturing the end-to-end transport-layer characteristics of a path. However, these metrics alone cannot reflect short-term and transient disruptions caused by environmental factors such as blockage in a timely manner. In such cases, bandwidth degradation may occur before the congestion window can adapt. To address this limitation, we incorporate a short-term link-layer throughput predictor based on local link-layer observations as well. This predictor captures sudden and transient performance changes that are not immediately reflected in transport-layer metrics. By combining estimated end-to-end loss-free bandwidth with predicted short-term link layer throughput, we obtain a more responsive and accurate ranking score for each path. The detailed design is described in Sec. IV-A.

Another challenge in heterogeneous multipath environments is out-of-order packet delivery. Differences in RTT across paths can cause earlier packets sent over a slower path to arrive later than subsequently transmitted packets sent over a

faster path, leading to reordering and throughput degradation if traditional single-path transport protocol congestion control mechanisms are adopted. To mitigate this issue, our scheduler continues to transmit packets over the highest-ranked available path as long as its ranking remains stable. Sustained transmission on a single high-quality path increases the likelihood of in-order delivery and reduces reordering overhead. The scheduler reevaluates path rankings when significant events occur, including: (1) Changes in RTT or loss rate: These delay-based metrics indicate evolving path conditions that may reduce the suitability of the current path. Since these metrics evolve relatively smoothly, estimated end-to-end loss-free bandwidth typically remains stable for a short duration, enabling continuous transmission. (2) Sudden drops in predicted link layer throughput: This signals environment-related disruption of the wireless channel, prompting the scheduler to avoid the affected path. (3) Significant congestion window reduction: This reflects reduced available bandwidth and also helps guard against false positives from throughput prediction. Through this adaptive ranking mechanism, *LAMP5G* achieves robust and efficient scheduling in heterogeneous and disruption-prone 5G multipath environments.

#### B. Path Management for Seamless Handover

Traditional single-path handovers suffer from packet loss and retransmissions, congestion window reductions, and out-of-order delivery. These effects can significantly degrade transport-layer performance. In contrast, multipath transmission in cellular networks provides an opportunity to mitigate these issues by avoiding reliance on data forwarding mechanisms inherent in conventional handover procedures. Instead of performing a break-before-make transition, the UE can initiate a new path through the target gNB while maintaining the existing path through the source gNB (Fig. 2). With both paths active simultaneously, data transmission can continue without interruption, resulting in a more stable connection than a single-path solution. Once the source gNB becomes unsuitable for communication (e.g., due to low signal strength), the old path can be gracefully removed. Because data transmission continues over the newly established path, there is no need for packet forwarding from the source gNB to the target gNB. To support this approach, coordination with lower-layer mechanisms is required. In this paper, we design our solution based on NR-DC, specifically leveraging the inter-MN handover without an SN change procedure. We provide detailed design and implementation considerations in Sec. V.

### IV. DETAILED DESIGN OF *LAMP5G* COMPONENTS

#### A. Link-Throughput-Aware Path Ranking Scheduler

An effective path-quality metric is essential for multipath scheduling. Following [24], *LAMP5G* characterizes each path using its “Estimated Loss-Free End-to-End Bandwidth”, computed from the congestion window  $cwnd_i$ , Round-trip time (RTT)  $RTT_i$ , and loss probability  $p_i$  as  $\frac{cwnd_i(1-p_i)}{RTT_i}$ . This favors paths with higher available bandwidth and lower

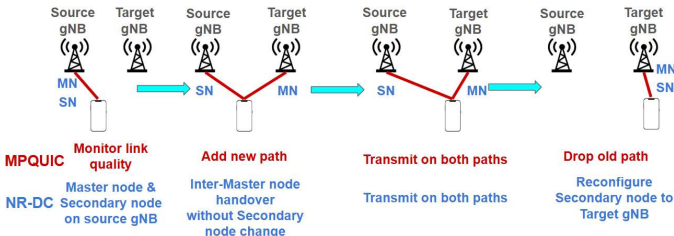


Fig. 2: Seamless Handover in LAMP5G.

loss. However, environmental impairments can reduce wireless capacity before transport-layer metrics react. To capture such short-term changes, LAMP5G incorporates AI-based 5G throughput prediction [25], [26], which uses link-layer indicators such as signal strength, UE mobility, and radio type. The scheduler then ranks each path using the minimum of the transport-layer estimate and the predicted throughput:

$$\min \left( \frac{cwnd_i(1-p_i)}{RTT_i}, T_i \right),$$

where  $T_i$  denotes the predicted throughput of path  $i$ .

Since an AI-based throughput predictor cannot achieve perfect accuracy, the impact of prediction errors must be considered. We analyze false predictions from two perspectives: (1) Overestimation, where the predicted throughput is significantly higher than the true available bandwidth. In this case, the “Estimated Loss-Free End-to-End Bandwidth” becomes the limiting factor for the ranking score. Because the scheduler selects the minimum value between the two metrics, overestimation in the predicted throughput does not adversely affect scheduling decisions. (2) Underestimation, where the predicted throughput is significantly lower than the true available bandwidth. In this scenario, the ranking score of an otherwise high-quality path may be incorrectly reduced, potentially causing the scheduler to select an alternate suboptimal path. However, our scheduling strategy restricts the amount of data transmitted on each path by its available congestion window. This prevents excessive transmission on a lower-ranked path. The congestion window on the better (under-estimated path) will still be open and thus enable multipath transport to continue to utilize that path. This limits the impact of temporary underestimation by the throughput predictor. The congestion control protocol’s window-based adaptation and packet transmissions naturally complement the short-term link-layer predictor despite potential inaccuracies.

## V. PATH MANAGEMENT FOR SEAMLESS HANDOVER

LAMP5G flexibly manages MPQUIC paths to mitigate disruptions caused by traditional handover procedures. Figure 2 illustrates the multipath path management process during handovers with LAMP5G. Consider a UE moving from a source gNB to a target gNB. Initially, the UE continuously monitors the link quality for both gNBs. Once the link quality of the target gNB becomes suitable for data transmission, the UE initiates the creation of a new path through the target gNB. In MPQUIC, path establishment is performed by probing with a PATH\_CHALLENGE frame to the peer. In our setting, the UE can probe this frame via the existing path through

the source gNB to establish the new path. After creating the new path, both paths are utilized concurrently, as long as their link qualities remain acceptable. When the link quality of the source gNB degrades below a usable threshold, the corresponding path is removed. Similar to path creation, a PATH\_ABANDON frame is sent by the UE to notify the peer to drop that path.

### A. Triggering Path Addition & Removal

A critical design question that significantly affects efficiency is determining when to add or remove paths. Ideally, a path through the target gNB should be established as early as possible once the link quality is sufficient to support data transmissions, while the path through the source gNB should be maintained as long as its link quality remains acceptable. The path addition and removal mechanism in LAMP5G is independent of the specific triggering strategy used for handover decisions, allowing LAMP5G to support a wide range of triggering policies. Here, we provide two such policies for path addition and removal:

- (1) Event-Driven Triggering. Handover decisions are traditionally driven by RRC-layer events based on UE signal-quality measurements. For example, A3 triggers handover when a neighboring cell’s signal strength exceeds the serving cell’s by a configured offset. LAMP5G adopts a similar event-driven strategy for path management: A4 triggers path addition when a neighboring cell exceeds a predefined signal threshold, while A2 triggers path removal when the serving gNB’s signal quality falls below a threshold.
- (2) Predicted-Throughput-Based Triggering. The throughput predictor provides an estimate of each path’s link quality, which LAMP5G uses for path management. The UE adds a new path when the predicted throughput of a target link exceeds a predefined threshold, and removes an existing path when its predicted throughput falls below a threshold.

### B. Lower Layer Support for Multipath Seamless Handover

Although LAMP5G focuses on transport-layer path management, lower-layer support is needed to enable such path transitions. We present one 3GPP-compliant example based on NR-DC. As shown in Fig. 2, the UE initially transmits through a gNB where both the MN and SN are colocated, which is allowed by TS 37.340 [2]. When the target gNB satisfies the required link-quality conditions, an inter-MN handover without SN change is performed: the MN moves from the source to the target gNB, while the SN remains at the source. During this process, the UE continues transmitting over the SN path, avoiding data interruption. After MN handover completes, a new transport-layer path is established through the target MN, allowing both paths to be used for higher aggregate throughput. When the source-gNB link degrades, the UE removes the corresponding transport-layer path, releases the SN connection, and reassigns the SN to the target or another gNB. Although this procedure is defined by 3GPP, current open-source implementations do not support “Inter-MN Handover without SN change”. We leave its implementation to future work.

TABLE I: Completion Time Statistics over Heterogeneous NR-DC

Method	Mean (s)	Std (s)	Min (s)	Max (s)
<i>LAMP5G</i>	29.932	1.099	28.931	31.756
Dual-QUIC baseline	69.860	17.267	50.777	93.080

### C. Congestion Window for New Path after Handover

A key challenge during handover is setting the congestion window for a newly added path, whose initial window is much smaller than that of the existing path. This can cause the scheduler to favor the existing path and under-utilize the new one. To address this, *LAMP5G* adjusts the new path’s initial congestion window. Since the two paths may share much of the path beyond the source and target gNBs, and possibly the same bottleneck from the 5G core to the data network, we treat the existing path’s congestion window at path addition as the total window budget and split it evenly across the two paths. This allows the new path to ramp up faster while avoiding short-term under-utilization. Per-path congestion control then adapts to the actual bandwidth-delay product if the paths have different bottleneck capacities.

## VI. EVALUATION

### A. *LAMP5G* in Heterogeneous Cellular Environment

In this section, we experimentally evaluate *LAMP5G* under diverse network conditions to understand the impact of path heterogeneity on its behavior.

#### 1) Heterogeneous Cellular Environment with NR-DC

In this experiment, we evaluate *LAMP5G* in a simulated heterogeneous cellular environment. We deploy an NR-DC-based testbed on CloudLab using free5GC [27] as the core network and freeRANUE [28] for the simulated UE and RAN. As shown in Fig. 3a, the core, MN, SN, UE, and data network (DN) run on separate nodes, with the MPQUIC server on the DN and client on the UE. We emulate heterogeneous sub-6 GHz links using Linux tc: the MN–UE path has 300 Mbps bandwidth, 5 ms delay, and 0.04% loss, while the SN–UE path has 100 Mbps bandwidth, 8 ms delay, and 0.1% loss.

To understand the benefits of deploying a multipath transport protocol in this cellular network configuration, we compare *LAMP5G* with a dual-QUIC baseline that uses two independent single-path QUIC connections, one over the MN path and one over the SN path. In the experiment, both approaches are used to transfer a 1 GB file. For MPQUIC, *LAMP5G* performs path scheduling across the two paths. For the dual-QUIC baseline, which serves as a simple parallel-transfer comparison, the file is evenly divided into two halves, and each half is transmitted over one of the two QUIC connections.

Table I reports file-transfer completion times over five runs. MPQUIC with *LAMP5G* consistently outperforms the dual-QUIC baseline, reducing completion time by 57% on average. Fig. 4 shows a representative throughput comparison. In dual-QUIC, the faster MN-path transfer finishes around 26 s and its QUIC connection closes, leaving the remaining data to the slower SN path. In contrast, MPQUIC keeps both paths active until transfer completion, improving resource utilization and

aggregate throughput. Even on the MN path alone, MPQUIC with *LAMP5G* achieves 28% higher average throughput than dual-QUIC, demonstrating the benefit of coordinated scheduling in heterogeneous cellular networks.

### B. Flexible Handover with Multipath

To show that *LAMP5G* avoids data loss and the resulting congestion-control reaction during handover, we compare it with traditional inter-gNB handover using the OpenAirInterface (OAI) gNB implementation [29]. We run two experiments: a standard N2-based break-before-make handover from the source to target gNB, and *LAMP5G*’s seamless handover, which dynamically manages paths between the two gNBs.

To emulate N2-based handover under realistic 5G conditions, we deploy OAI gNB and UE implementations with RFSimulator on CloudLab. The setup includes one OAI core connected to two OAI gNBs, as shown in Fig. 3b, serving as the source and target gNBs. Because RFSimulator cannot adjust signal quality after initialization, we emulate mobility by modeling the UE moving from the source toward the target gNB. Signal strength is derived using the 3GPP TR 38.901 Urban Microcell (UMi) path-loss model [30], and handover is triggered by the A3 event. The gNBs are placed 400 m apart with 250 m coverage radii, creating a 100 m overlap region; the UE moves at 8.3 m/s.

For *LAMP5G*, we use the same setup but deploy two OAI UE simulator instances to emulate simultaneous connectivity to different gNBs, since the current OAI UE simulator does not support multiple gNB connections within one UE instance. Rather than performing traditional handover, *LAMP5G* dynamically adds and removes MPQUIC paths using the event-based strategy in Sec. V-A, based on the modeled signal strength. Both UE–gNB connections remain active during path updates, emulating seamless data continuity similar to NR-DC “inter-MN handover without SN change,” which is not yet supported by existing open-source implementations. This design focuses on seamless transport-layer handover.

In both experiments, the client runs at the UE and the server in the DN, with the UE downloading a 1.5 GB file. The baseline uses single-path QUIC with CUBIC. As shown in Fig. 5a, the UE initially downloads through the source gNB. When the A3 condition is met at around 43 s, handover reduces throughput to nearly zero for about one second, followed by roughly three seconds of recovery. Fig. 5c shows a congestion-window reduction to about 70% of its pre-handover level, reflecting CUBIC’s packet-loss recovery.

For seamless multipath handover, we use MPQUIC with *LAMP5G*. As shown in Fig. 5b, *LAMP5G* adds Path 2 through the target gNB at around 29 s, when the target signal exceeds the join threshold. Both paths then carry traffic concurrently. When the source signal falls below the drop threshold, Path 1 is gracefully removed at around 50 s, and transmission continues over Path 2 without interruption. Fig. 5c shows that, upon path addition, Path 1’s congestion window is redistributed across both paths. Since CUBIC growth depends on the historical max window, Path 1 initially grows faster than

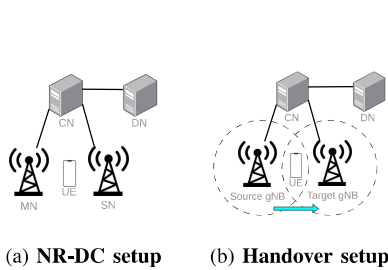
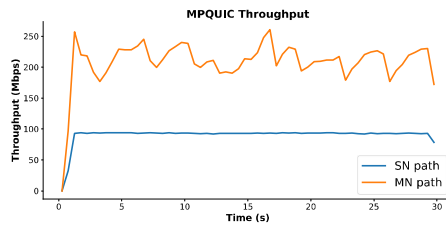
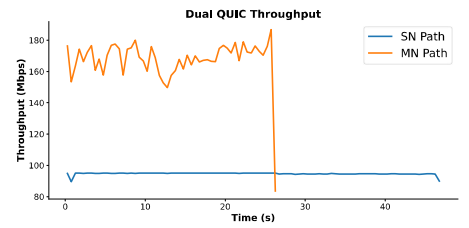


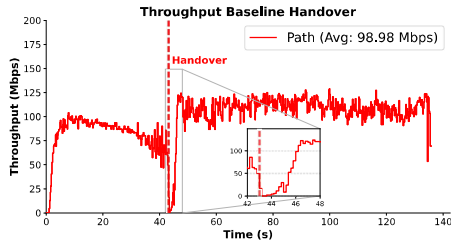
Fig. 3: Experiment setup.



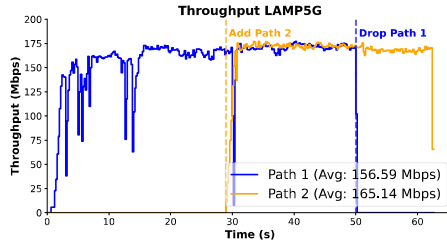
(a) MPQUIC w/ *LAMP5G*



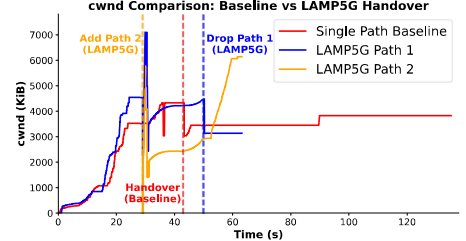
(b) Two concurrent single-path QUIC transfers



(a) Single-path QUIC throughput



(b) *LAMP5G* throughput



(c) Single-path QUIC vs *LAMP5G* cwnd

Fig. 5: Single-path vs. *LAMP5G* performance during handover.

Path 2. After Path 1 is removed, Path 2 continues to grow and sustain the transfer. Overall, seamless multipath handover improves average throughput by around 60% over the baseline.

## VII. CONCLUSION

In this paper, we present *LAMP5G*, a link-aware MPQUIC scheduler for heterogeneous and disruption-prone 5G cellular networks. *LAMP5G* uses bandwidth-aware path ranking with an AI-based throughput predictor to adapt proactively to environmental impairments, especially in high-frequency wireless networks. It also introduces seamless multipath handover to mitigate break-before-make disruptions. Evaluations show that *LAMP5G* reduces completion time by around 57% compared with a dual-QUIC connection baseline, demonstrating its effectiveness in dynamic 5G environments.

## VIII. ACKNOWLEDGEMENT

This work was supported by the generous US NSF grant CNS-2434009. The authors thank the reviewers for their helpful comments.

## REFERENCES

- [1] M. Ghoshal *et al.*, "Performance of cellular networks on the wheels," in *Proceedings of IMC'23*, 2023, pp. 678–695.
- [2] 3GPP, "Evolved universal terrestrial radio access (e-utra) and nr: multi-connectivity; overall description; stage-2," Tech. Rep. TS 37.340, 2025.
- [3] J. Drogo *et al.*, "INDIGO: Orchestration across multi-operator, multi-vendor, 5G networks during civilian disasters," in *Proceedings of the IEEE WF-PST*. IEEE, 2025.
- [4] 3GPP, "Network Sharing; Architecture and Functional Description," 3GPP, Technical Specification TS 23.251, Dec. 2023, defines Multi-Operator Core Network (MOCN) and GWCN architectures.
- [5] V. Jain *et al.*, "L25gc: a low latency 5g core network based on high-performance nfv platforms," in *Proceedings of SIGCOMM'22*, 2022.
- [6] Z. Liu *et al.*, "M2ho: Mitigating the adverse effects of 5g handovers on tcp," in *Proceedings of the 30th Mobicom*, 2024, pp. 1089–1103.
- [7] C. Paasch *et al.*, "Multipath TCP in the Linux kernel," <https://www.multipath-tcp.org>, 2013.
- [8] Q. De Coninck *et al.*, "Multipath QUIC: Design and evaluation," in *CoNEXT*, 2017, pp. 160–166.
- [9] Y. Liu *et al.*, "Managing multiple paths for a quic connection," Internet Engineering Task Force, Internet-Draft draft-ietf-quic-multipath-19, 2026, work in Progress.
- [10] Google, "SPDY," <https://www.chromium.org/spdy/spdy-whitepaper/>.
- [11] H. Yu *et al.*, "M&M: Machine-Learning Enhanced MPQUIC for Resilient Communication with mmWave Networks," in *MILCOM'24*.
- [12] A. Narayanan *et al.*, "A first look at commercial 5g performance on smartphones," in *Proceedings of The Web Conference 2020*, 2020.
- [13] Signals Research Group, "A Global Perspective of 5G Network Performance," Qualcomm Technologies, Inc., Technical Report, 2019.
- [14] K. Heimann *et al.*, "Experimental 5g mmwave beam tracking testbed for evaluation of vehicular communications," in *5GWF'19*. IEEE, 2019.
- [15] K. Larsson *et al.*, "High-speed beam tracking demonstrated using a 28 ghz 5g trial system," in *VTC-Fall'17*. IEEE, 2017, pp. 1–5.
- [16] H. Zhao *et al.*, "28 ghz millimeter wave cellular communication measurements for reflection and penetration loss in and around buildings in new york city," in *ICC'13*. IEEE, 2013.
- [17] K. W. Cho *et al.*, "mmWall: A steerable, transmissive metamaterial surface for nextg mmwave networks," in *NSDI'23*, 2023, pp. 1647–1665.
- [18] Z. Gao *et al.*, "Mambas: Maneuvering analog multi-user beamforming using an array of subarrays in mmwave networks," in *MOBICOM'24*.
- [19] S. K. Saha *et al.*, "Musher: An agile multipath-tcp scheduler for dual-band 802.11 ad/ac wireless lans," in *Mobicom'19*, 2019, pp. 1–16.
- [20] A. Zhang *et al.*, "Habitus: Boosting mobile immersive content delivery through full-body pose tracking and multipath networking," in *NSDI'24*.
- [21] A. Hassan *et al.*, "Vivisecting mobility management in 5g cellular networks," in *SIGCOMM'22*, 2022, pp. 86–100.
- [22] M. Prakash, "Resource allocation for heterogeneous wireless networks of devices with multi-connectivity," 2022.
- [23] M. G. Kibria *et al.*, "Next generation new radio small cell enhancement: Architectural options, functionality and performance aspects," *IEEE Wireless Communications*, vol. 25, no. 4, pp. 120–128, 2018.
- [24] V. Sharma *et al.*, "A transport protocol to exploit multipath diversity in wireless networks," *IEEE/ACM TON*, vol. 20, no. 4, 2012.
- [25] A. Narayanan *et al.*, "Lumos5g: Mapping and predicting commercial mmwave 5g throughput," in *IMC'20*, 2020, pp. 176–193.
- [26] S. Liu *et al.*, "Prophet: Prediction of 5g bandwidth using event-driven causal transformer," *Proceedings of POMACS'25*, 2025.
- [27] "free5GC: An open-source 5g core," Online, <https://free5gc.org/>.
- [28] "free-RAN-UE: A next-generation open-source 5g ran/ue simulator," Online; <https://free-ran-ue.github.io/>, 2025, accessed: 2026-02-23.
- [29] OpenAirInterface Software Alliance, "Openairinterface: Open-source 5g ran including gnb," 2026, <https://openairinterface.org/>.
- [30] 3GPP, "Study on Channel Model for Frequencies from 0.5 to 100 GHz," 3GPP, Technical Report TR 38.901, 2025.