IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 28 September 2023, accepted 11 October 2023, date of publication 18 October 2023, date of current version 31 October 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3325702

RESEARCH ARTICLE

Multipath Transport Analysis Over Cellular and LEO Access for Aerial Vehicles

AYGÜN BALTACI[1,2], KAUSHIK CHAVALI[2], MIKE KOSEK[2], NITINDER MOHAN[2], DOMINIC A. SCHUPKE[1], (Senior Member, IEEE), AND JÖRG OTT[2]

[1]Airbus Central Research and Technology, 82024 Taufkirchen, Germany
[2]Chair of Connected Mobility, Technical University of Munich, 85748 Garching, Germany

Corresponding author: Aygün Baltaci (ayguen.baltaci@airbus.com)

ABSTRACT Recent industrial advancements introduce novel safety-critical applications for commercial networks. Remote Piloting (RP) Aerial Vehicles (AVs) is an example application, where reliable wireless connectivity is key to ensure safe operations in the sky. Jointly utilizing cellular and satellite networks can enable robust Multipath (MP) communications; however, their usage must be orchestrated efficiently toward application requirements. In this work, we investigate the MP communications performance of cellular and Low-Earth-Orbit (LEO) satellite links with respect to the Quality-of-Service (QoS) requirements of RP operations. Using MP-Transmission Control Protocol (MPTCP) and MP-Datagram Congestion Control Protocol (MP-DCCP), we evaluate various transport layer configurations to efficiently orchestrate both links and to support the application requirements. For this purpose, we develop an end-to-end MP emulation testbed that can provide means to realistically emulate cellular and LEO links with MPTCP and MP-DCCP. We run bi-directional RP traffic over our testbed and measure the MP performance using different schedulers and Congestion Control (CC) algorithms. The results show that the flow size largely influences the individual path utilization due to high LEO link-layer losses. Moreover, excessive retransmissions occur on the MPTCP layer due to Head-of-Line (HoL) blocking from asymmetric link conditions. Using MP-DCCP without retransmissions helps avoid late arrivals and can meet the 99.999% communication reliability demand.

INDEX TERMS Multipath emulator, multipath communications, cellular communications, satellite communications, LTE, 5G, 6G, LEO, GEO, non-terrestrial networks, aerial networks, eVTOL, UAV.

I. INTRODUCTION

Safety-critical aerial applications demand seamless and reliable wireless connectivity to ensure safe operations in the sky. Remote Piloting (RP) is a significant use-case for future aerial applications, where a remote pilot located on ground operates an Aerial Vehicle (AV) by relying on wireless communications. Although wireless ecosystem offers state-of-the-art technologies to provide seamless connectivity, the unpredictable Radio Frequency (RF) nature of wireless communications poses challenges toward meeting the stringent Quality of Service (QoS) requirements of safety-critical applications. In this regard, RP operations demand ultra-reliable connectivity, up to 99.999%, and in addition, low-latency video delivery is essential for remote pilot to timely sense the flight environment and maneuver the vehicle promptly [1].

The conventional wireless communication technologies are not designed to provide such high reliability [1] and hinging on a single network for a safety-critical application poses shared risks from various aspects, such as single point of failure. However, Multipath (MP) connectivity can be a promising method to increase communication reliability for such use cases [2], [3], [4] using diversity from different aspects such as link technology, network and RF channel. Although it has indications toward increased Size, Weight and Power (SWaP) requirements as well as potential data overhead, MP communications can introduce a number of benefits for safety-critical use cases. Compared to single-path communications, the usage of multiple wireless links can aggregate the overall throughput, improve communication reliability and reduce the end-to-end latency. In addition,

The associate editor coordinating the review of this manuscript and approving it for publication was Shu Xiao.

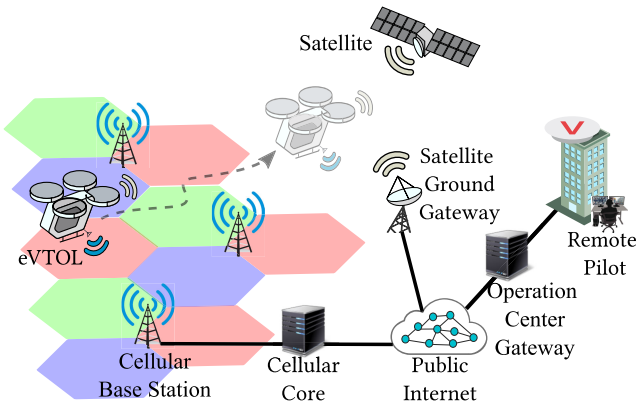

FIGURE 1. The considered multipath communications scenario in this study. An AV is equipped with cellular and LEO satellite terminals; hence, it can utilize both networks simultaneously. A pilot is located on ground and operates the vehicle remotely. While AV sends a live video stream to the pilot, it receives control commands via wireless links.

having network diversity is significant to achieve disjoint end-to-end paths, which can help avoid single point of failures. However, the underlying wireless links must be orchestrated effectively to avoid artifacts from link heterogeneity and to efficiently meet the QoS demands.

The joint use of cellular and satellite links has particular advantages and the potential to enable future safety-critical use cases in the sky [5]. It can also be considered for different mobility applications such as railway communications, where MP schemes are also considered [6]. While ground-based cellular networks can provide high-throughput connectivity in urban areas, satellite connectivity can complement cellular networks to avoid coverage holes and to achieve ultra-reliable connectivity. The latest 3rd Generation Partnership Project (3GPP) standards include the Non-Terrestrial Networks (NTNs) in the study items [7], [8], and emergent Low-earth Orbit (LEO) constellations, such as Starlink [9] and OneWeb [10], provide high-throughput services over LEO links. Although various studies evaluate the potential of cellular networks toward next-generation aerial use cases [3], [4], [11], [12], a research gap exists in understanding the joint performance of cellular and satellite networks in a MP fashion to evaluate their suitability toward safety-critical aerial applications. In addition, suitable MP transport algorithms should be explored to efficiently orchestrate cellular and LEO links. Furthermore, investigations of the feasibility of satellite links for Unmanned Aerial Vehicles (UAVs) are currently scarce, since conducting hardware-based measurements with satellite terminals are constrained by the limited SWaP capabilities of the UAVs. Hence, emulation-based studies are necessary to be able to evaluate the joint performance of cellular and satellite communications for aerial applications.

In this work, we study the MP communications performance of cellular and LEO links when they are jointly orchestrated for the RP operations of AVs. For this purpose, we develop an end-to-end MP testbed that emulates cellular and LEO links based on real-world traces, thereby representing an end-to-end connectivity. We run representative bi-directional RP traffic that comprises AV control and video data exchange. Employing Multipath Transmission Control Protocol (MPTCP) and Multipath Datagram Congestion Control Protocol (MP-DCCP) as the transport protocols, we evaluate and compare the suitability of reliability-centric versus best-effort transmission toward the QoS requirements of RP use case. Our aim is to find out whether the joint use of cellular and LEO links can meet the RP application requirements, and determine suitable MP transport configurations that can efficiently orchestrate these heterogeneous links. Therefore, we study different MP schedulers and Congestion Control (CC) algorithms along with the influence of transport-layer retransmissions to dissect their individual effect on achievable MP networking and application-layer performance. We characterize the cellular link based on AV-based Long Term Evolution (LTE) measurements [13], and the satellite link is based on our measurements with Starlink since LEO constellations are more suitable for low-latency communications rather than other satellite orbits at farther distances to the Earth.

We aim at answering the following questions throughout this study:

1) What is the achievable MP networking performance over cellular and LEO networks toward the QoS requirements of the RP operations? Can their joint use meet the QoS requirements of the video and control traffic?
2) What are the most suitable MP transport configurations to efficiently orchestrate cellular and LEO links? Is a reliability-centric protocol like MPTCP or a User Datagram Protocol (UDP)-like best-effort transmission more favorable for the RP scenario?
3) What are the MP performance bottlenecks and how can they be avoided in an optimal MP transport protocol?

The results show that high LEO link-layer losses yields aggressive CC decisions, and the LEO link stays underutilized for large flows (10 Mbps video traffic) with MPTCP. Nevertheless, High Definition (HD)-resolution video transfer with playback latency <300 ms can be achieved up to 90% of the time even with MPTCP. Using MP-DCCP without retransmission helps in avoding late arrivals, and 99.999% communication reliability can be achieved when both paths are utilized in a redundant manner.

Our work contributes to literature not only with an in-depth evaluation of cellular and LEO links in a MP fashion toward a safety-critical use case, but also with an insightful comparison between the capabilities of MPTCP and MP-DCCP to efficiently handle the heterogeneity between both links. Moreover, we introduce a novel MP testbed that can realistically emulate cellular and LEO links with MP transport protocols to facilitate future research.

We organize the rest of the paper as follows: We discuss the background of RP operations and related work in Section II. Next, we detail our MP communications testbed platform, cellular link and LEO link setup along with MP transport

layer configurations in Section III. We analyze the emulation results in terms of single-link performance, MP networking and application perforance for our scenario in Section IV. After discussing our findings and limitations in Section V, Section VI concludes the study with key takeaways.

II. RELATED WORK

We begin this section by describing the RP scenario and its QoS requirements. We also provide background regarding relevant MP transport protocols, schedulers and CCs. Afterwards, we summarize the previous works in regards to MP communications studies for UAVs as well as MP networking testbeds.

A. BACKGROUND

1) THE RP SCENARIO

RP operations essentially consists of an AV, a remote pilot, and a set of wireless links to control the AV as well as the wired backhauls [14], as shown in Figure 1. In this work, the term *AV* comprises the electrified aerial platforms that operate at low-altitude airspace below 1 km altitude [1] such as electric Vertical Take-off and Landings (eVTOLs) and UAVs. Although this scenario also includes an Unmanned Traffic Management (UTM) system to coordinate the airspace, its connectivity requirements are outside the scope of this study. The remote pilot controls the AV from a ground operation center, which connects to the AV via cellular and satellite access networks to send control traffic and receive video stream through public Internet. We consider cellular and LEO links together since their high-throughput, low-latency link characteristics are promising toward the QoS requirements of RP [1]. They also bring link, network and technology diversity to the scenario. In the rest of the paper, *uplink* channel refers to the traffic from an AV to a remote pilot, and vice versa for the *downlink* channel.

2) THE QoS REQUIREMENTS

Data rate requirements of the video stream can vary between 10 and 100 Mbps depending on the number of cameras and minimum video quality. Control traffic demands low bitrates between 0.25 Mbps and 1 Mbps [1]. Furthermore, low-latency communications is essential for both streams to ensure safe operations and thus, an upper bound of 300 ms end-to-end latency is required [15]. Lastly, operating an AV from ground is a safety-critical application in nature, and reliable connectivity is one of the major components to ensure safe operations in the sky. Therefore, a communication reliability of 10^{-5} is required [1]. A single wireless link cannot provide such high reliability, mainly due to the unpredictable RF propagation. It is also beyond the level of reliability that wireless communication standards promote for outdoor applications [1]. Hence, this requirement is one of the fundamental reasons for studying MP communications in this scenario. Latency and communication reliability requirements are applicable to both downlink and uplink channels.

3) MULTIPATH TRANSPORT

MPTCP is the multipath extension of TCP, which is a full-duplex and connection-oriented transport standard with features such as packet loss recovery, flow control, and in-order packet delivery [16]. TCP does not assume reliability on the lower layers in order to support any reliable and unreliable connectivity services on the internet. MPTCP extends the TCP architecture by enabling the simultaneous use of multiple end-to-end paths for different QoS targets such as improved throughput [17]. It creates multiple subflows across the available paths and selects suitable paths for each data packet mainly based on the scheduling and CC decisions [17].

MP-DCCP provides unreliable transport scheme and is based on the Datagram Congestion Control Protocol (DCCP) standard [18]. Unlike MPTCP, it is a connectionless service (e.g., UDP) and can be rather suitable for services that demand low-latency connectivity. It also has scheduling and CC algorithms to handle the underlying paths. However, it does not have error recovery, and in-order delivery is only optional [18].

Congestion Controls

Congestion Controls (CCs) have an essential role in MP transport protocols to use the available paths efficiently avoiding congestion bottlenecks. They are designed to improve throughput of a MP flow, ensure fairness among flows and balance the congestion among all the available paths [19]. While uncoupled CCs treat each underlying subflow independently, coupled CCs can dynamically adapt their overall agresiveness by considering all the subflows in a joint manner. The following CCs are relevant to our study:

1) **BAlanced LInked Adaptation (BALIA)** aims to balance the available network resources among the flows and is responsive to network changes [20]. It increases the Congestion Window (CWND) with a complex function that is based on the number of sent bytes over an Round Trip Time (RTT) period after each received acknowledgement packet. CWND reduction function is dynamically determined based on the aggresiveness factor. in the order of after every packet loss.
2) **Bottleneck Bandwidth and Round-trip Propagation Time (BBR)** is designed to keep the bottleneck saturated without creating a congestion. It measures the delivery rate and propagation delay to estimate the path capacity and RTT, and it applies pacing to control the sending rate [21]. It halves the CWND after every packet loss.
3) **Cubic** is a loss-based, uncoupled CC that employs a cubic function to quickly increase CWND after loss events. In case of a packet loss, it applies a multiplicative decrease to the CWND [22].

4) **NewReno** is also an uncoupled CC and it takes packet losses as congestion signal. it essentialy works by multiplicatively decreasing and additively increasing the CWND during packet loss and recovery events, respectively [23].
5) **Opportunistic Linked-Increase Algorithm (OLIA)** is essentially similar to BALIA however, it is not as responsive as BALIA to the changes in network [20]. OLIA increases the CWND based on the quality of paths, which is determined by the number of transmitted bytes as well as link RTT [24]. It also halves the CWND after every packet loss.
6) **Weighted Vegas (wVegas)** is a delay-based CC algorithm and it estimates the link queuing delay to configure the CWND size. It monitors the RTT fluctuations as well as the change in MP aggresiveness factor to adjust the CWND accordingly [23].

Schedulers also have a primary role in MP transport to allocate the data segments across the underlying subflows while taking the available CWND of each subflow into account. Schedulers are triggered either when they receive data from an application or when an acknowledgement packet frees up the CWND of a subflow. Below are the scheduling algorithms we used in our study:

1) **BLocking ESTimation (BLEST)** scheduler is designed to avoid Head-of-Line (HoL) blocking by employing CWND and link RTT information to proactively perform scheduling decisions. It estimates the sending window occupation time of each segment to decide on which subflow to schedule the segment for transmission [25].
2) **Cheapest Path First (CPF)** scheduler works by assigning a cost value for each underlying link and uses the path with the lowest cost as long as the CWND of the link with the lowest is available.
3) **Lowest RTT (LowRTT)** scheduler basically works by measuring the RTT of each link and selecting the link with lowest RTT as long as the CWND of that link is available [26].
4) **Redundant** scheduler duplicates the application data over all the available links and sends them in a best-effort manner. This helps quickly recover from packet losses and minimize latency at the expense of data overhead.
5) **Round Robin** scheduler aims to treat each path fairly and uses each available one after the other as long as their CWND are available [26].

In the next section, we provide insight regarding the related studies from literature and the novel contributions of our work.

B. LITERATURE REVIEW

Space-air-ground networks recently got attention in multiple domains and are studied for diverse research subjects. While a number of works envision a fully-integrated space-air-ground architecture for terrestrial use cases [27], [28], [29], internet service providers already investigate the potential of emergent satellite services to complement their terrestrial networks and submarine cable infrastructure [30], [31], [32], [33]. Hence, the majority of the studies toward space-air-ground networks focus on the *operation* or *management* aspects of these integrated networks. Whereas in our work, we consider the joint use of space (LEO) and ground (cellular) segments from an *end-user* perspective, where the aerial segment is a user device rather than a network provider. We take the existing space and terrestrial networks as is and aim to utilize them in a MP fashion in order to meet the connectivity demands of a safety-critical aerial use case.

A number of studies analyze the potential of cellular networks for AVs and a few papers evaluate the performance gain from cellular links in a MP fashion. Several works also consider MPTCP in aerial scenarios with different wireless link combinations rather than cellular and LEO links together. The scope of our work differs from those studies since we analyze cellular and LEO links over MPTCP and MP-DCCP toward the stringent connectivity requirements of the RP operations. As MP-DCCP is recently developed, we did not find relevant work toward aerial applications. Lastly, we also did not notice any MP simulation or emulation testbed in literature that can emulate cellular and satellite links simultaneously and orchestrate them via an MP transport protocol.

Table 1 presents a brief summary of the scope of the related studies compared with our work, and we provide further details in the following. Utilizing multiple links redundantly is particularly effective in improving the communication reliability in drone applications. In this regard, the studies [3], [4], [11] propose to use multiple cellular links to achieve this goal. Performing real-life measurements with a drone, the authors of [4] highlight that utilizing dual connectivity via different Mobile Network Operators (MNOs) can achieve 99.9% communication reliability compared with 97.6% reliability of a single-link connectivity. In addition, the work, [36], proposed UAV-based MP video streaming using dual cellular links for forest fire surveillance operations. They proposed a system that comprises a Raspberry Pi module, a GPS antenna and two cellular modems and they distribute the video data over each cellular link. Their field tests show that up to 10 Mbps video stream could be delivered to the ground with sufficient video quality. In [3], the authors take a different approach and combine a public and a private cellular links for UAV maritime and rescue applications. They orchestrate the links with MPTCP to achieve high reliability. After setting up a hardware-in-the-loop experiment, they utilize the MPTCP's Lowest RTT (LowRTT) and redundant schedulers. The results show that while the LowRTT scheduler helps improve the communication range and the data rate, redundant scheduler minimizes the RTT. Overall, although these studies consider different multilink approaches, none of them includes a satellite link to evaluate the potential gain from link and technology diversity.

TABLE 1. Comparison between the scope of related works and our study.

Article	MP Scheme	Studied Wireless Links	Research Goal
[3]	MPTCP	Dual Cellular	Improve communication availability and robustness for in-car connectivity.
[4]	Redundant transmission	Dual Cellular	Improve communication reliability and latency for RP operations.
[11]	Redundant transmission	Dual Cellular	Improve communication reliability for RP operations.
[12]	NECTOR	Dual Cellular & GEO Satellite	Improve communication resiliency for UAV connectivity.
[34]	MPTCP	Cellular and Satellite	Evaluate the suitability of BBR CC over heterogeneous paths.
[35]	MPTCP	Cellular and WiFi	Evaluate the trade-off between coupled versus uncoupled CCs for high-speed railway networks.
Our work	MPTCP and MP-DCCP	Cellular and LEO Satellite	Evaluate suitable MP transport configurations to orchestrate heterogeneous links and compare the trade-off between reliable and unreliable transmission for RP operations.

The study, [12], presents a novel protocol, `NECTOR`, which is based on network coding using UDP with two LTE links and one satellite link. The receiver controls the packet reception rate with a torrent-based methodology, and they improve communication reliability by employing network coding. However, their links are emulated with basic parameters, not reflecting the representative behaviors of cellular and satellite links. Nevertheless, compared with MPTCP, `NECTOR` reduces the required datagram size by 11.2% to recover packets and achieves at least 5-6% higher average throughput than MPTCP. Although this protocol operates on the application layer, it can be an alternative to MP transport protocols to achieve reliable MP communication for RP operations and can be considered in a future work for aerial applications.

As for MP testbed platforms, the authors of [37] develop an integrated AV and network simulator called `FlyNetSim` by combining the `Ardupilot` and `ns-3` simulation platforms. Similar to our work, they support multiple wireless links with different technologies: LTE, Wireless Fidelity (WiFi) and Device-to-device (D2D); however the networking capabilities of `FlyNetSim` depends on the features provided in the `ns-3` simulator and it does not simulate satellite links. In addition, `FlyNetSim` does not model MP transport layer. Another work, [38], provides a wireless network emulator, `Collosseum`, which can emulate cellular and WiFi links as well as RF channels using up to 128 Software Defined Radios (SDRs). Although the paper emphasizes its capabilities on emulating RF signals on the physical layer, it does not mention about the capabilities on higher layers. The work in [39] is a multi-node testbed platform and does provide actual cellular connectivity over multiple MNOs. Their default use case is to provide robust connectivity to different voting locations in Norway. Hence, they don't use multiple links at the same time: The secondary link is triggered in case the primary link fails. Compared to our platform, they don't have satellite link emulation capabilities. In addition, their end devices are at fixed locations, thus the testbed does not include channel fluctuations due to mobility. Lastly, the work in [40] develops a multilink simulation platform with flight physics capabilities. Even though the simulator enables the use of cellular and satellite links for different aerial platforms, its simulation capabilities are constrained up to the MAC layer.

The authors of [34] consider the MPTCP performance over paths with large latency differences, such as cellular and GEO-stationary satellite links. Setting up an emulation environment over Mininet emulator with path latencies between 10 ms and 1000 ms, they found out that the Bottleneck Bandwidth and Round-trip Propagation Time (BBR) CC is able to maximize the achievable throughput over heterogeneous links. Relatedly, the authors of [35] also consider various CCs in their work and analyze the MPTCP performance for high-speed railway use cases over cellular and WiFi links. Comparing with the Linked Increases Algorithm (LIA) CC, a simple uncoupled CC that treats each subflow as independent TCP connection outperforms in achieved throughput, CWND and RTT performance. Lastly, the study, [41], developed an MPTCP path selection algorithm for UAV swarms based on a matching algorithm between UAV data traffic types and the underlying wireless networks. After defining *fitness scores* based on the criticality of the data as well as the capabilities of the wireless links, the scheduler matches the fitness scores of data services and the links. In a simulation study, they showed that their algorithm could achieve stable throughput and 20% throughput gain compared to Opportunistic Linked-Increase Algorithm (OLIA).

All in all, our study contributes the following novel aspects compared to the previous work: 1. We study cellular and LEO networks in a MP fashion toward the connectivity demands of future RP operations, 2. We include MPTCP and MP-DCCP in the same study to compare the trade-off between reliable and unreliable transmission on MP-level for the RP data traffic, 3. We investigate MP transport configurations and evaluate different CCs and schedulers to find out suitable settings for the scenario, and 4. We develop a heterogeneous MP testbed that can emulate cellular and satellite links with MPTCP and MP-DCCP as transport protocols. The next section describes our MP emulation testbed setup as well as the requirements of the RP use case.

III. MULTIPATH EMULATOR SETUP AND THE RP SCENARIO

In this section, we first describe our experiment design methodology and give an overview of our MP emulation environment. We describe the architecture of the testbed and explain how we perform the cellular and satellite link emulations. Following, we characterize the cellular and satellite links, the emulated data traffic and the MP transport layer configurations evaluated in our measurements.

In the measurement setup, our overarching goal is to create a representative RP scenario that is operated over cellular and LEO connectivity in a MP fashion. For this purpose, we develop a MP testbed that has cellular and LEO link emulation capabilities based on real-life traces. We create a realistic bi-directional RP traffic by generating control commands on downlink and HD-resolution AV video traffic on uplink [15]. After setting up MPTCP and MP-DCCP as transport protocols, we run a set of measurements with different scheduler and CC combinations with the aim to analyze their trade-off and find the most optimal transport configurations to meet the QoS demands of the application.

A. MULTIPATH TESTBED SETUP

Figure 2 shows our MP emulation testbed architecture. It is composed of the following modules: 1. A client (left-most) and a server (right-most) that represent the end-users, 2. A cellular access emulator based on `MoonGen` (pink) [42], and 3. A satellite access emulator based on `OpenSAND` (blue) [43], [44]. Both link emulators are previously evaluated in literature [42], [45] and employed in other works to emulate cellular and satellite links realistically [46], [47], [48]. The client and server represent an AV and a RP, respectively, and Both are MPTCP- and MP-DCCP-capable, which allow them to orchestrate bi-directional traffic over cellular and satellite paths. They are connected to the link emulators via proxy gateways. We set up `netem` instances [49] before the proxy gateways to emulate link delays, e.g., from wired backhauls. The client has multi-homing capability with access to the interfaces of both links. As the server does not support multi-homing, it relies on a multi-homed gateway that connects it to the cellular and satellite networks.

The *Cellular emulator* is based on [42], which emulates an LTE link using the MoonGen traffic generator [50] augmented with a latency queue data structure. It can create bi-directional traffic with four processing threads (two for each direction). Frames in the latency queue are forwarded to a transmission ring, which are then scheduled for packet transmission in First In First Out (FIFO) order [42]. Furthermore, the cellular emulator supports heterogeneous uplink and downlink rates, latency, packet losses, network buffer, concealed loss recovery at the link layer utilizing the Hybrid Automatic Repeat Request (HARQ) and and power-saving features such as the Discontinuous Reception (DRX). We used an LTE emulator rather than 5G since the real-life traces we collected with a drone are with the LTE technology due to unpredictable and insufficient 5G coverage in the air. Nevertheless, the emulator can be extended to 5G by modeling the appropriate link parameters.

The Satellite emulator is based on `OpenSAND` [43] and emulates an end-to-end satellite communication system. It creates bi-directional wireless links between a satellite and a gateway, as well as between a User Equipment (UE) and a satellite. OpenSAND works by emulating the Digital Video Broadcasting - Return Channel via Satellite (DVB-RCS2) wireless protocol stack along with IPv4/v6 and ethernet connectivity [43]. It also takes propagation delay and link attenuation into account. It can create star and mesh topologies, and also supports multiple gateways as well as multiple spot beams [43]. Our satellite link emulation topology consists of four components: 1. *Satellite terminal* that represents a UE, 2. A *satellite* in orbit, 3. *Satellite GW*, which is a satellite gateway operated at a ground station of the respective *satellite* component, and 4. A *vSwitch*, which creates a bidirectional connectivity between each emulated satellite component.

1) HARDWARE AND SOFTWARE SETUP

Our MP testbed operates on a single workstation, powered by a 24-core Xeon CPU at 3.4 GHz. The workstation has 128 GB RAM to ensure enough memory for running parallel applications. These capabilities exceed the hardware setup of the original implementations of individual link emulators [42], [44]. During emulation runs, we observed an average of 12% CPU utilization with $<$1 GB of RAM usage. We use a total of 4 Network Interface Cards (NICs): 2 x Intel I350 1G and 2 x Intel 82599 10 G, since each downlink and uplink direction in the cellular emulator requires two NICs. Only NICs with hardware timestamping feature are compatible with this setup due to the Data Plane Development Kit (DPDK) environment, where the cellular emulator runs. The DPDK framework offers a polled-mode operation to ensure advanced control over packet timings [51].

We set up the emulator using Ubuntu 18.04 release. We employ the MPTCP v0.95.3 and MP-DCCP v0.3 kernels. While MP-DCCP v0.3 is latest version at the time of conducting this study, we use the MPTCP v0.95.3 since it is deployed with a variety of schedulers and CC options, and hence, it is more suitable for research works. Whereas, the latest MPTCP v1 is still a work-in-progress and currently only supports the LowRTT scheduler without coupled CCs [52], [53]. Furthermore, we develop a containerized environment by creating 10 Linux network namespaces each with its function for executing different modules of our emulator. We also use the MoonGen [50] packet processing tools since the cellular emulator relies on it for packet generation [42].

2) EXTENSIONS TO ACCESS EMULATORS

We extended the cellular emulator by adding Handover (HO) capabilities. We implemented the HOs with a delay function during packet processing in MoonGen. Once a HO event

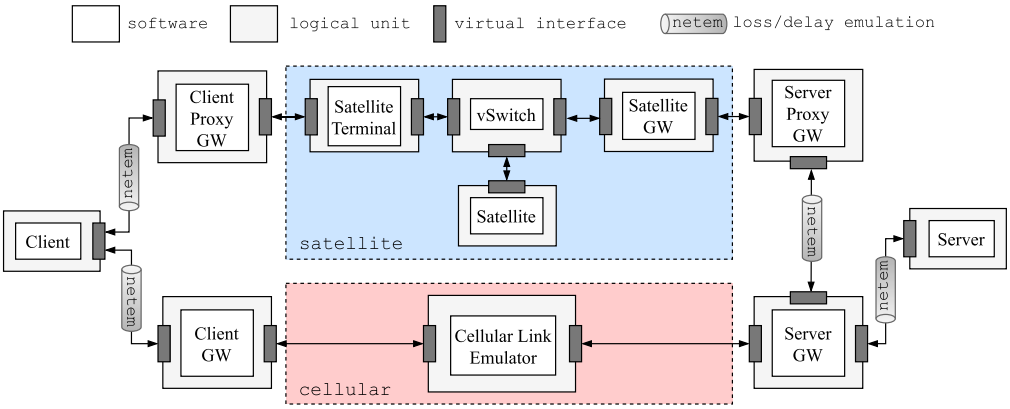

FIGURE 2. Architecture of the multipath emulation testbed developed for this study. The testbed is operated on a single workstation via physical and virtual interfaces and support bidirectional multipath traffic using MPTCP and MP-DCCP between client and server.

TABLE 2. Main differences between MPTCP and MP-DCCP protocols [18].

Feature	MPTCP	MP-DCCP
Connection-oriented	✓	✓
Reliable data transfer	✓	x
Congestion control	✓	✓
Flow control	✓	x
Packet loss handling	retransmission	report only
Reordering	✓	optional

is triggered, the emulator halts packet processing for the given duration and upon HO completion, it forwards the accumulated frames according to the configured rate. We also introduced time-varied link emulation capabilities to reflect the RF propagation effects on the link behavior.

As for the satellite emulator [44], we extended it by introducing time-varied link characteristics to reflect the link dynamics for end-to-end latency.

3) MULTIPATH CAPABILITIES OF THE EMULATOR

We set up the MPTCP and MP-DCCP as transport protocols in the testbed. We selected MPTCP since the protocol is mature, widely deployed [54] and its underlying robustness is favorable for our safety-critical RP scenario. However, MPTCP relies on TCP's reliability mechanism, which can cause performance artifacts such as HoL blocking [55] that can introduce bottlenecks toward ensuring low-latency delivery – especially for the video traffic. Therefore, we also utilized the MP-DCCP protocol to experiment unreliable transmission in comparison to MPTCP. As MP-DCCP is based on UDP protocol [56], it can be rather suitable for ensuring low-latency video traffic. Table 2 describes the main differences between MPTCP and MP-DCCP. Transport-layer retransmissions are one of the main differentiators between these protocols. The flow control mechanism in MPTCP influences the packet sending rate based on the available receive window [17].

4) SCALABILITY

The testbed supports multiple flows as well as background traffic during measurements. Using multiple cellular links are constrained by the hardware capabilities since each bi-directional cellular link demands 4 NICs. The satellite emulator, OpenSAND, supports up to 5 satellites and 2 ground gateways on a single instance [43]. Multiple UAVs can also be created by ensuring dedicated namespaces for each of them. In addition, a set of cellular and satellite links must be created for each UAV. Although direct emulation of a UAV flight is not possible, UAV motion can be represented by modeling the influence of UAV movement on the link performance such as data rate, latency and packet losses. Therefore, we introduce time-varying link properties to both cellular and satellite emulators. In addition, other transport protocols such as Multipath QUIC (MP-QUIC) can be enabled by installing the application on user space or newer versions of MPTCP can be utilized by installing it in the kernel space.

We publish the source code of our MP testbed along with the measurement tools, and our extensions to the cellular and satellite emulators in [57]. We also provide the necessary recipes to clone our testbed to other platforms.

B. RP SCENARIO SETUP IN THE TESTBED

In this section, we detail the way we utilize our MP testbed to create a representative end-to-end RP scenario. Figure 3 illustrates the overall system model we created in the testbed. Subsequently, we elaborate on the building elements of the scenario in detail.

1) COLLECTION OF THE REAL-LIFE MEASUREMENTS

We collected the LTE measurement in our drone measurements with public network operators in an urban environment [13]. The environment was surrounded with tall buildings and hence, the RF channel was affected by multipath and shadowing effects.

The measurements covered up to 120 m height and consisted of a total of 90 flights. We collected pcap traces on both ends to compute data rate, latency and packet losses. We used the `QCSuper` [58] to collect LTE-layer information in order to detect HO events. These measurements helped us identify the differences of cellular network performance in the air compared with the ground. Hence, we used our collected traces to reflect the influence of drone flight on the cellular link.

We also used real-life measurements to model the satellite link. Our goal is to mimic the behavior of a LEO link as close as possible in an emulation environment. We measured the network performance of a standard Starlink and run measurements for 48 hours using a standard Starlink terminal. The client Dish was located in Garching, Germany and the server in an Amazon Web Service (AWS) data center in Frankfurt, Germany. As the dish was located on a rooftop, the RF channel was mainly free-space along with minor reflections from the roof. Similar to the LTE measurements, we characterized the data rate, latency and packet losses using the collected pcap traces. We plan to make the Starlink dataset public upon acceptance.

We validated the link behaviors of cellular and satellite emulators against real-life traces before running our measurements. We included more details regarding our validation work in Appendix A.

2) EMULATOR LINK SETUP

We configure the cellular and LEO link parameters such as data rate, end-to-end latency, Packet Error Rate (PER), LTE HO frequency and Handover Execution Time (HET) based on our experimental measurements. We modeled the cellular HOs based on the frequency and the duration of HOs from collected traces. Modeling these parameters in the emulator are significant to reflect the differences of cellular network performance in the air compared with the ground [13]. As we set up the cellular and LEO links based on our real-life measurements, the emulated links already include the effects from AV mobility, background traffic as well as the other influences from core networks and the public Internet. We included the link dynamics in terms of link capacity and end-to-end latency, especially on uplink to account for time-varying link conditions for the high-rate (10 Mbps) video traffic, whereas the dynamics are less relevant for the low-rate (1 Mbps) control traffic. We modeled constant capacity on the LEO link due to a software limitation with the modulation scheme in `OpenSAND` v5, which has been discussed in recent releases [59]. As the LEO link capacity is abundant (62/18 Mbps on downlink/uplink) compared with the data traffic volume (1/10 Mbps on downlink/uplink), the influence of capacity fluctuations can be abstracted.

3) CHARACTERIZATION OF THE WIRELESS LINKS

Figure 4 shows the modeled data rate and latency on LTE and LEO links based on our experimental measurements. From our drone flight dataset, the average data rate on LTE link fluctuates between ≈15 and 45 Mbps on both directions. As for end-to-end latency, LTE uplink is on average ≈53 ms and observes some latency spikes, which goes as high as 2900 ms. These spikes are correlated with the increased HET outliers while the drone flies in the air [13]. On downlink, we estimate the mean latency to be 45 ms during measurements. We measure the mean HO duration as 20.01 ms with a standard deviation of 195.13 ms. Mean HO frequency is 0.05 Hz with a variance of 0.042 Hz. Appendix B provides more details regarding our HO model in the air. Lastly, we measure the average PER as 0.006% on LTE in both directions.

As for the LEO link, mean capacity are ≈62 and 18 Mbps on downlink and uplink, respectively, based on our Starlink measurements. Latency on both channels are fairly symmetric and fluctuate between 12 and 38 ms, as illustrated in Figure 4. Such low end-to-end latencies are achieved since our client and server are located in nearby regions in Germany during our measureements. This is a representative scenario for RP operations since such operations consider short-range (<80 km distance) regional flights [1, Sec. II-B]. Finally, we measure the average PER on the LEO link as 0.17%. It is two orders of magnitude higher than that of the LTE and can be related to satellite HOs. This finding is aligned with the results shown by a recent study [60].

Modeled LTE and LEO links are asymmetric in terms of not only link capacity, but also communication reliability. LEO link-layer losses significantly influence the MP orchestration performance and create a heterogeneous MP conditions for the transport protocols, as we analyze in detail in the next section.

4) THE RP DATA TRAFFIC GENERATION

The RP traffic comprises commands to control the drone on downlink (from remote pilot to drone) and a high-quality video traffic on uplink (from drone to remote pilot). While control commands are low- and constant-rate data traffic, video stream is high-rate and bursty [15]. We simultaneously send bidirectional traffic in the testbed to realistically evaluate the scenario, and to find out whether individual data traffic influence one another due to, for instance, sharing the same network resources. Based on the data rate requirements in Section II-A, we set the control and video traffic rates to 1 Mbps and 10 Mbps, respectively. While 1 Mbps represents an upper bound for control commands, 10 Mbps is sufficient to stream a HD-quality video [61]. We use `iPerf2` to generate the control traffic, and we utilize a `Gstreamer`-based streaming appliction to generate Real-time Transport Protocol (RTP) video traffic with MPTCP. We used `iPerf2` instead of `iPerf3` since `iPerf2` provided more stable throughput pattern. In MP-DCCP, we utilize a modified `iPerf3` tool that is particularly adapted for MP-DCCP [62]. As there is no video application available for MP-DCCP at the time of conducting this study, we generate a representative

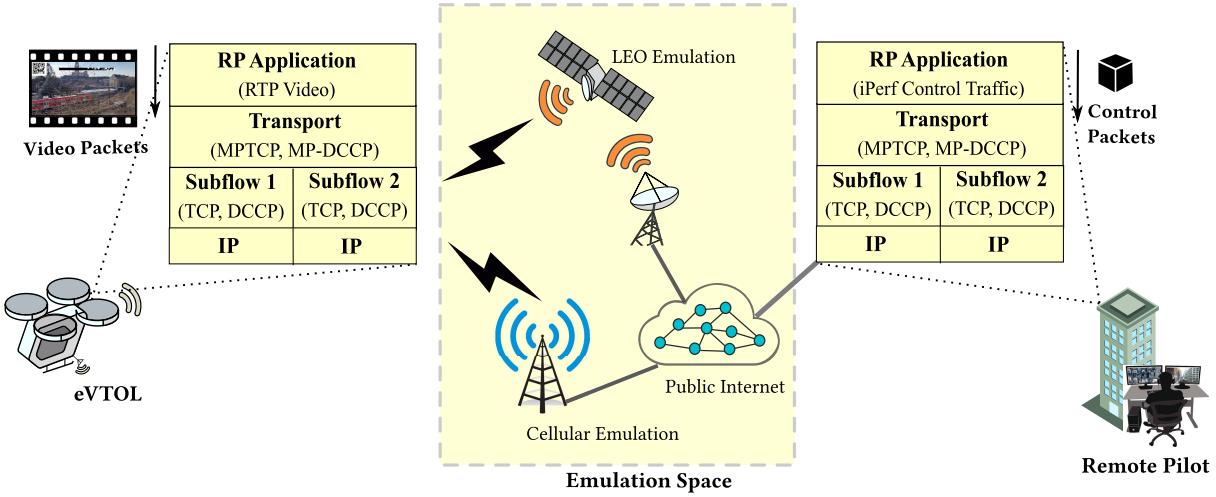

FIGURE 3. Overall system model of the RP scenario in our MP testbed. While AV and remote pilot create the video and control data traffic at each end, the MP transport schedules the generated data packets to the underlying wireless links. The emulated cellular and LEO links deliver the traffic to the end users.

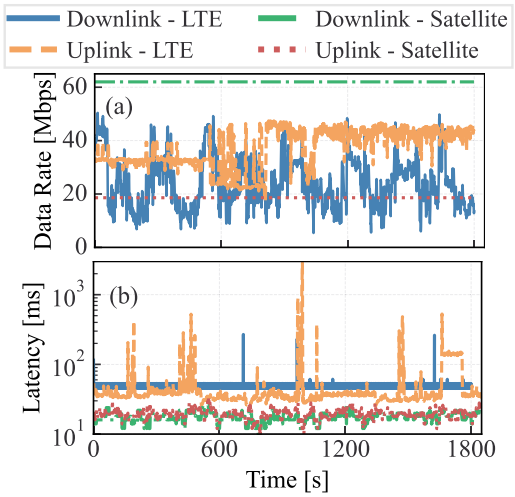

FIGURE 4. Emulated data rate and one-way latency on the LTE and LEO links in the testbed. The LTE link configuration is based on the collected traces from previous drone measurements in the air. We model the LEO link with real life traces we collect from a standard Starlink dish in static conditions on ground.

10 Mbps flow with `iPerf3` for the video traffic. We use constant bitrate streaming for the video traffic in order not to combine the effects of CC for adaptive bitrate streaming on top of the transport layer CC. For repeatability and reproducibility, we use a pre-recorded video that contains similar motions like in a drone flight. We set Gstreamer's RTP jitter buffer to 150 ms to accommodate late arrivals. We selected 150 ms to maintain the playback latency below the application requirement of 300 ms [15].

5) MULTIPATH TRANSPORT LAYER SETUP

We evaluate MPTCP and MP-DCCP to compare the trade-off between reliable versus unreliable MP transmission on transport layer. By studying these protocols, we aim to find out the achievable MP performance from two contrary perspectives: Maximizing the communication reliability versus minimizing the end-to-end latency. We focus on the schedulers that can serve this aim. Thereby, we employ the LowRTT and BLocking ESTimation (BLEST) schedulers to minimize the end-to-end latency, and the redundant scheduler to maximize the communication realibility with MPTCP. The LowRTT scheduler is designed to select an available path with the lowest RTT estimate, and BLEST aims to minimize the Out-of-order (OFO) packet arrivals by optimizing the MPTCP send window [25]. The redundant scheduler employs best-effort delivery over all the available links in a redundant manner and hence, it can push the overall reliability to a maximum on MP level. In MP-DCCP, we utilize the Cheapest Path First (CPF) and redundant schedulers. We use CPF instead of LowRTT because the LowRTT in MP-DCCP persistently uses the link with the lowest RTT based on our observations from measurements. Even if the link with the lowest RTT experiences congestion, it waits till the CWND of that link becomes available rather using the link with higher RTT. However, the LowRTT in MPTCP can use a link with a higher RTT if the congestion window of the link with the lowest RTT becomes full [26]. Therefore, we rather utilize the CPF scheduler in MP-DCCP to achieve a fair comparison between MPTCP and MP-DCCP. In CPF, each link is assigned with a priority (or a cost), and it schedules the packets on a link with the highest priority. If the CWND of the link with the highest priority becomes full, CPF still schedules the packets on the next link with less priority [56]. In our scenario, this MP-DCCP scheduler functions more similarly to the LowRTT of MPTCP if we assign a higher priority to the LEO link than the LTE link. As the measured RTT on the LEO link is always lower than that of LTE, CPF works the same way as LowRTT with LEO having

the highest priority. Hence, we employ the CPF in our study.

In regard to CCs, we aim to evaluate the ones that take different metrics to detect congestion, and find out the most favorable CC algorithms to handle the heterogeneity between cellular and LEO links. For this reason, we select Cubic, NewReno and weighted Vegas (wVegas) with MPTCP. Cubic and NewReno are uncoupled CCs, so they treat each underlying TCP subflow independently taking packet losses as congestion signal [22]. Although Cubic function is more favorable than NewReno's additive/multiplicative functions to quickly recover from congestion events, we still leverage NewReno since it is the only common CC algorithm in MPTCP and MP-DCCP. Using NewReno enable us to perform fair comparison between the two transport protocols. wVegas is a coupled CC and rather estimates the queing delay to detect path congestion [23]. In MP-DCCP, we used the NewReno (known as *Congestion Control Identifier 2 (CCID2)*) and Bottleneck Bandwidth and Round-trip Propagation Time (BBR) (known as *CCID5*). BBR is based on the bandwidth and RTT estimates, and it aims to maximize the data rate while reducing the queuing delay and bufferbloat [21]. In addition, we also analyze the effect of transport layer retransmissions on the MP networking performance by comparing MPTCP and MP-DCCP under similar configurations. This aspect is significant in understanding the effect of HoL blocking in this scenario. We set packet reordering engine to *fixed* in MP-DCCP, which ensures in-order arrival. We enable reordering to have a fair comparison with MPTCP.

We run a total of 24 measurements to test different schedulers and CCs combinations with both protocols. Each test runs for 30 minutes, which is a reasonable upper bound for a typical drone flight [63]. Next section elaborates on the measurement results in a single link as well as multilink level using the MP transport protocols, and also gives insight on the achievable application performance.

IV. MEASUREMENT RESULTS

We begin this section by providing insights on the achieved single link performance with TCP and DCCP in terms of goodput and RTT. Next, we evaluate the MP networking performance using MPTCP and MP-DCCP protocols. We analyze the achieved path utilization, MPTCP RTT and retransmissions, and compare loss- versus model-based CC, and the influence of transport-layer retransmissions.

Lastly, we evaluate the application-level packet losses and achieved video delivery performance in terms of Frames Per Second (FPS), playback latency and received video quality at different video bitrates with MPTCP.

A. SINGLE LINK PERFORMANCE

This section analyzes how TCP and DCCP treat individual links with the RP traffic over single path. Evaluating single link performance is essential to be able to correlate the individual contributions of the links to the measured MP networking performance. Figure 5 shows the achieved goodput and RTT over LTE and LEO links using NewReno CC.

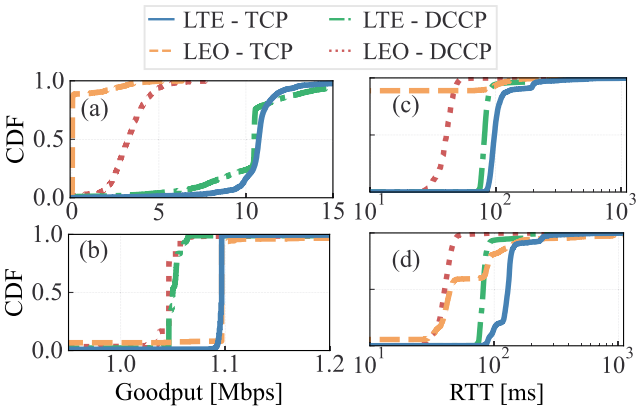

FIGURE 5. Achieved single-link goodput (a)(b) and RTT (c)(d) results using NewReno CC. Top and bottom figures correspond to uplink and downlink channels, respectively. The LEO link is exposed to congestion triggers due to high link-layer losses with video traffic.

In Figure 5 (a), while the LTE link can maintain 10 Mbps traffic, the rate reduces by ≈70% with DCCP and 95% with TCP on the LEO link. This further reduction is due to RTP video CC taking conservative sending bitrate decisions in addition to the reduction in congestion window due to the LEO link losses. We have a dedicated section regarding how the RTP influence the achievable goodput in Appendix C. Beside RTP, excessive retransmissions on TCP due to link-layer losses also influence the goodput reduction.

In Figure 5 (b), both protocols can maintain the low-rate control traffic over both links. The goodput is slightly higher with TCP potentially due to the use of different `iPerf` applications with DCCP and TCP.

As for RTT results in Figure 5 (c), RTT in LTE is around 100 ms and is slightly higher with TCP due to retransmissions. The RTT of the LEO link is not recorded up to 90% of the time due to the goodput bottleneck we observe in Figure 5 (a). Nevertheless, the RTT is ≈50 ms less than that of LTE, mainly because of the differences in the link-layer latencies. Lastly, in Figure 5 (d), RTT on LTE downlink with TCP is ≈10 ms higher than DCCP possibly due to additional video acknowledgement traffic.

Overall, these results show that the behavior of the cellular and LEO links are heterogeneous in terms of not only packet loss rates and goodput, but also RTT. High link-layer losses on the LEO is the root cause of this heterogeneity.

Now that we assessed the single link performance, we investigate the achieved MP orchestration performance from different transport-layer perspectives in the next section.

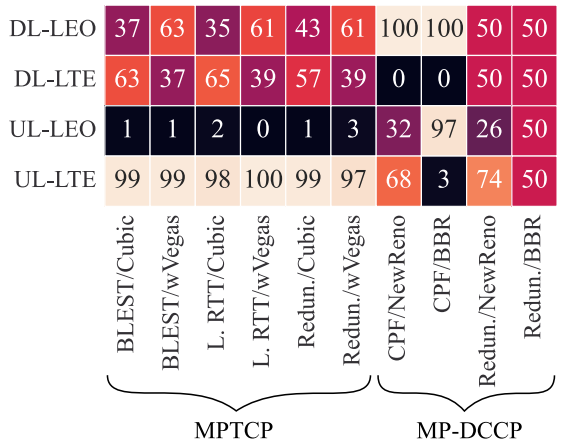

FIGURE 6. Comparison of path utilization (in %) using different scheduler and CCs. In MPTCP, LEO link is congested on uplink due to high link layer losses while in MP-DCCP, LEO link is always utilized on downlink due to its lower RTT.

Takeaway

High LEO link-layer losses cause aggressive CC behavior for the video traffic and limit the achievable goodput, whereas the low-rate control traffic is not affected. Video acknowledgement traffic can increase the observed end-to-end latency for the control traffic on downlink.

B. MULTIPATH NETWORKING PERFORMANCE

We start this section by providing a comprehensive analysis on the achieved path utilization results with all the tested schedulers and CCs combinations to evaluate their MP orchestration performance. Subsequently, we narrow down our analysis with particular schedulers and CCs that provide the most insightful findings and takeaways in each subsection.

1) PATH UTILIZATION

In this section, we compare the achieved path utilization on MP level over each link with different scheduler and CC configurations. The aim is to find out how the path selection decisions differ with the selected CC and scheduling algorithms. We also evaluate how the size of the traffic flow influences the utilization of individual links. We measure path utilization by comparing the scheduled packets on the sendings queues of individual links in MPTCP. MP-DCCP path utilization results are estimated based on the transmitted packets over each link.

Figure 6 presents the measured path utilization over individual links. Firstly, the size of the traffic flow largely influences the path utilization. In MPTCP, while the data traffic is more fairly distributed on downlink with control traffic (1 Mbps), the video flow (10 Mbps) is almost always routed over the LTE link. Due to high link losses, the LEO

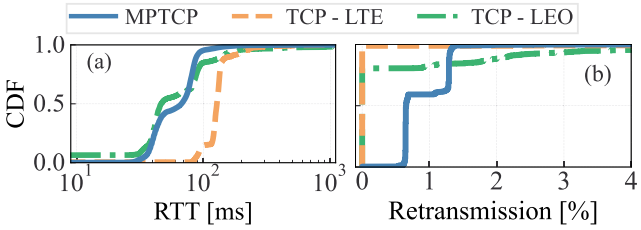

FIGURE 7. Achieved RTT and retransmission performance with MPTCP vs TCP using BLEST scheduler with wVegas CC. While MPTCP-level RTT lies in between individual performance of the links, retransmissions on MPTCP occur more frequently than on TCP-level due to HoL blocking.

Congestion Window (CWND) experiences bottlenecks and in addition, the RTP protocol lowers the video sending bitrate if LEO link is used. Whereas, we measure that an `iPerf` traffic at 10 Mbps could achieve ≈20% utilization on the LEO link instead of 3% with RTP traffic, as can be seen in Figure 6. As the RTP network quality reports to the video application triggers congestion alerts, LEO path utilization further reduces below 20% (more details are in Appendix C). On downlink, LEO path utilization reaches up to 60% with the wVegas CC since wVegas performs load balancing between available links using queuing delay [64]. Other CCs favor the LTE link more due to its lower link losses.

In MP-DCCP on downlink, the LEO link is always selected since its CWND does not fill up with the 1 Mbps control traffic. Whereas with 10 Mbps flow on uplink, LTE link utilization increases to 32% using NewReno since NewReno is sensitive to LEO link-layer losses and detects congestion. Lastly, the redundant scheduler can achieve fair utilization with BBR as it takes bandwidth and delay estimates into account rather than packet losses.

These outcomes highlight the challenge of achieving a fair utilization over each link, especially for the video traffic, although we tested various schedulers and CCs with distinct properties. As the LEO link has a lower RTT and a higher packet losses than LTE, this contradiction usually creates a binary decision for the link selection. Consequently, it becomes challenging to balance the path utilization.

2) MPTCP RTT AND RETRANSMISSION PERFORMANCE

We analyze RTTs and retransmissions with MPTCP to find out how efficiently it can orchestrate the LTE and LEO links together. We analyze the downlink channel since path utilization over both links is more fairly distributed compared to uplink.

Figure 7 (a) compares the achieved RTT performance at MPTCP versus TCP. We show results from BLEST scheduler with wVegas CC to highlight the main takeaways, and we include the results from all the scheduler and CC combinations in Appendix D. The average RTT on MPTCP-level is ≈80 ms. MPTCP-level RTT is slightly improved compared with the TCP-level on LTE since part of the control traffic is routed over the LEO that has lower link delay.

In Figure 7 (b), MPTCP is exposed to retransmissions more frequent than that of TCP on individual links. While MPTCP-level retransmissions occur 99.67% of the time, it is 1.22% and 19.59% with LTE and LEO on TCP-level, respectively. This is mainly due to the RTT heterogeneity observed between LTE and LEO links. This heterogeneity introduces late arrivals at the receiver and therefore, HoL blocking occurs at the MPTCP receive buffer. HoL blocking creates a communication overhead on MP-level due to excessive retransmissions, and it becomes a challenge for MPTCP to orchestrate the LTE and LEO links together.

Figure 8 (a) and (b) compare the MPTCP-level RTT and retransmission performance with respect to different scheduler and CCs. In general, the distribution of the data points are similar with BLEST and redundant schedulers, except that the BLEST is exposed to more RTT outliers >300 ms. Average RTT with the redundant scheduler is ≈10 ms less than that of BLEST since redundant scheduling helps in recovering the packet losses quicker, and thus it mitigates HoL blocking. Nonetheless, it comes at the expense of data overhead and reduced goodput. Although BLEST is also designed to minimize the HoL blocking [25], its performance gain is limited compared with the redundant scheduler in this scenario. Delay-based wVegas CC reduces RTT compared with the loss-based Cubic CCs. Using wVegas with redundant scheduler reduces the mean RTT by 40 ms compared with other configurations. This combination also manages to minimize the RTT outliers down to 400 ms, which goes as high as 10 s with other combinations.

In Figure 8 (b), using the redundant scheduler decreases the retransmission rate by ≈0.25% on average since out-of-order arrivals can be resolved quicker. Selection of the CCs does not largely influence the retransmissions and the mean retransmission rates stay between 0.8-0.85%. Although wVegas relies on delay measurements for congestion detection, it cannot reduce the retransmissions that are induced due to link latency heterogeneity. This is because the LEO link has lower link delay than LTE but contrarily, it is also exposed higher link losses.

In Figure 8 (c), we show the correlation between the RTT and retransmission performance using BLEST scheduler with wVegas CC since other configurations also have similar distributions. Dashed lines show the median for RTT and retransmission metrics. Majority of the data are accumulated within 30-110 ms RTT along with 0-3% retransmissions. On a first sight, no direct correlation exists between retransmissions and RTT. Nevertheless, retransmission outliers above 3% occur if the RTT is above its median value. Hence, those excessive retransmissions degrade the RTTs performance. In addition, the other RTT effect is that its outliers above 200 ms occur due to link-layer latency fluctuations on the LTE link. Hence, they are not correlated with increased retransmissions.

In conclusion, MPTCP is exposed to HoL blocking while orchestrating the LTE and LEO links due to link heterogeneity. Although most of the CCs and schedulers we tested do not improve the situation, using redundant scheduling helps minimize not only the retransmissions, but also the RTT outliers above 300 ms. The wVegas CC also helps reduce the average RTT by 19 ms, which is ≈22% less compared with that of Cubic.

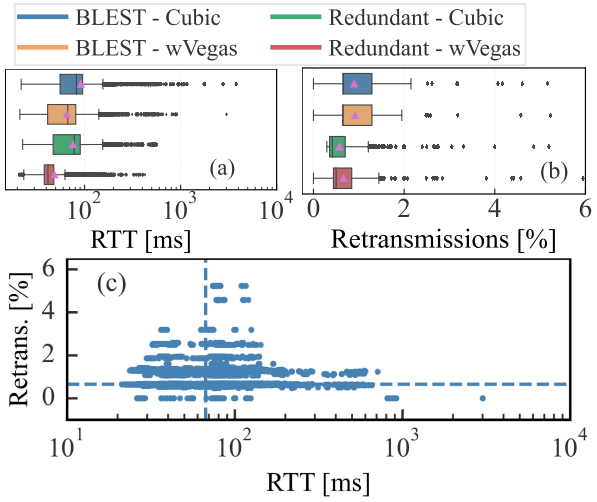

FIGURE 8. Comparison of the MPTCP-level RTT and retransmission performance in (a) and (b) along with the correlation between RTT and retransmissions in (c) using wVegas CC and BLEST scheduler. Vertical and horizontal lines in (c) represent the median values for RTT and retransmissions. Redundant scheduler handles the HoL blocking quicker while RTT outliers above 300 ms and excessive retransmissions occur more frequently with other configurations.

3) INFLUENCE OF CONGESTION CONTROL

As we observed that the path utilization is largely influenced by the selected CC in Subsubsection IV-B1, we have a detailed look on how a CC algorithm influences the goodput and delay performance. For this purpose, we compare the loss-based NewReno versus model-based BBR using MP-DCCP. We especially selected these CCs since they take different link parameters into account for congestion detection. We use the CPF scheduler to evaluate the achievable minimum delay, and we focus on the video traffic since it is more challenging for CCs to orchestrate it compared with the control traffic. We performed the evaluations for goodput and One-way Delay (OWD) metrics. *OWD* represents the end-to-end latency in one direction and we use this metric with MP-DCCP since MP-DCCP is based on unreliable transmission.

Figure 9 demonstrates the achieved goodput and OWD performance on uplink channel with NewReno (a)(c) and BBR (b)(d). Firstly, the target 10 Mbps goodput can be achieved with both configurations. The difference is the amount of traffic that is routed over each link, which is similar to our observation in path utilization results. While BBR largely prefers the LEO link, NewReno allocates the ≈2/3rd of traffic to LTE. Hence, BBR is a better option to avoid congestion on the LEO link. It can also potentially achieve higher overall goodput on MP-level (depending on CWND

performance of individual links at higher traffic rates) since it can better utilize the aggregated capacity from both links. Whereas, the achievable goodput with NewReno is majorly determined by the LTE link capacity only.

As for OWD, the results are largely influenced by the utilization of the LEO path, since its link-layer delay is lower than that of LTE. Then, BBR favors the LEO link more than LTE. Consequently, the mean OWD on MP-level is reduced down to 23 ms with BBR, which is improved by ≈24% compared to that of NewReno. Therefore, this finding highlights that the usage of BBR not only helps in maximizing the overall achievable goodput on MP-level, but also in minimizing the OWD. Whereas, NewReno achieves a fairer traffic distribution over individual paths.

4) INFLUENCE OF TRANSPORT LAYER RETRANSMISSIONS

As the MPTCP retransmission analysis highlighted the occurence of HoL blocking due to link heterogeneity, we investigate how transport-layer retransmissions affect the overall MP networking in detail. As retransmissions cannot be completely ruled out in MPTCP, we design a scenario by comparing MPTCP (with retransmissions) and MP-DCCP (without retransmissions) using similar scheduler and CC configurations. We set up a 10 Mbps `iPerf` traffic with both protocols using the NewReno CC since it is the common CC in both transport protocols. We employ the Round-robin (RR) and redundant schedulers in this scenario to avoid the effect of complex path selection algorithms particular to other schedulers. Beside retransmissions, MPTCP also has other reliability-centric mechanisms that MP-DCCP does not employ, however, their influence on path selection can be negligible in this scenario.

Figure 10 compares the achieved goodput and OWD performance with MPTCP and MP-DCCP. In Figure 10 (a), all the configurations can achieve an average of 10 Mbps, except that the average rate stays around 8 Mbps with MP-DCCP RR. This is related to the higher path utilization of the LEO link in MP-DCCP and consequently, its CWND experiences bottleneck. LEO utilization is 42% and 16% with RR in MP-DCCP and MPTCP, respectively.

As for OWD in Figure 10 (b), MP-DCCP RR manages to reduce the OWD down to 35 ms, which is 19% less than the average OWD achieved with other combinations. This is mainly due to the retransmissions with MPTCP that lead to increased OWD outliers. This can be clearly observed in Figure 10 (c), which presents the OWD outliers above 300 ms. For MPTCP RR, we observe a significant number of data points above 1000 ms OWD. Whereas in MP-DCCP RR, all the measured delays stay less than 1000 ms. This bottlenecks disappears when MPTCP employs the redundant scheduler since it can resolve retransmissions quicker by sending packets over both links simultaneously. Hence, employing a transport protocol without retransmissions is especially helpful for avoiding late arrivals with respect to the delay requirements of the RP application. In addition, the default MPTCP Retransmission Timeout (RTO) values are not suitable for the application-level delay requirements of RP, as we elaborate in the next section.

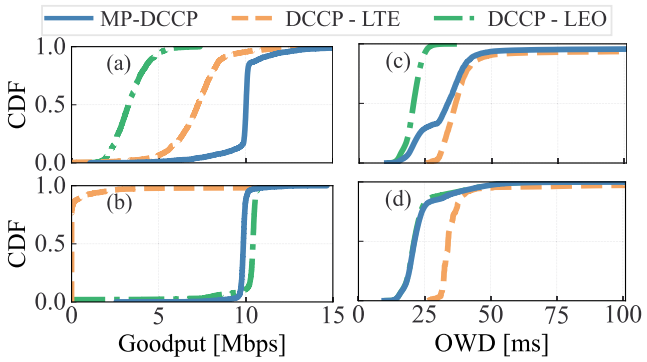

FIGURE 9. Comparison of the loss-based NewReno (figures at the top) and model-based BBR performance (figures at the bottom). As the BBR is less sensitive to LEO link layer losses, it maximizes the usage of the LEO link and in turn, this improves the OWD performance by ≈24%.

Takeaway

It is challenging to achieve a fair path utilization between LTE and LEO due to link heterogeneity. Also, size of the traffic flow significantly influences the path utilization. For CCs, using BBR helps improve the overall achievable goodput and reduces the OWD by 24%. In MPTCP, redundant scheduler is the most favorable option in to minimize HoL blocking and to avoid late arrivals. Lastly, Avoiding transport-layer retransmissions with MP-DCCP improves the OWD performance by 19%.

C. RP APPLICATION PERFORMANCE

Now that we analyzed the MP networking performance in detail, we investigate the achievable application-layer performance for the RP traffic in this section. We analyze the essential application metrics for control and video traffic. As the achieved goodput and OWD results suffice the control traffic requirements, we particularly focus on the communication reliability metric for the control flow. We aim to find out whether the 99.999% reliability demand can be met in particular configurations. As for the video traffic, we evaluate the video delivery performance in terms of FPS, playback latency and the received video quality using the Structural SIMilarity (SSIM) metric.

SSIM evaluates the quality of received video frames by measuring the degradation of the luminance, contrast, and structure information [65]. It has a range between 0 and 1 (1 being the best quality), and it is a subjective metric that is determined based on human perception. We set the minimum quality threshold to 0.5 based on our visual evaluations on the received video quality to correctly sense the environment shown in the video.

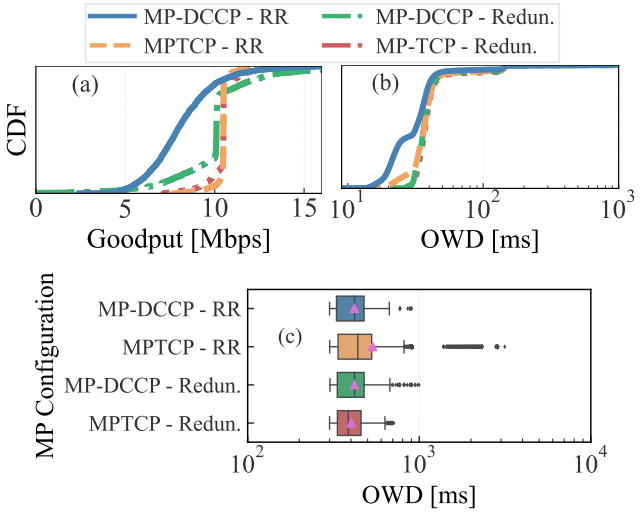

FIGURE 10. Comparison of the MPTCP and MP-DCCP due to transport layer retransmissions in terms of goodput (a), OWD (b) and OWD outliers above 300 ms (c) using NewReno CC. OWDs above 1000 ms with MPTCP RR are the direct outcome of transport retransmissions and hence, they do not occur without retransmissions in MP-DCCP.

As an MP-DCCP video pipeline is not available at the time conducting our study, we provide this analysis using MPTCP. Although MPTCP may not be the primary choice for video delivery in our scenario due to HoL blocking, it provides a lower bound on achievable video delivery performance, especially in terms of playback latency.

1) APPLICATION-LEVEL PACKET LOSS ANALYSIS

In this section, we analyze the packet loss performance for control traffic using two schedulers: 1. LowRTT to analyze the achievable communication reliability while minimizing the end-to-end latency, 2. The redundant scheduler to find out the upper bound for achievable reliability when using both wireless paths in a best-effort manner. We again utilize the CPF scheduler with MP-DCCP since its behavior is more similar to the LowRTT scheduler of MPTCP in our scenario (as described in Subsubsection III-B5). We again use the NewReno CC with both MPTCP and MP-DCCP.

As the RP application treats late arrivals after the latency threshold (300 ms) as packet loss, we derive the application-level losses based on the end-to-end latency bound of the use case [1]. Therefore, with MPTCP, we compute the packet losses based on the retransmissions that occur after 300 ms since the first transmission of a packet. This method gives an upper-bound on packet loss estimations since retransmissions can occur not only due to packet losses, but also because of out-of-order arrivals. For MP-DCCP, we compute the packet losses including the late arrivals after 300 ms threshold using the transport sequence numbers.

We present the packet loss results on MP-level in Table 3. With LowRTT/CPF schedulers, while MPTCP can achieve 0.02% loss rate, it stays an order of magnitude higher with MP-DCCP. This is mainly related to the path utilization of the LEO link in each MP protocol. As MP-DCCP always uses the LEO link (as shown in Figure 6), it is more affected by the high LEO link-layer losses. Whereas, MPTCP utilizes the LEO link around 30%.

TABLE 3. Application-level packet loss analysis.

Feature	LowRTT/CPF	Redundant
MPTCP (LowRTT)	0.02149%	0.02235%
MP-DCCP (CPF)	0.2045%	0.0006%

In regard to redundant schedulers, we observe a contrary situation. MPTCP does not benefit from redundant scheduling and the packet loss rate still stays ≈0.02%. Although redundant transmission minimizes the HoL blocking in MPTCP, late arrivals still exist with respect to the application latency threshold. Indeed, the TCP retransmission scheme does not bring significant benefit to the RP application due to the default RTO values. It already takes 200 ms to trigger the first RTO [66], which then makes it challenging to meet the application-level latency bound of 300 ms. This indicates that a maximum of one retransmission can merely bring benefit for the application, and any further retransmissions are actually waste of resources. Therefore, RTO as well as the number of retransmissions should be optimized in MPTCP with respect to the perceived link-layer latency of the LTE and LEO to improve the achievable communication reliability.

As for MP-DCCP, redundant scheduling manages to reduce the packet losses by three orders of magnitude and achieves 99.9994% communication reliability. This finding further supports the innecessity of transport-layer retransmissions for the RP scenario and suggests that using two paths in a redundant manner can meet the reliability requirements. In conclusion, while only MP-DCCP with redundant scheduling can meet the stringent 99.999% reliability demand of the use case, MPTCP may achieve improved results if RTO values are fine-tuned for the scenario.

2) VIDEO DELIVERY ANALYSIS

This section presents the achieved video delivery performance in MPTCP using wVegas CC with redundant and BLEST schedulers since they are the outstanding configurations to minimize the delay and HoL blocking, as we discussed in Subsubsection IV-B2. We perform two different bitrates at 10 Mbps and 5 Mbps to highlight the tradeoff between bitrate and playback latency.

We present the measured video delivery performance in Figure 11. Firstly, the default 30 FPS can be maintained during the entire video playback session. It experiences frequent but minor drops at the warm-up phase. Then, minor drops occur at the times, when playback latency spikes are observed due to the LTE link latency spikes.

In Figure 11 (b), we observe the tradeoff between the playback latency and video bitrate. When the video bitrate is set to 10 Mbps, the playback exceeds the latency threshold

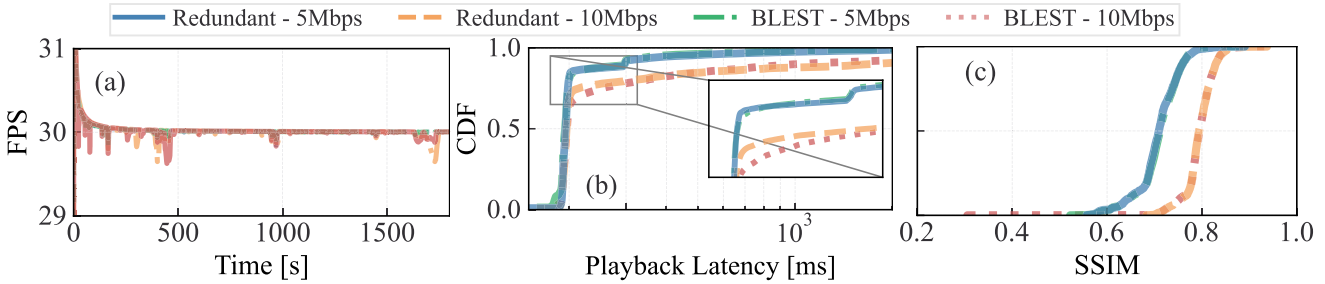

FIGURE 11. Measured FPS (a), playback latency (b) and SSIM (received video quality) (c) performance. Playback latency experiences spikes up to 10 s when video bitrate is set to 10 Mbps due to network latency spikes. This behavior is smoothened if video bitrate is reduced to 5 Mbps. Also, running the video at 5 Mbps lowers the SSIM ≈0.1 on average.

of 300 ms up to 20% of the time, whereas it is around 10% at 5 Mbps. Additionaly, the redundant scheduler can slightly reduce the latency outliers >300 ms compared to BLEST. This is correlated with the RTT performance of the redundant scheduler. We also noticed 3 playback latency spikes that are correlated with the LTE network latency spikes (as presented in Figure 4 (b)). Although those network latency spikes are between ≈300-1000ms, the playback latency hits as high as 10 s when the bitrate was set 10 Mbps. This is very likely related to bufferbloating due to link congestion and out-of-order packet arrivals. Nevertheless, playback latency spikes do not go above 1000 ms when we set the bitrate to 5 Mbps.

The received frame quality is also correlated with the bitrate as shown in Figure 11 (c). Although SSIM reduces by ≈0.1 on average at 5 Mbps bitrate, measured frame quality is maintained above the minimum threshold of 0.5 more than 95% of the time in any configuration. Overall, MPTCP could manage to support the video delivery within the application latency threshold as well as sufficient quality up to 90% of the time.

Takeaway

Using LTE and LEO in a redundant way without retransmissions can achieve the 99.999% communication reliability requirement for the control traffic. Video delivery with stable FPS and sufficient received frame quality can be achieved at 10 Mbps. However, network latency spikes cause severe disruptions on playback latency. Reducing the bitrate to 5 Mbps can ensure to meet 300 ms latency threshold up to 90% of time and is still sufficient to achieve HD 1080p video resolution [61].

V. DISCUSSION

A. WHAT ARE THE MOST SUITABLE MP TRANSPORT CONFIGURATIONS WITH CELLULAR AND LEO LINKS FOR RP APPLICATION REQUIREMENTS?

The MP measurement analysis highlights the particular challenge when it comes to utilizing the LEO path efficiently along with a cellular link due to the large heterogeneity in link-layer packet losses. A loss-based NewReno CC versus model-based BBR CC behaves completely opposite in terms of invididual path utilization. Hence, there is an open room for fine-tuning a CC scheme that can utilize the both paths in a more fair way. Nevertheless, using BBR can become a favored choice to avoid goodput bottlenecks and to minimize the overall OWD in this scenario. As for schedulers, although redundant is not an efficient method, it is the only configuration that can meet the communication reliability demand of the control traffic. Hence, it remains as the only choice for the RP scenario. Lastly, transport-layer retransmissions become unnecessary in our scenario due to the stringent application delay requirements with respect to the default RTO in MPTCP. Fine-tuning the RTO may bring benefit [67], [68] however, there is only a small room for fine-tuning since the link-layer RTT are already within 60-100 ms range.

B. CAN MPTCP BE USED TO ORCHESTRATE CELLULAR AND LEO LINKS TOGETHER?

The MP networking analysis in Subsection IV-B reveal that the performance gain from using LEO link is limited with large flows (video traffic in our case), mainly due to high link losses creating CWND bottlenecks and HoL blocking. LEO and LTE hold opposite link characteristics not only in terms of link losses, but also RTTs. This brings difficulties for CC algorithms to handle both links in an efficient way. Overall, the reliability-centric architecture of TCP challenges the co-orchestration of LTE and LEO links together.

C. PERFORMANCE TRADE-OFF BETWEEN SINGLE-PATH AND MP CONNECTIVITY

In general, MP connectivity can be beneficial to boost communication performance toward multiple QoS metrics. However, we observe a trade-off between single-path and MP connectivity for different performance metrics in our scenario. As for goodput, MP connectivity can aggregate the available link capacity from both links to maximize the achievable goodput. However, this is not the case in MPTCP in our scenario since the LEO link is largely congested while running video traffic. As we analyze the single-path results in Figure 5, the LEO link cannot cope with the data rate demands of the video traffic independent of the underlying MPTCP

configurations. Therefore, the MP-level goodput and latency performance is largely determined by the capabilities of the LTE link. In MP-DCCP, on the other hand, MP can improve the overall goodput especially with BBR since it reduces the congestion sensitivity on the LEO path by using bandwidth and delay parameters rather than packet loss. As for latency, the benefit of MP connectivity is relative with respect to the considered single-path scenario. Compared with single-path LTE, MP connectivity improve the end-to-end latency by opportunistically using the LEO path with lower delay. However, the LEO single-path can perform better than MP in minimizing the latency as long as it can handle congestions. Lastly, and most importantly, the benefits of MP connectivity become particularly evident when it comes to improving the communication reliability. While single-path connectivity can provide a maximum of 99-99.9% reliability, simultaneous use of the cellular and LEO improves the reliability by two orders of magnitude. Achieving such high reliability has a vital role in enabling safe RP operations. In summary, the main benefits of MP connectivity for our scenario are the improvement in communication reliability as well as increased MP capacity.

D. GENERALITY OF THE MP PERFORMANCE RESULTS

Link-level RTT and reliability are the two vital metrics that influence the MP transport protocols with schedulers and CCs. We set up the link parameters in our emulation scenario based on experimental measurements to ensure that our findings are applicable in real-life. Nevertheless, LEO link latency can vary depending on the distance between end users [60]. If we change the location of the server for our measurement, LEO latency can become higher than LTE, and in turn, this can influence the delay-based scheduler and CC algorithms. Similarly, the LTE latency can also be influenced by the client-server distance. In regard to LEO link-layer losses, we assume that the measured high losses are related to the satellite HOs, however this cannot be proven since no low-layer information are available from Starlink. Based on this assumption, LEO link losses are less likely to alter in the future. However, if another LEO constellation is used (e.g., OneWeb), both RTT and link reliability can differ. Whereas for the cellular link, LTE is a mature standard by now and the measured 99.9% reliability is aligned with the 3GPP specifications [1, Table 18]. Overall, while our findings are valid for a real-measurement case with LTE and LEO for RP operations in close vicinity, the MP performance may differ for different server locations or LEO links.

E. INDICATIONS OF THE TAKEAWAYS IN MEASUREMENT ANALYSIS FOR FUTURE RESEARCH DIRECTIONS

1. Path utilization results highlight unfairness with the tested CC and scheduler configurations. Especially the sensitivity of the CCs toward packet losses or link latency play a key role in these results. Future research can investigate the trade-off between packet loss and RTT metrics in designing a novel CC for the co-orchestration of cellular and LEO links to better balance the path utilization. 2. Although only redundant scheduler could achieve 99.999% communication reliability, optimizing the retransmission settings can bring the possibility of increasing the reliability with retransmissions. In turn, this can enable employing more resource-efficient scheduling schemes other than redundant transmission. 3. Video delivery results point to the performance bottleneck with playback latency. Even though this is mainly related to the MPTCP retransmissions, RTP also play a particular role in this metric due to the network latency spikes that occur on the LTE link. Therefore, future research can also investigate the video application to smoothen the effects of network latency spikes on the video playback. 4. Lastly, as the control and video traffic have different QoS priority (reliability versus latency), this heterogeneity brings a challenge of fitting a single CC and scheduler that can maximize the performance for each traffic. Different scheduling and congestion algorithm techniques can be considered for each traffic individually. Even cross-layer optimization techniques can be studied to optimize the congestion and scheduling decisions in a joint manner.

VI. CONCLUSION

In this work, we presented our findings with different MP transport schemes and configurations to support the RP applications over cellular and LEO links. Overall, the results highlight that HD-video stream along with 99.999% reliable control traffic can be achieved in particular MP transport configurations. Hence, using single cellular and LEO links in a MP fashion can potentially suffice the QoS requirements and enable the future RP applications in the sky. Yet, there is an open room for developing CC and scheduler algorithms to optimize path selection and achieve fairer utilization. Such algorithm can also be studied to prioritize reliability on MP-level while utilizing the links in a more efficient way. MPTCP is also challenged in orchestrating these heterogeneous links, especially due to high loss rates on the LEO link. Hence, a UDP-based protocol such as MP-DCCP is more suitable not only to minimize the end-to-end latency, but also to avoid unnecessary retransmissions. Future work can also take different directions such as including the cellular link performance at higher altitudes than 120 m for eVTOL flights, the effect of mobility and flight attitude on LEO link performance and optimizing RTP video application toward network latency spikes. Investigation toward the formation of eVTOL/drone swarms to enable air-to-air connectivity as another connectivity option in a MP connectivity scenario can also be considered. Evaluation regarding the potential benefits of MP communications toward signal jamming and interference are also important research directions toward enabling secure operations in the sky.

APPENDIX A
VALIDATION OF THE EMULATED LINK SETTINGS

In order to verify the accuracy of the link emulations in our testbed, we compared the emulated link behaviors

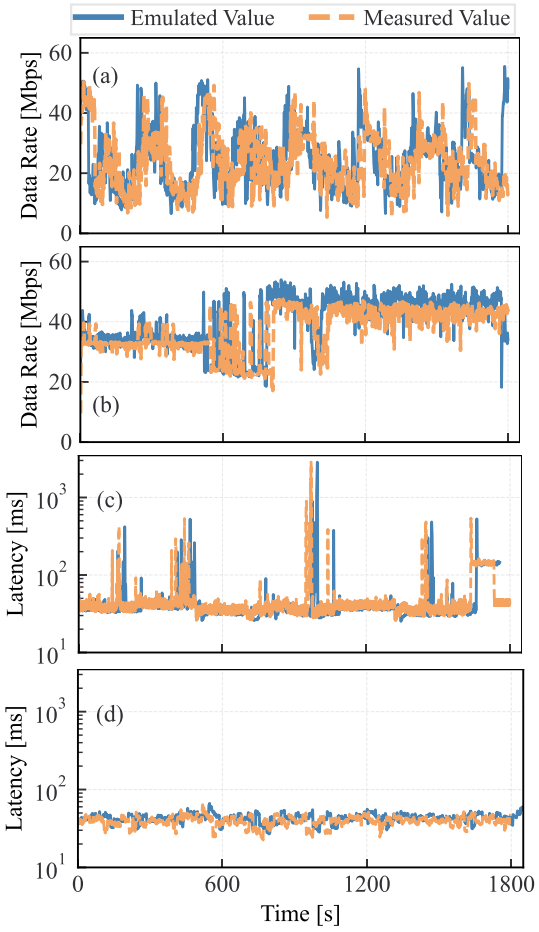

FIGURE 12. Comparison of the measured versus emulated link dynamics in terms of: (a) LTE link capacity on downlink, (b) LTE link capacity on uplink, (c) LTE link latency on uplink, and (d) LEO RTT. Correlation between emulated and measured traces are visually perceivable and the Pearson correlation analysis validates it.

against real-life measurements in terms of link capacity and end-to-end latency. Figure 12 compares the emulation link dynamics to real-life traces. We collected the data for emulated traces using `iPerf` tool. There is a slight time-shift in emulation versus measurement values due to the warm-up phase of the emulator until individual links are set up. In each figure, we can visually observe the correlation between the emulated link dynamics versus the actual ones occured during experimental measurements.

We further performed Pearson correlation analysis to validate our findings. When each experimental and emulated dataset are time-wise aligned, the Pearson correlation are 0.98, 0.94, 0.71 and 0.88 for LTE downlink capacity, LTE uplink capacity, LTE uplink latency and LEO RTT, respectively. Reduced correlation in LTE uplink latency is due to the slight differences between emulated and measured latency when latency spikes occur. Nevertheless, any value above 0.7 represents a strong and positive correlation between the emulated and actual link dynamics [69].

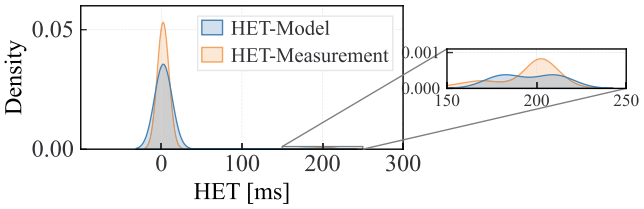

FIGURE 13. The distributions of the measured HET during drone flights with LTE and the resulting distribution model. During measurements, the mean HET was 20.01 ms with a 195.13 ms of standard deviation. The HET is primarily clustered around 1.9 ms, 3.8 ms, 9.38 ms and 200 ms. We model the HET with a Gaussian Mixture Model with a cluster size of 4.

APPENDIX B
LTE HANDOVER MODELING FOR AERIAL TRACES

We modeled the cellular HOs based on two parameters: I. Handover Execution Time (HET), II. HO frequency. While *Handover Execution Time (HET)* describes the time duration between the reception of the *RRCConnectionReconfiguration* packet from the source Base Station (BS) and the transmission of the *RRCConnectionReconfigurationComplete* packet at the target BS [70], HO frequency specifies how frequent a HO occurs. We model the HET based on the HO dataset we collected in the air [71]. Modeling these parameters in the emulator are significant in order to reflect the difference of the cellular network performance in the air compared to the ground [13]. During measurements, we observed that the mean HET is 20.01 ms with a standard deviation of 195.13 ms. The durations are primarily clustered around 1.9 ms, 3.8 ms, and 200 ms. We use the statistical DIP test [72] to test unimodality of the dataset, which measures the unimodality by the maximum difference over all sample points between the empirical distribution and unimodal distribution function. The DIP test tests the p-value of the collected HET dataset. The result is 0.032, which is less than the threshold value of 0.05, and hence, the data distribution is not unimodal. We could fit the HET distribution using a Gaussian mixture model with a cluster size of 4, and the resulting model with the generated HET is shown in Figure 13. We use the Expectation-Maximization algorithm to learn the parameters of a mixture model. A cluster size of 4 is determined to avoid overfitting/underfitting problem. We use the resulting distribution model as shown in Figure 13 to generate HETs over the LTE link in the emulator.

APPENDIX C
INFLUENCE OF THE RTP PROTOCOL ON MPTCP

In order to find out how the RTP protocol can influence the MP layer, we performed emulations on the uplink channels using iPerf and video traffic. The aim is to compare how path utilization varies when RTP protocol is used versus when a constant bitrate traffic is sent. This comparison highlights the individual contribution of the RTP protocol on path utilization beside the MPTCP. Therefore, we set both `iPerf` and the video traffic at 10 Mbps, and performed emulations over

TABLE 4. Statistical insights from achieved RTT performance with all the scheduler and CC combinations.

Configuration	TCP-LTE		TCP-LEO		MPTCP	
	Median [ms]	90^{th} Percentile [ms]	Median [ms]	90^{th} Percentile [ms]	Median [ms]	90^{th} Percentile [ms]
BLEST-BALIA	126.29	170.63	45.25	110.70	75.67	100.76
BLEST-Cubic	125.02	182.96	46.20	96.31	82.41	118.62
BLEST-OLIA	124.27	181.04	42.32	90.97	82.92	121.76
BLEST-wVegas	127.57	163.65	46.31	137.63	67.01	88.22
LowRTT-BALIA	111.00	154.13	45.37	96.11	80.43	105.09
LowRTT-Cubic	118.37	178.42	44.66	92.12	84.34	120.40
LowRTT-OLIA	128.84	167.41	45.60	100.56	79.01	108.30
LowRTT-wVegas	124.76	148.72	46.00	137.80	68.35	88.38
Redundant-BALIA	132.63	184.07	46.24	137.92	67.52	108.05
Redundant-Cubic	123.55	169.29	47.33	136.81	78.06	107.40
Redundant-OLIA	129.66	180.22	43.49	147.56	75.49	99.79
Redundant-wVegas	127.76	170.66	44.52	142.08	41.66	80.76

TABLE 5. Statistical insights from achieved retransmission performance with all the scheduler and CC combinations.

Configuration	TCP-LTE		TCP-LEO		MPTCP	
	Median [%]	90^{th} Percentile [%]	Median [%]	90^{th} Percentile [%]	Median [%]	90^{th} Percentile [%]
BLEST-BALIA	0	0	0	1.92	0.65	1.30
BLEST-Cubic	0	0	0	2.65	0.66	1.32
BLEST-OLIA	0	0	0	2.81	0.68	1.34
BLEST-wVegas	0	0	0	2.17	0.65	1.30
LowRTT-BALIA	0	0	0	2.94	0.65	1.31
LowRTT-Cubic	0	0	0	1.69	0.66	1.32
LowRTT-OLIA	0	0	0	3.23	0.65	1.31
LowRTT-wVegas	0	0	0	2.17	0.66	1.30
Redundant-BALIA	0	0	0	2.08	0.51	0.96
Redundant-Cubic	0	0	0	2.06	0.42	0.82
Redundant-OLIA	0	0	0	3.85	0.69	1.36
Redundant-wVegas	0	0	0	1.67	0.50	1.00

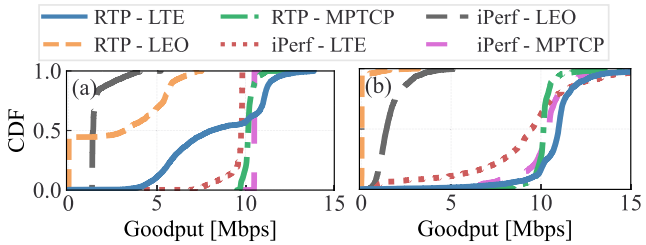

FIGURE 14. Comparison of the TCP and MPTCP performance of `iPerf` and RTP video traffic running at 10 Mbps over lossless emulated links in (a) and with link losses in (b). Overall, these figures show that running the RTP protocol on top of MPTCP creates different link utilization rates compared to a constant traffic since the video application also runs its own congestion mechanism. RTP protocol is more conservative toward using the LEO link due to its high link loss rate.

the LTE and LEO links with and without packet losses to evaluate achievable throughput over both links in different link conditions. Lastly, we derived the achieved goodput with `iPerf` and video traffic at TCP as well as MPTCP level.

Figure 14 shows the measured goodput performance in this scenario. When there are no emulated link losses in (a), MPTCP-level goodput performance of both traffic are at similar levels. However, the RTP traffic experiences more goodput fluctuations on TCP level.

In (b), high-loss rates on the LEO link limits achievable average throughput to 2 Mbps for the control traffic, and running RTP protocol over MPTCP creates even more conservative scenario, and the video traffic is almost completely obselete on the LEO link. This is very likely due to RTP network quality reports to the video application, which triggers congestion alert on the LEO link.

In summary, we observe contradicting behavior between the RTP and constant bitrate `iPerf` traffic in both figures. While the video traffic utilizes the LEO link more than that of `iPerf` in the absence of link losses, it goes vice versa when high link loss (0.17%) is introduced to the LEO link.

APPENDIX D
FURTHER INSIGHTS FROM ALL SCHEDULER AND CONGESTION CONTROL OPTIONS

This section provides further insights from our MP analysis by including all the available scheduler and CCs. We also

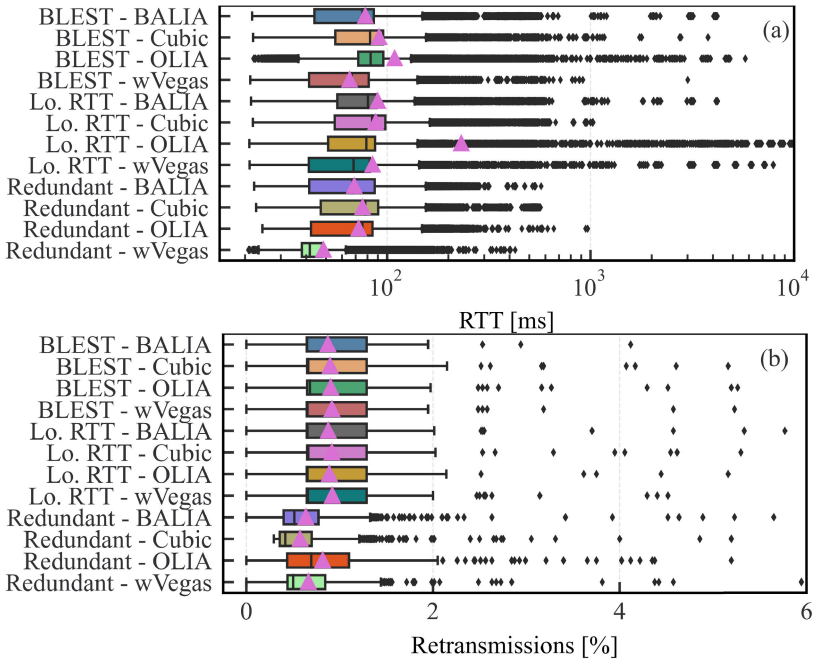

FIGURE 15. MPTCP-level RTT and Retransmission Distributions of All the Scheduler and CC Combinations. Purple triangles represent the achieved mean values.

include the OLIA and BAlanced LInked Adaptation (BALIA) CCs in this analysis. We particularly compare the TCP- and MPTCP-level RTT and retransmission performance of each scheduler and CC combination on Downlink (DL) channel in Table 4 and Table 5. Hence, these tables complement the results in Figure 7. We show the median and 90^{th} percentiles for RTT and retransmission values in Table 4 and Table 5 instead of mean values to avoid bias from outliers.

In Table 4, we observe that the MPTCP-level median RTT generally lies in between the TCP-level median RTTs in all the scheduler and CC combinations except the combination of redundant scheduler with the wVegas CC. This exception is due to the quick resolution of retransmissions on MPTCP-level and the wVegas, a delay-based CC, being less sensitive to LEO link losses compared to other loss-based CCs. While MPTCP-level median RTT usually is between 67 ms and 85 ms, it can go as high as 133 ms and 47 ms on the TCP-level over the LTE and the LEO links, respectively. With respect to 90^{th} percentiles, we notice that MPTCP RTT can outperform the TCP-level RTTs in particular scheduler and CC combinations, especially with the wVegas CC. Whereas, the RTT on TCP-LTE is consistently the worst in the 90^{th} percentile, mainly due to its higher link RTT compared to that of LEO.

In regard to retransmission results in Table 5, only MPTCP-level median retransmission is above 0% and it lies between 0.4% and 0.68%. 90^{th} percentile retransmissions on TCP-LEO is approximately two times higher compared to MPTCP-level. In addition, while MPTCP-level retransmissions occur between 97.0% and 100.0% of the time in any scheduler and CC combination, it reduces to the range of 0.9%-1.3% in TCP-LTE and 11.5% and 25.2% in TCP-LEO. These results also highlight that not only BLEST-wVegas combination (in Subsubsection IV-B2) but also all the tested combinations suffer from HoL blocking. Therefore, MPTCP-level retransmissions are higher than TCP-level retransmissions.

We also provide MPTCP-level RTT and retransmission distribution graphs for each scheduler and CC combination in Figure 15 (complementing Figure 8 (a) and Figure 8 (b)). Purple triangles represent the mean RTTs. In Figure 15 (a), frequent RTT outliers occur with the OLIA CC, especially above 1 s. In turn, its mean RTT is poorer compared to other CCs. Only exception to this observation is the redundant scheduler with OLIA since the redundant transmission can enable fast recovery from retransmissions. Contrarily, wVegas CC is able to minimize the mean RTT compared to other CCs with any scheduler combination. Using wVegas can improve the mean RTT as high as ≈20 ms with respect to other CCs.

As for Figure 15 (b), while any configuration with the BLEST and LowRTT scheduler achieve mean retransmission of ≈0.85%, using redundant scheduler can improve it by 0.2%.

ACKNOWLEDGMENT

The authors would like to thank Simon Zelenski with his support during the setup of the multipath emulation testbed, and Hendrik Cech for the provision of the Starlink network measurement traces.

REFERENCES

[1] A. Baltaci, E. Dinc, M. Ozger, A. Alabbasi, C. Cavdar, and D. Schupke, "A survey of wireless networks for future aerial communications (FACOM)," *IEEE Commun. Surveys Tuts.*, vol. 23, no. 4, pp. 2833–2884, 4th Quart., 2021, doi: 10.1109/COMST.2021.3103044.

[2] A. Volkert, H. Hackbarth, J. Lieb, and S. Kern, "Flight tests of ranges and latencies of a threefold redundant C2 multi-link solution for small drones in VLL airspace," in *Proc. Integr. Commun., Navigat. Surveill. Conf. (ICNS)*, 2019, pp. 1–14, doi: 10.1109/ICNSURV.2019.8735265.

[3] J. Güldenring, L. Koring, P. Gorczak, and C. Wietfeld, "Heterogeneous multilink aggregation for reliable UAV communication in maritime search and rescue missions," in *Proc. Int. Conf. Wireless Mobile Comput., Netw. Commun. (WiMob)*, Oct. 2019, pp. 215–220, doi: 10.1109/WiMOB.2019.8923123.

[4] R. Amorim, I. Z. Kovacs, J. Wigard, G. Pocovi, T. B. Sorensen, and P. Mogensen, "Improving drone's command and control link reliability through dual-network connectivity," in *Proc. IEEE 89th Veh. Technol. Conf. (VTC-Spring)*, Apr. 2019, pp. 1–6, doi: 10.1109/VTC-Spring.2019.8746579.

[5] R. Jewett. (2023). *EchoStar and SkyFive Demonstrate Satellite and Terrestrial Drone Command and Control With Airbus*. Accessed: Jul. 1, 2023. [Online]. Available: https://www.satellitetoday.com/mobility/2023/05/26/echostar-and-skyfive-demonstrate-satellite-and-terrestrial-drone-command-and-control-with-airbus/

[6] Icomera. (2023). *SureWAN Aggregation Protocol*. Accessed: Jul. 1, 2023. [Online]. Available: https://www.icomera.com/solutions/surewan-aggregation-protocol/

[7] *Study on New Radio (NR) to Support Non Terrestrial Networks (Release 15)*, document TR38.811, 3GPP, 2018. Accessed: Jan. 16, 2021. [Online]. Available: https://www.3gpp.org/ftp//specs/archive/38_series/38.811/

[8] *Unmanned Aerial System (UAS) Support in 3GPP (Release 17)*, document TS 22.125, 2019. Accessed: Jan. 11, 2021. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/22_series/22.125/

[9] Starlink. (2023). *STARLINK Specifications*. Accessed: Jun. 21, 2023. [Online]. Available: https://www.starlink.com/legal/documents/DOC-1002-69942-69

[10] OneWeb. (2022). *Space is the Future—OneWeb's Mission Continues*. Accessed: Jun. 21, 2023. [Online]. Available: https://assets.oneweb.net/s3fs-public/2022-08/AnnualReport_2022.pdf

[11] J. Säe, R. Wirén, J. Kauppi, J. Torsner, S. Andreev, and M. Valkama, "Reliability of UAV connectivity in dual-MNO networks: A performance measurement campaign," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, Jun. 2020, pp. 1–5, doi: 10.1109/ICCWorkshops49005.2020.9145477.

[12] P. Saxena, T. Dreibholz, H. Skinnemoen, Ö. Alay, M. A. Vazquez-Castro, S. Ferlin, and G. Acar, "Resilient hybrid SatCom and terrestrial networking for unmanned aerial vehicles," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, Jul. 2020, pp. 418–423, doi: 10.1109/INFOCOMWKSHPS50562.2020.9162943.

[13] A. Baltaci, H. Cech, N. Mohan, F. Geyer, V. Bajpai, J. Ott, and D. Schupke, "Analyzing real-time video delivery over cellular networks for remote piloting aerial vehicles," in *Proc. IMC*. New York, NY, USA: Association for Computing Machinery, 2022, pp. 98–112, doi: 10.1145/3517745.3561465.

[14] "Unmanned aircraft systems (UAS)," Int. Civil Aviation Org., Canada, Tech. Rep. Cir 328 AN/190, 2012. Accessed: Jul. 1, 2023. [Online]. Available: https://www.icao.int/meetings/uas/documents/circular%20328_en.pdf

[15] A. Baltaci, M. Klügel, F. Geyer, S. Duhovnikov, V. Bajpai, J. Ott, and D. Schupke, "Experimental UAV data traffic modeling and network performance analysis," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2021, pp. 1–10, doi: 10.1109/INFOCOM42981.2021.9488878.

[16] L. Parziale, D. T. Britt, C. Davis, J. Forrester, W. Liu, C. Matthews, and N. Rosselot, "TCP/IP tutorial and technical overview," IBM Corp., Tech. Rep., 2006. [Online]. Available: https://www.redbooks.ibm.com/redbooks/pdfs/gg243376.pdf

[17] A. Ford, C. Raiciu, M. Handley, O. Bonaventure, and C. Paasch, *TCP Extensions for Multipath Operation With Multiple Addresses*, document RFC 8684, 2020. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6824

[18] M. Amend, A. Brunstrom, A. Kassler, V. Rakocevic, and S. Johnson, "DCCP extensions for multipath operation with multiple addresses," IETF, USA, Tech. Rep. Internet-Draft, 2023. Accessed: May 25, 2023. [Online]. Available: https://datatracker.ietf.org/doc/draft-ietf-tsvwg-multipath-dccp/

[19] C. Xu, J. Zhao, and G.-M. Muntean, "Congestion control design for multipath transport protocols: A survey," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 4, pp. 2948–2969, 4th Quart., 2016.

[20] A. Walid, J. Hwang, Q. Peng, and S. Low. (2014). *Balia (Balanced Linked Adaptation) A New MPTCP Congestion Control Algorithm*. Accessed: Sep. 22, 2023. [Online]. Available: https://www.ietf.org/proceedings/90/slides/slides-90-mptcp-0.pdf

[21] D. Scholz, B. Jaeger, L. Schwaighofer, D. Raumer, F. Geyer, and G. Carle, "Towards a deeper understanding of TCP BBR congestion control," in *Proc. IFIP Netw. Conf. (IFIP Netw.) Workshops*, 2018, pp. 1–9.

[22] S. Ha, I. Rhee, and L. Xu, "CUBIC: A new TCP-friendly high-speed TCP variant," *ACM SIGOPS Operating Syst. Rev.*, vol. 42, no. 5, pp. 64–74, Jul. 2008, doi: 10.1145/1400097.1400105.

[23] B. Y. L. Kimura and A. A. F. Loureiro, "MPTCP Linux kernel congestion controls," 2018, *arXiv:1812.03210*.

[24] L. Khalili. *Congestion Control of Multipath TCP: Problems and Solutions*. Accessed: Sep. 22, 2023. [Online]. Available: https://www.ietf.org/proceedings/87/slides/slides-87-iccrg-7.pdf

[25] P. Hurtig, K.-J. Grinnemo, A. Brunstrom, S. Ferlin, Ö. Alay, and N. Kuhn, "Low-latency scheduling in MPTCP," *IEEE/ACM Trans. Netw.*, vol. 27, no. 1, pp. 302–315, Feb. 2019.

[26] C. Paasch, S. Ferlin, O. Alay, and O. Bonaventure, "Experimental evaluation of multipath TCP schedulers," in *Proc. ACM SIGCOMM Workshop Capacity Sharing Workshop (CSWS)*. New York, NY, USA: Association for Computing Machinery, Aug. 2014, pp. 27–32, doi: 10.1145/2630088.2631977.

[27] S. Li, Q. Chen, Z. Li, W. Meng, and C. Li, "Civil aircraft assisted space-air-ground integrated networks: Architecture design and coverage analysis," *China Commun.*, vol. 19, no. 1, pp. 29–39, Jan. 2022.

[28] L. Wei, J. Shuai, Y. Liu, Y. Wang, and L. Zhang, "Service customized space-air-ground integrated network for immersive media: Architecture, key technologies, and prospects," *China Commun.*, vol. 19, no. 1, pp. 1–13, Jan. 2022.

[29] M. H. Eiza and A. Raschellà, "A hybrid SDN-based architecture for secure and QoS aware routing in space-air-ground integrated networks (SAGINs)," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2023, pp. 1–6.

[30] A. Burkitt-Gray. (2022). *T-Mobile and SpaceX Working on Messaging-From-Space Plan*. Accessed: Aug. 8, 2023. [Online]. Available: https://www.capacitymedia.com/article/2ajibesk7e5xljax6f1mo/news/t-mobile-and-spacex-working-on-messaging-from-space-plan

[31] A. Burkitt-Gray. (2022). *T-Mobile and SpaceX Working on Messaging-From-Space Plan*. Accessed: Aug. 8, 2023. [Online]. Available: https://www.capacitymedia.com/article/2ajibesk7e5xljax6f1mo/news/t-mobile-and-spacex-working-on-messaging-from-space-plan

[32] Apple. (2023). *Use Emergency SOS Via Satellite on Your iPhone 14*. Accessed: Aug. 8, 2023. [Online]. Available: https://support.apple.com/en-us/HT213426

[33] F. Royal, R. Hartani, and A. Mendes. (2020). *Defining the Synergies Between LEO Satellite Constellations and Submarine Cables*. Accessed: Aug. 8, 2023. [Online]. Available: https://xonapartners.com/wp-content/uploads/2020/12/Defining-the-Synergies-between-LEO-Satellite-Constellations-and-Submarine-Cables.pdf

[34] P. Austria, C. H. Park, J.-Y. Jo, Y. Kim, R. Sundaresan, and K. Pham, "BBR congestion control analysis with multipath TCP (MPTCP) and asymmetrical latency subflow," in *Proc. IEEE 12th Annu. Comput. Commun. Workshop Conf. (CCWC)*, Jan. 2022, pp. 1065–1069.

[35] C. Yu, W. Quan, N. Cheng, S. Chen, and H. Zhang, "Coupled or uncoupled? Multi-path TCP congestion control for high-speed railway networks," in *Proc. IEEE/CIC Int. Conf. Commun. China (ICCC)*, Aug. 2019, pp. 612–617.

[36] K. Nihei, N. Kai, Y. Maruyama, T. Yamashita, D. Kanetomo, T. Kitahara, M. Maruyama, T. Ohki, K. Kusin, and H. Segah, "Forest fire surveillance using live video streaming from UAV via multiple LTE networks," in *Proc. IEEE 19th Annu. Consum. Commun. Netw. Conf. (CCNC)*, Jan. 2022, pp. 465–468.

[37] S. Baidya, Z. Shaikh, and M. Levorato, "FlyNetSim: An open source synchronized UAV network simulator based on ns-3 and ardupilot," in *Proc. 21st ACM Int. Conf. Model., Anal. Simul. Wireless Mobile Syst. (MSWIM)*. New York, NY, USA: Association for Computing Machinery, Oct. 2018, pp. 37–45, doi: 10.1145/3242102.3242118.

[38] T. Melodia, S. Basagni, K. R. Chowdhury, A. Gosain, M. Polese, P. Johari, and L. Bonati, "Colosseum, the world's largest wireless network emulator," in *Proc. 27th Annu. Int. Conf. Mobile Comput. Netw. (MobiCom)*. New York, NY, USA: Association for Computing Machinery, Oct. 2021, pp. 860–861, doi: 10.1145/3447993.3488032.

[39] A. Kvalbein, D. Baltrunas, K. Evensen, J. Xiang, A. Elmokashfi, and S. Ferlin-Oliveira, "The Nornet edge platform for mobile broadband measurements," *Comput. Netw.*, vol. 61, pp. 88–101, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128613004490

[40] A. Baltaci, F. Geyer, M. Klügel, and D. Schupke, "Demo: A combined multipath connectivity and flight simulation framework," in *Proc. 25th ACM Int. Conf. Modeling, Anal. Simulation Wireless Mobile Syst. (MSWIM)*. New York, NY, USA: Association for Computing Machinery, 2022, pp. 17–21, doi: 10.1145/3551662.3560925.

[41] R. M. N. Chirwa and A. P. Lauf, "Performance improvement of transmission in unmanned aerial systems using multipath TCP," in *Proc. IEEE Int. Symp. Signal Process. Inf. Technol. (ISSPIT)*, Dec. 2014, pp. 000019–000024.

[42] L. Stratmann, B. Walker, and V. A. Vu, "Realistic emulation of LTE with MoonGen and DPDK," in *Proc. WiNTECH*. New York, NY, USA: Association for Computing Machinery, 2020, pp. 87–94, doi: 10.1145/3411276.3412183.

[43] OpenSAND. (2022). *OpenSAND Provides an Easy and Flexible Way to Emulate an End-to-End Satellite Communication System*. Accessed: Sep. 25, 2022. [Online]. Available: https://www.opensand.org/index.html

[44] M. Kosek, H. Cech, V. Bajpai, and J. Ott, "Exploring proxying QUIC and HTTP/3 for satellite communication," in *Proc. IFIP Netw. Conf. (IFIP Networking)*, Jun. 2022, pp. 1–9.

[45] A. Auger, E. Lochin, and N. Kuhn, "Making trustable satellite experiments: An application to a VoIP scenario," in *Proc. IEEE 89th Veh. Technol. Conf. (VTC-Spring)*, Apr. 2019, pp. 1–5.

[46] V. A. Vu and M. Akselrod, "An experiment of dual-LTE MPTCP with in-car voice assistant," in *Proc. IEEE 93rd Veh. Technol. Conf. (VTC-Spring)*, Apr. 2021, pp. 1–5.

[47] V. A. Vu and J. Wolff, "Supporting delay-sensitive applications with multipath QUIC and forward erasure correction," in *Proc. 17th ACM Symp. QoS Secur. Wireless Mobile Netw. (Q2SWinet)*. New York, NY, USA: Association for Computing Machinery, Nov. 2021, pp. 95–103, doi: 10.1145/3479242.3487312.

[48] OpenSAND. (2023). *OpenSAND Publications*. Accessed: Jun. 17, 2023. [Online]. Available: https://www.opensand.org/references.html

[49] Netem. *NetEm—Network Emulator*. Accessed: Sep. 25, 2022. [Online]. Available: https://www.linux.org/docs/man8/tc-netem.html

[50] P. Emmerich, S. Gallenmüller, D. Raumer, F. Wohlfart, and G. Carle, "MoonGen: A scriptable high-speed packet generator," in *Proc. Internet Meas. Conf. (IMC)*, Tokyo, Japan, Oct. 2015, pp. 275–287.

[51] DPDK. (2020). *Data Plane Development Kit—Programmer's Guide*. Accessed: Jul. 9, 2023. [Online]. Available: https://fast.dpdk.org/doc/pdf-guides-17.11/prog_guide-17.11.pdf

[52] MPTCP. (2023). *Packet Scheduler, Path Manager and Congestion Control*. Accessed: Aug. 9, 2023. [Online]. Available: https://github.com/multipath-tcp/mptcp_net-next/issues/418

[53] MPTCP. (2023). *Packet Scheduler, Path Manager and Congestion Control*. Accessed: Aug. 9, 2023. [Online]. Available: https://github.com/multipath-tcp/mptcp_net-next/issues/300

[54] F. Aschenbrenner, T. Shreedhar, O. Gasser, N. Mohan, and J. Ott, "From single lane to highways: Analyzing the adoption of multipath TCP in the Internet," in *Proc. IFIP Netw. Conf. (IFIP Networking)*, Jun. 2021, pp. 1–9.

[55] Y.-S. Lim, E. M. Nahum, D. Towsley, and R. J. Gibbens, "ECF: An MPTCP path scheduler to manage heterogeneous paths," in *Proc. 13th Int. Conf. Emerg. Netw. EXperiments Technol. (CoNEXT)*. New York, NY, USA: Association for Computing Machinery, Nov. 2017, pp. 147–159, doi: 10.1145/3143361.3143376.

[56] M. Amend, E. Bogenfeld, M. Cvjetkovic, V. Rakocevic, M. Pieska, A. Kassler, and A. Brunstrom, "A framework for multiaccess support for unreliable Internet traffic using multipath DCCP," in *Proc. IEEE 44th Conf. Local Comput. Netw. (LCN)*, Los Alamitos, CA, USA, Oct. 2019, pp. 316–323, doi: 10.1109/LCN44214.2019.8990746.

[57] K. Chavali. (2023). Accessed: Aug. 11, 2023. *LTE and SATCOM Joint Emulator*. [Online]. Available: https://github.com/KaushikChavali/cellular-satcom-emulator

[58] Marin, B. Michau, and C. Shawn. (2022). *Qcsuper*. Accessed: Jul. 1, 2023. [Online]. Available: https://github.com/P1sec/QCSuper

[59] OpenSAND. (2022). *Blog—Latest Announcements*. Accessed: Jun. 25, 2023. [Online]. Available: https://www.opensand.org/blog.html

[60] F. Michel, M. Trevisan, D. Giordano, and O. Bonaventure, "A first look at starlink performance," in *Proc. 22nd ACM Internet Meas. Conf. (IMC)*. New York, NY, USA: Association for Computing Machinery, Oct. 2022, pp. 130–136, doi: 10.1145/3517745.3561416.

[61] Y. Help. (2022). *System Requirements*. Accessed: Jul. 1, 2022. [Online]. Available: https://support.google.com/youtube/answer/78358?hl=en

[62] D. Telekom. (2022). *Tools—iPerf3 With MP-DCCP/DCCP Support*. Accessed: Jun. 11, 2023. [Online]. Available: https://multipath-dccp.org/tools.html

[63] *How to Estimate the Maximum and Recommended Flight Times of a UAS, UAV, or Drone System*, RRPLSUA Center, Louisiana State Univ. Agricult. Center, USA, 2022. Accessed: Oct. 13, 2022. [Online]. Available: https://www.lsuagcenter.com/NR/rdonlyres/24FEF5D5-FD37-4DF2-8323-974243633CC7/105242/3469Drones.pdf

[64] Y. Cao, M. Xu, and X. Fu, "Delay-based congestion control for multipath TCP," in *Proc. 20th IEEE Int. Conf. Netw. Protocols (ICNP)*, Oct. 2012, pp. 1–10.

[65] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[66] M. Pracucci. (2018). *Linux TCP RTO MIN, TCP RTO MAX and the TCP Retries2 Sysctl*. Accessed: Jul. 11, 2023. [Online]. Available: https://pracucci.com/linux-tcp-rto-min-max-and-tcp-retries2.html

[67] L. Li, K. Xu, D. Wang, C. Peng, Q. Xiao, and R. Mijumbi, "A measurement study on TCP behaviors in HSPA+ networks on high-speed rails," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2015, pp. 2731–2739.

[68] M. A. K. Fard, K. A. Bakar, S. Karamizadeh, and R. H. Foladizadeh, "Improve TCP performance over mobile ad hoc network by retransmission timeout adjustment," in *Proc. IEEE 3rd Int. Conf. Commun. Softw. Netw.*, May 2011, pp. 437–441.

[69] D. Mindrila and P. Balentyne. (2013). *Scatterplots and Correlation*. Accessed: Jun. 18, 2023. [Online]. Available: https://www.westga.edu/academics/research/vrc/assets/docs/scatterplots_and_correlation_notes.pdf

[70] *Evolved Universal Terrestrial Radio Access (E-Utra); Study on Latency Reduction Techniques for LTE (Release 14)*, document TR 36.881, 3GPP, 2016. Accessed: Oct. 27, 2021. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/36_series/36.881/

[71] K. Chavali, "Feasibility analysis of multipath TCP for the connectivity of remote piloting operations of drones and flying taxis," M.S. thesis, Dept. Inform., Tech. Univ. Munich, Germany, 2022.

[72] J. A. Hartigan and P. M. Hartigan, "The dip test of unimodality," *Ann. Statist.*, vol. 13, no. 1, pp. 70–84, Mar. 1985. [Online]. Available: http://www.jstor.org/stable/2241144

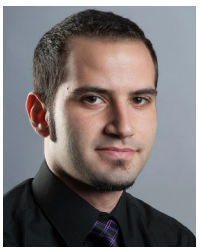

AYGÜN BALTACI received the bachelor's degree in electrical engineering from the University of North Florida, in 2016, and the M.Sc. degree in communications engineering from the Technical University of Munich (TUM), Germany, in 2018, where he is currently pursuing the Ph.D. degree. He is a Research Scientist with Airbus Central Research and Technology. His research interests include heterogeneous networking for reliable communications, multipath transport, and reliable and robust networking in aerospace networks.

KAUSHIK CHAVALI received the B.Eng. degree in computer engineering from Gujarat Technological University, India, in 2015, and the M.Sc. degree in computer science from the Technical University of Munich, Germany, in 2022. He was a Junior Research Fellow with the Aeronautical Development Agency, Bengaluru, from 2017 to 2018. His research interests include next-generation wireless networks, multipath transport protocols, and internet measurements.

MIKE KOSEK is currently a Researcher with the Technical University of Munich (TUM), Germany. His research interests include internet architecture in general, and transport protocol standardization, development, and deployment, in particular.

NITINDER MOHAN received the M.Tech. degree (Hons.) from the Indraprastha Institute of Information Technology Delhi (IIIT-D), India, in 2015, and the Ph.D. degree (as a Marie Curie ITN Fellow) from the Department of Computer Science, University of Helsinki, Finland, in 2019. He is currently a Postdoctoral Researcher with the Chair of Connected Mobility, Technical University of Munich, Germany. His research interests include edge computing, next-generation networked applications, and large-scale internet measurements. He has been awarded the Outstanding Ph.D. Dissertation Award by the IEEE Technical Committee on Scalable Computing (TCSC).

DOMINIC A. SCHUPKE (Senior Member, IEEE) received the Dr.-Ing. degree (summa cum laude) in electrical engineering and information technology from RWTH Aachen, Imperial College London, and the Technical University of Munich (TUM). He is currently a Research Leader in reliable communication networks, with a focus on wireless communications with Airbus Central Research and Technology, Munich, Germany. He is also a Lecturer in network planning with the TUM. Prior to Airbus Central Research and Technology, he was with Nokia, Siemens, and TUM. He is the author or coauthor of more than 150 journal articles and conference papers (Google Scholar H-index 34). His research interest includes aerospace networks.

JÖRG OTT has been the Chair of Connected Mobility, School of Computation, Information and Technology, Technical University of Munich, since 2015. He explores edge and in-network computing as well as mobile and decentralized (cloudless) services, with an emphasis on privacy by design. His research interests include network architecture, (transport) protocol design, and algorithms for connecting mobile nodes to the internet and to each other, covering the spectrum from delay-tolerant to real-time networking.