This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

Evaluation of MTIP, MPTCP and MPQUIC for a Time-Sensitive Application

Delia Rico*, Francisco Luque-Schempp* and Pedro Merino*
*ITIS Software, University of Malaga, Spain. Email: {deliarico,schempp,pmerino}@uma.es

***Abstract*—This paper compares the behavior of the Multi-connection Tactile Internet Protocol (MTIP) with two other well-known multipath protocols, Multipath TCP (MPTCP) and Multipath QUIC. The evaluation focuses on the effectiveness of the protocols in scenarios involving the remote control of a quadruped robot during emergency rescue operations. The study aims to determine how these protocols perform under various network conditions, including stable connections, link failures, and typical 5G network impairments. MTIP leverages network programmability to dynamically adapt to changing conditions, maintaining low latency without compromising the expected reliability for time-sensitive applications. MTIP dynamic path management provides significantly better results for the target use case compared to the other protocols, especially in adverse environments. This paper contributes to understanding how multipath communication can enhance the execution of time-sensitive applications.**

***Index Terms*—Multipath Protocols, Time-Sensitive Applications, MPTCP, Multipath QUIC, MTIP, 5G, Tactile Internet**

I. INTRODUCTION

Time-sensitive applications encompass critical environments where even small delays can significantly affect operation or user experience. This term is commonly applied to Time-Sensitive Networking (TSN) or real-time protocols. However, it extends to various applications that require low-latency and reliable communication, regardless of the underlying technology. These applications have become increasingly critical in several industries, from healthcare to industrial automation, relying on connectivity that allows accurate and timely responses [1]. In particular, these characteristics are especially vital for time-critical situations and subsequently offer substantial advantages in terms of safety and operational efficiency. For instance, the remote operation of robots or drones during emergency operations, where human presence is either impractical or dangerous. In addition, the concept of the Tactile Internet goes further by introducing real-time haptic feedback, including not only traditional traffic transmission but also tactile and control signals with very demanding requirements. In this context, 5G Ultra-Reliable Low Latency Communications (URLLC) emerges as a key enabler, meeting the stringent requirements demanded for such applications.

Multipath transport protocols are also a promising solution to improve the reliability and latency of communications in these applications [2]. The use of multiple network paths simultaneously allows protocols to introduce redundancy and increase data throughput, minimizing disruptions in environments where connection failures could be fatal, and addressing connectivity needs even under variable network conditions. Therefore, in combination with the advanced programmability of 5G [3] and future 6G networks, multipath protocols enable greater flexibility and robustness, enhancing network performance. Multipath TCP (MPTCP) [4] and Multipath QUIC (MPQUIC) [5] are positioned as alternatives capable of leveraging various paths to enable continuous communications, which makes them particularly attractive for demanding applications. For even more challenging applications that demand ultra-low latency and high reliability, where the Tactile Internet is proposed as a key enabler of real-time and haptic feedback, the Multi-connection Tactile Internet Protocol (MTIP) has been developed [6]. MTIP introduces dynamic route management and continuously adapts to network conditions, providing the robustness and flexibility needed for this purpose.

The aim of this paper is to assess the use of multipath transport protocols for time-sensitive applications, specifically the remote control of a quadruped Unmanned Ground Vehicle (UGV) in scenarios corresponding to rescue conditions. The implementation details of each protocol (MPTCP, MPQUIC and MTIP), the testbed setup, and key metrics are discussed. In addition, the evaluation examines how each protocol handles varying network conditions in different scenarios, including stable connections, link failures, and dynamic environments (i.e., latency fluctuations and congestion). Although MPTCP and MPQUIC are the most advanced and widely adopted multipath transport protocols for most Internet traffic, their strong performance does not necessarily extend to all types of traffic. In particular, they may fail to fully meet the strict latency deadlines and reliability constraints of time-sensitive applications. This paper presents the first evaluation comparing these general-purpose multipath protocols with a multipath protocol specifically designed for time-critical scenarios (i.e., MTIP). The results highlight the advantages and limitations of these multipath protocols in the context of such applications, where traditional communication protocols like TCP and UDP, although foundational for Internet communication, fail to meet the stringent demands.

The rest of the article is organized as follows. Section II provides background on time-sensitive applications and multipath transport protocols. Section III details the evaluation description, including the use case, testbed setup, scenarios, and evaluation metrics. Then, Section IV presents the evaluation results for various conditions, such as stable links, link failure, and additional 5G network impairments. In Section 5, we discuss the findings, focusing on latency, reliability, and resource utilization. Finally, Section VI concludes summarizing key insights and potential implications for future research.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

II. BACKGROUND

A. Time-Sensitive Applications

The growing need for the use of time-sensitive applications in critical sectors has increased the demand for ultra-low latency and high-reliability communication networks. These kinds of applications (e.g., industrial automation, telemedicine, and collaborative robotics) require flexible network architectures capable of delivering rapid response times and high availability. For example, minor delays or connection losses can have serious consequences in an emergency situation such as a natural disaster or a rescue mission, affecting both safety and operational performance. Thus, meeting these strict requirements is crucial.

Within this scope, the Tactile Internet introduces a shift in the way humans interact with machines and digital environments, emerging as a fundamental enabler for applications where real-time control and haptic feedback are essential. The IEEE standards position Tactile Internet applications [1] a step beyond traditional data transmission, incorporating stringent requirements for ultra-low latency and high reliability to support functionalities such as remote surgery, autonomous driving, and collaborative robotics. Subsequently, the imposed traffic requirements require advanced communication networks capable of maintaining such communications, that is, resilient and flexible enough to support them in various situations.

In addition, 5G URLLC is specifically designed to support mission-critical communications and plays a crucial role by providing the low latency, ultra-high reliability, and deterministic behavior required for time-sensitive applications. On the one hand, by leveraging advanced techniques such as network slicing, edge computing, and enhanced radio resource management, 5G URLLC ensures seamless and reliable connectivity for Tactile Internet applications. On the other hand, programmability and reconfigurability in 5G and the forthcoming 6G networks are key to optimizing resource utilization and traffic prioritization, aligning network performance with the unique demands of various applications. These features bring high levels of control and flexibility, enabling networks to dynamically adjust based on application-specific requirements. The integration of programmability into transport protocols is a critical factor in this paradigm shift. It allows for the maintenance of low latency and high reliability, which are crucial for real-time applications.

B. Multipath Transport Protocols

Traditional transport protocols are not designed to meet stringent requirements [7] and new emerging challenges [8]. On the one hand, TCP focuses on the reliability of the delivery through retransmissions, which can introduce significant delays in the communication. UDP, on the other hand, provides low-latency communication but lacks reliability mechanisms, making it unsuitable for applications where data loss cannot be tolerated. Consequently, multipath communication protocols have emerged as key technologies, which enable the simultaneous use of multiple network paths, thereby enhancing redundancy and optimizing data flow under adverse network conditions to improve both reliability and performance.

First, Multipath TCP (MPTCP) [4] is one of the most studied and implemented multipath protocols. It extends traditional TCP by allowing multiple subflows within a single connection, enabling data to be transmitted over multiple paths. This improves fault tolerance and can enhance throughput, especially in heterogeneous network environments. Second, Multipath QUIC (MPQUIC) [5] is a prominent protocol that extends the QUIC protocol to support multiple paths. QUIC, originally developed by Google, is a modern transport protocol designed for low-latency internet communication. It integrates features such as encryption, congestion control, and connection migration, making it well suited for real-time applications. Multipath QUIC builds on these features by allowing the use of multiple network paths and, therefore, enhancing performance and reliability. The multi-connectivity capability enables applications to maintain stable connections by automatically reallocating traffic over available paths in case of failure or congestion. This provides crucial benefits such as reducing the impact of network variability on latency and improving availability, enabling continuous communication in situations where network stability may be compromised. Recent efforts such as MPR-QUIC [9] have improved video streaming over MPQUIC through partially reliable transmission and priority-based, deadline-aware schedulers. However, video traffic optimizations do not necessarily meet the low-latency and high-reliability needs of time-sensitive applications. Therefore, despite the advances made by MPTCP, which pioneered the use of multiple paths at the transport layer, and Multipath QUIC, which enhances low-latency communication by avoiding strict acknowledgment mechanisms, there is still a need for protocols specifically tailored to transport deadline constrained reliable low-latency data [2], as required by these applications.

In this context, the Multi-connection Tactile Internet Protocol (MTIP) has been developed [6], [10] to provide a flexible and adaptable communication solution that can meet the ultra-low latency and high reliability requirements of applications such as remote robot control. MTIP is described in detail in our previous work [6], but its key aspects are summarized here. MTIP establishes a link between two endpoints using multiple sublinks, where each sublink represents an independent end-to-end connection, such as a 5G connections. These sublinks are dynamically managed to enhance the operation, ensuring that the requirements of time-sensitive applications are met. Rather than relying on strict acknowledgment-based communication, MTIP prioritizes latency and deadline adherence while improving reliability through packet duplication across sublinks. However, it avoids using duplication when it is not necessary, to minimize resource consumption, such as energy, which is particularly critical for battery-powered devices.

To make intelligent sublink selection decisions, MTIP gathers and processes contextual information from two primary sources. On one side, it provides an Application Programming Interface (API) that allows applications to specify key communication parameters, such as the maximum allowable one-way latency and the importance of resource efficiency. This API enables applications to fine-tune their communication preferences and even override MTIP's automated selection process by directly specifying the number of active sublinks.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

On the other side, MTIP continuously monitors network conditions using both data packets and dedicated control packets, which are used to perform continuous network measurements and synchronize network state between endpoints, ensuring an accurate and up-to-date view of the available connections. Leveraging this real-time information, MTIP dynamically selects sublinks for transmission. The sending algorithm continuously evaluates both application input and network conditions to determine which and how many sublinks should be used to transmit the data. On the receiving end, the algorithm processes incoming packets to comply with the application-defined deadline. Packets that exceed the deadline are discarded, as are redundant duplicates of previously received packets. To support this process efficiently while keeping overhead low, MTIP employs a compact header structure based on a UDP-like format with two additional fields: sequence numbers, which allow the identification and elimination of duplicate packets, and timestamps, which enable time and latency tracking. MTIP continuously updates and synchronizes its sending algorithm and sublink selection tables to adapt to changing network conditions. This ensures the timely and reliable delivery of time-sensitive traffic while maintaining efficient resource utilization.

Evaluating the performance of multipath protocols [11], [12] in demanding environments requires a variety of tests covering different network conditions. These conditions must simulate the extreme scenarios that devices may face in realistic environments and provide insight into the robustness and adaptability of protocols. In this paper, we focus on the remote control of a quadruped robot in simulated rescue operations, a scenario that demands real-time responsiveness, high reliability, and adaptability to changing network conditions. Using the Victoria Network at the University of Malaga [13], we analyze the efficiency of multipath transport protocols for time-sensitive applications. In particular, the Multi-connection Tactile Internet Protocol (MTIP) compared to Multipath TCP (MPTCP) and Multipath QUIC (MPQUIC). This evaluation not only assesses the ability of each protocol to support time-sensitive applications, but also explores how network programmability can enhance deterministic multi-connectivity solutions, offering practical insights into their suitability for real-time remote control in critical applications.

III. EVALUATION APPROACH

A. Use case

The use case under consideration is framed within the context of the 5G+TACTILE project, which aims to enable low-latency, high-reliability communication in a live 5G network environment. This use case involves the remote control of a quadruped UGV in emergency rescue operations, where real-time responsiveness, consistent reliability, and the ability to adapt to fluctuating network conditions are essential. Specifically, it employs the Ghost Vision 60 robot[1], which features precise control mechanisms, including variable speed (up to 3 m/s), long-range capabilities (up to 10 km), considerable power runtime (maximum of 3 hours), and obstacle avoidance, enabling rapid adaptation to new environments and navigation of complex terrain. These capabilities allow the robot to quickly adapt to dynamic environments and traverse challenging terrains, making it an ideal candidate for emergency scenarios. However, to support this, the underlying network must be able to meet the required performance levels. One key enabler considered in the project is multi-connectivity, as it enhances reliability by enabling simultaneous connections across multiple network paths, thus reducing the risk of packet loss or delays due to issues with a single network path.

Regarding the specifics of the application and its traffic pattern, it closely resembles a typical remote control application, where commands such as motion control consist of small packet sizes (less than 200 bytes) transmitted every 30 milliseconds (ms). The application has a critical constraint of maintaining a maximum one-way end-to-end delay of 15 ms, making delayed packets more detrimental to the system than packet loss [1]. However, the system also poses a high reliability constraint of 99% for correctly received orders, further emphasizing the need for flexible, robust, and resilient communication protocols. In this context, we consider two well-known multipath protocols, MPTCP and MPQUIC, alongside the newly proposed MTIP protocol, specifically designed for applications of this nature. The testbed, scenario, and metrics used to evaluate these protocols are presented in the following subsections.

B. Testbed setup

The testbed depicted in Figure 1 is set to evaluate the three protocols: MPTCP, MPQUIC, and MTIP. The setup consists of a computer acting as a remote controller that sends orders to a computer connected to the quadruped robot. These two computers are interconnected by two independent 5G paths. The 5G paths are established using two Nokia AirScale remote radio heads, one Micro (AWHQE) which operates in band 78 and one Pico (AWHQE) which operates at a different frequency in band 77, both in standalone mode (SA). There are two differentiated 5G network cores, one Polaris 5G and one Open5GS core. In addition, each core is configured with a dedicated slice, ensuring efficient and isolated data transmission that can exhibit different network conditions at any given time. This testbed is part of the Victoria Network at the University of Malaga [13], a private 5G network focused on testing and advancing toward 6G technologies. The network ensures precise clock synchronization between endpoints using dedicated network interfaces and GPS to provide a reference signal to a Precision Time Protocol (PTP) grandmaster clock, the ADVA FSP 150 EMS.

In the controller, a proxy takes remote control orders and encapsulates them in the corresponding protocol (MPTCP, MPQUIC, or MTIP) before sending them through the multiple paths. On the robot side, the orders are decapsulated and forwarded to the robot in their original format. The iteration between the proxies offers a programmable solution to evaluate the functioning and adaptation of each protocol under different network conditions for the remote control application under study.

---

[1] Ghost Vision 60: https://www.ghostrobotics.io/vision-60

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

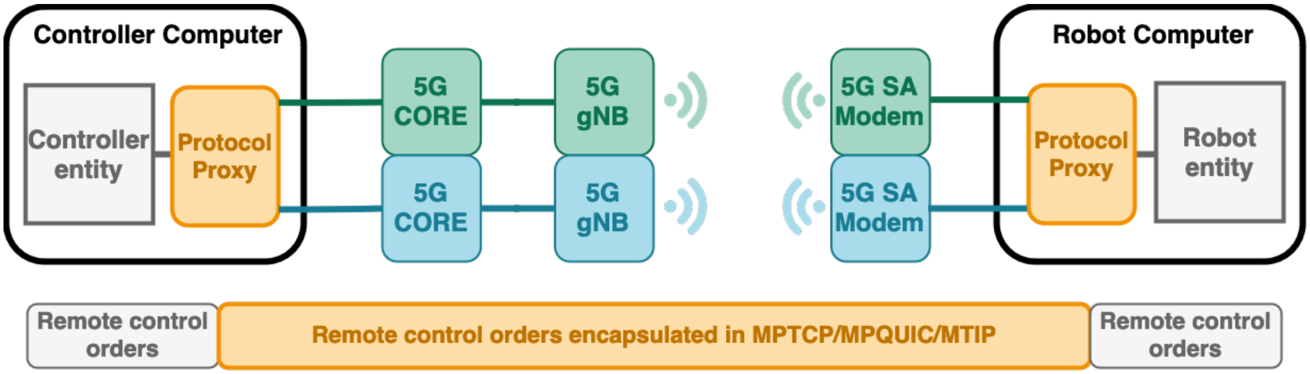

Fig. 1: Diagram of the testbed setup.

The specific implementation of the protocols used for the proxies is the Multipath TCP implementation included in the Linux OS kernel[2], the Multipath QUIC implementation created by the University of Louvain[3] and the Multi-connection Tactile Internet Protocol accessible on GitHub[4].

Regarding MTIP, its programmability enables the evaluation of two configurations: MTIP-A and MTIP-B. MTIP-A is configured to override the MTIP sending algorithm and consistently uses both paths to transmit duplicate data. In contrast, MTIP-B leverages the MTIP sending algorithm to dynamically adjust the number of active paths during communication, aiming to optimize network resource utilization and reduce energy consumption at the endpoints, albeit with some trade-offs in performance. MTIP-B follows the data exchange mechanism detailed by Rico et al. [6], [10]. In summary, it employs a rapid adaptation strategy that constantly re-evaluates path selection. If too many packets are lost within the defined observation period the number of active paths is increased; conversely, if too many packets are duplicated, the number of paths is reduced. This adaptive method is particularly beneficial in energy-constrained scenarios, such as battery-powered robotic systems or networks shared with other devices. MTIP configurations are selected to adhere to an end-to-end latency constraint of 15 ms. If a packet exceeds this latency threshold, it is discarded, prompting a re-evaluation of the MTIP sending algorithm to ensure compliance with the latency requirement. Note that MTIP-A is a configuration designed to achieve higher performance without focusing on efficient use of resources. This configuration can be used as a baseline for comparison with other configurations that make better use of available resources.

C. Scenarios

The evaluation focuses on four scenarios: a) stable links with no failures, b) one fatal link failure, c) delay and jitter impairments in the environment, and d) congestion incidents. The purpose is to study the adaptability of the protocols under different conditions. For scenario a) no special configuration is needed, but the normal functioning of both the 5G connections, which are generally reliable and error-free. In scenario b) there is a link failure configured in the middle of the connection such that alters the main link being used for the communication with a packet loss of 100%, to evaluate how the protocols reconfigure and adapt to this path loss. For scenario c), we have used our previous knowledge of the behavior of the 5G network examined in [14] to replicate a common traffic pattern of latency and jitter variation in 5G networks. Specifically, we have studied how buffers and timers operate in the network with the aim of replicating one of these patterns on both links. All different impairments are added to the topology using the Netem[5] network emulator, a Linux-based tool that enables controlled management of network performance (for example, degradation). It allows different parameters to be fine-tuned to replicate the characteristics of a real-world 5G network. We developed a script to apply the corresponding Netem rules to network interfaces concurrently with the scenario. Finally, scenario d) examines how congestion incidents in a path impact the protocols. These incidents are introduced in two ways. First, as in scenario c), we leverage prior knowledge of 5G network behavior and use the Netem tool to generate packet losses that reflect typical congestion patterns. Second, we introduce background video feedback, which acts as a typical source of congestion in the targeted use case. To generate this traffic, we use the Keysight IxChariot tool[6], which is a comprehensive network testing solution that emulates real-world application traffic, providing realistic data flows through the network. Specifically, we generate a video streaming flow over the Real-time Transport Protocol (RTP), encoded with H.264/AVC, with a bit rate of 45 Megabytes (representing, for instance, a 360 camera feed).

Note that all tests have been repeated a minimum of 25 times based on our experience in previous projects to obtain reliable and statistically significant KPI validation [15]. The exact latency and jitter levels used for the impairments in scenario c), as well as the detailed traffic characteristics of scenario d) are available on GitHub[4].

---

[2] MPTCP implementation: https://www.mptcp.dev/

[3] MPQUIC implementation: https://multipath-quic.org/

[4] MTIP repository – contains all materials related to the evaluation presented in this paper, including scenarios details, results logs, and Wireshark traces: https://github.com/deliarico/MTIP

[5] Netem: https://wiki.linuxfoundation.org/networking/netem

[6] IxChariot: https://keysight.com/se/en/products/network-test/ixchariot

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

D. *Metrics*

For the evaluation of the scenarios, we focus on two primary Key Performance Indicators (KPIs): one-way end-to-end latency and reliability (as the percentage of correct received orders). To calculate these KPIs accurately, packets have been timestamped at the application level, allowing for precise measurement of delays and reliability across the communication process. In addition to these metrics, we also analyze the utilization of different paths to assess how each protocol manages the use of available network resources.

IV. RESULTS

In this section, we present the results obtained from the evaluation, organized by the different scenarios. The detailed logs, results and additional figures for each iteration are available on GitHub for further review[4].

A. *Stable links with no failures*

In Figure 2 are shown the results of the different iterations with each protocol for scenario a). In general, we observe that there are negligible losses in all iterations with all protocols since the network is stable and can support the amount of packets being sent. Regarding latency, the network predominantly supports the 15-millisecond end-to-end delay limitation. However, MPTCP and MPQUIC do not include specific mechanisms to accommodate this type of application requirement and do not take action to address it. As a result, the application must handle the discarding of traffic that exceeds this threshold, without the ability to recover from it due to the latency constraint. In the case of MTIP, using both paths simultaneously (MTIP-A) provides the necessary redundancy to offer the best operation in terms of both latency and packet loss. This configuration could be recommended if resource consumption is not a critical concern. However, when resource waste is a concern, the MTIP-B configuration proves to be the most beneficial, triggering a reselection of the active paths when too many packets are late or lost (a further study of the path utilization can be found in Section V). This comes at the cost of some packet losses (0.08%), since packets that exceed the threshold are being discarded. However, this trade-off is acceptable in the requirements of the given use case (which tolerates some losses, as it is preferable to exceeding the latency deadline).

B. *Link failure*

Figure 3 illustrates the case of a fatal link failure. For this scenario, an additional upper section has been added to the MPTCP and MPQUIC graphs to highlight some latency peaks observed in their behavior. MPTCP and MPQUIC can recover from a fatal link failure, at the cost of a significant impact on the delay. MPTCP recovers with 0% packet loss, as its TCP mechanisms prioritize 100% reliability. However, this comes at the cost of latency peaks. Analyzing the individual executions, we observe a significant latency peak when the link failure occurs, followed by latency recovery in the subsequent seconds. For MPQUIC, two main behaviors are observed. In the first behavior (e.g. iterations 6, 9 and 12), the studied implementation of MPQUIC sometimes fails to recover from the

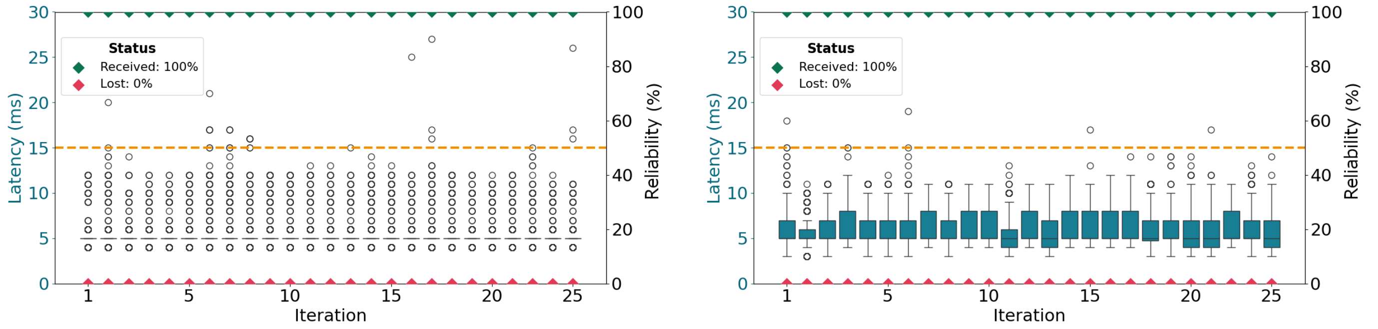

(a) MPTCP.

(b) MPQUIC.

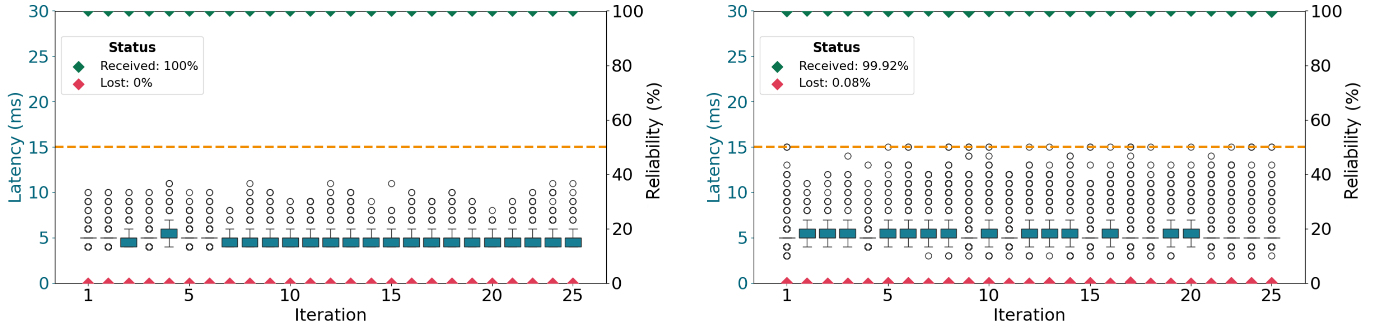

(c) MTIP-A.

(d) MTIP-B.

Fig. 2: Latency and reliability of scenario a) Stable links.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

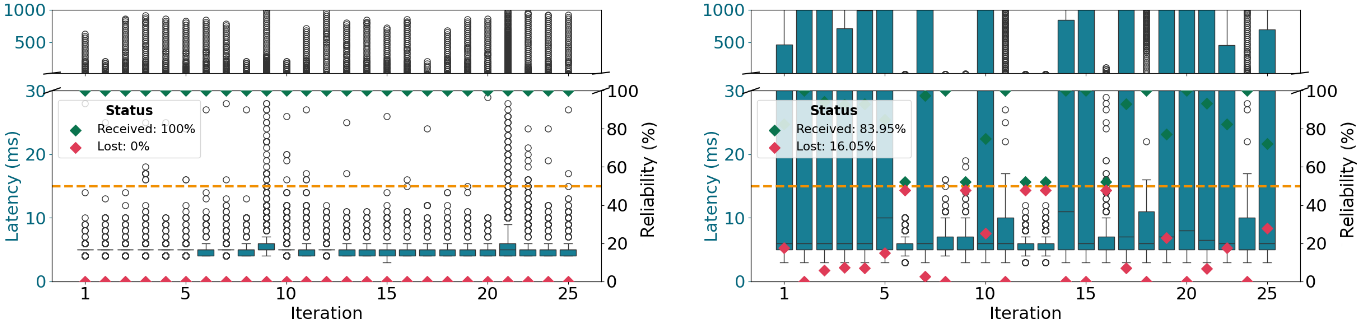

(a) MPTCP.

(b) MPQUIC.

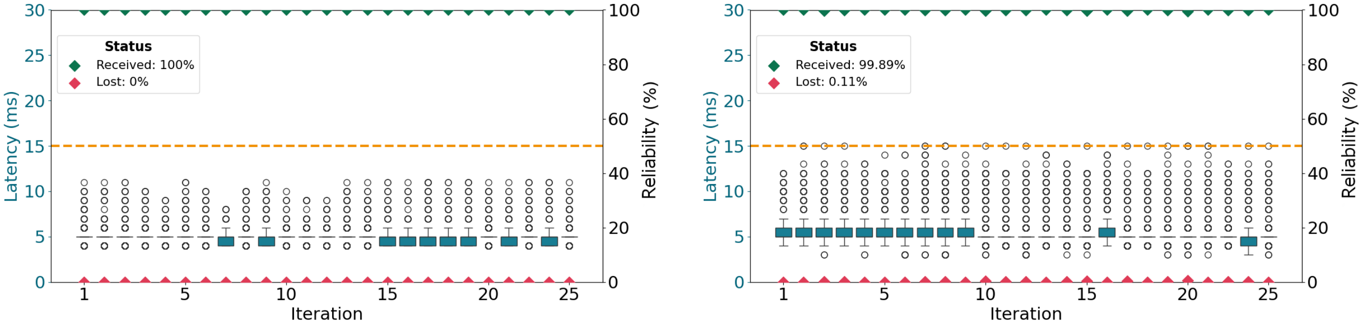

(c) MTIP-A.

(d) MTIP-B.

Fig. 3: Latency and reliability of scenario b) Link failure.

link failure, being unable to reroute the traffic to the other path. In the second behavior, observed in most of the remaining iterations, MPQUIC does recover, but it often experiences high delays, affecting overall communication. MTIP adheres to the configured 15-millisecond constraint. In MTIP-A, no losses are observed because both paths are used simultaneously, ensuring that if one path fails, the other continues to function as expected. In MTIP-B, a trade-off between recovery and packet loss is observed, with an average loss rate of 0.11%. This loss occurs during recovery time when the protocol detects packet losses and flexibly switches to the secondary path as well as when packets exceed the deadline threshold. In the context of the use case, the resulting reliability of 99.89% is considered an acceptable trade-off (refer to the 99% limit in Section III). However, more stringent use cases could be negatively affected by this.

C. *Latency and jitter impairments*

Figure 4 depicts the case of latency and jitter impairments introduced by the 5G network. This scenario can be considered an intermediate case between scenarios a) and b) with the objective of assessing how the protocols react to small variations in the latency and jitter. As observed in the previous scenario, MPTCP mechanisms result in high latencies as it recovers from latency fluctuations. In general, MPQUIC responds well to impairments; however, there are some iterations where changes in latency cause blocks that lead to latency peaks. MTIP seamlessly adapts to link variations. With MTIP-A, the adaptation is made through accepting the data from the path with lower latency and discarding the data from the other path. However, MTIP-B performs the adaptation in the selection of active paths to send data. In this scenario, MTIP-B achieves an average packet reception rate of 99.85% due to the adaptation and the discarding of late packets. As in scenario b) this loss rate is within the limits of our use case which requested a reliability of 99% to ensure the proper operation of the remotely controlled robot. However, it should be noted again that such losses could pose a challenge in use cases with stricter reliability demands.

D. *Congestion incidents*

Figure 5 illustrates the scenario of network congestion incidents caused by losses in the 5G network and concurrent background video feedback. The objective is to assess how the protocols handle congestion incidents caused by losses and sharing network resources with additional data flows.

MPTCP maintains 100% reliability in this scenario by adapting its path selection to avoid packet loss. However, this reliability comes at the cost of increased latency, as the path congestion caused by the external sources limits the protocol's ability to maintain low delay. MPQUIC is also significantly affected by the increased latency. Moreover, despite its low loss rates in many iterations, the congestion on the path results in connection instability, leading to temporary losses (for instance, in iterations 4, 8 and 12). MPQUIC's relatively low adaptability in handling this type of path congestion causes occasional performance degradation. In contrast, MTIP-A benefits from using both network paths simultaneously to

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

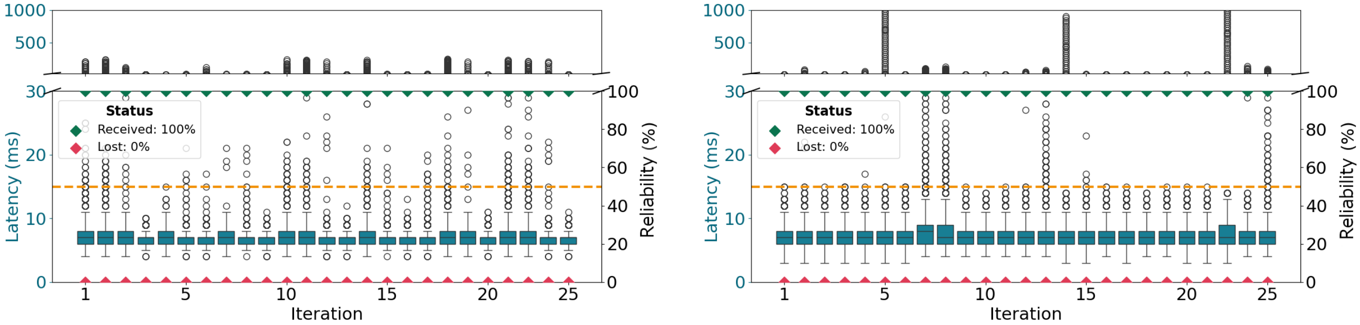

(a) MPTCP.

(b) MPQUIC.

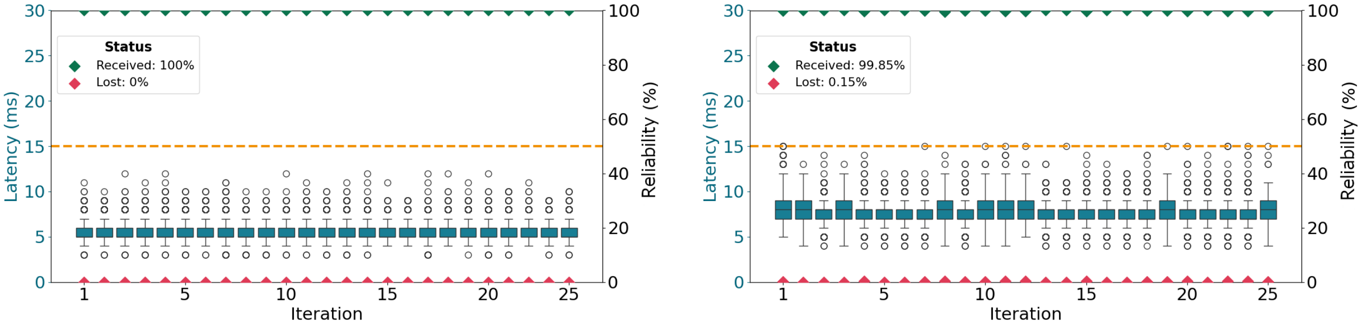

(c) MTIP-A.

(d) MTIP-B.

Fig. 4: Latency and reliability of scenario c) Latency and jitter impairments.

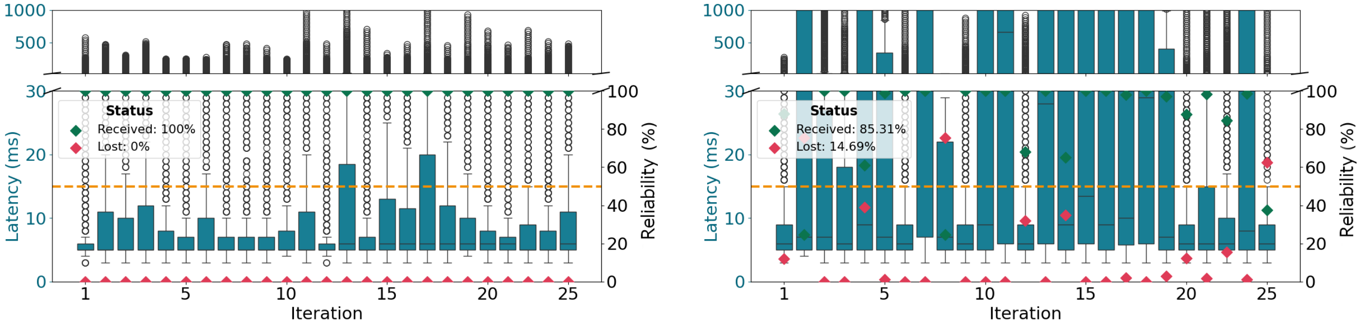

(a) MPTCP.

(b) MPQUIC.

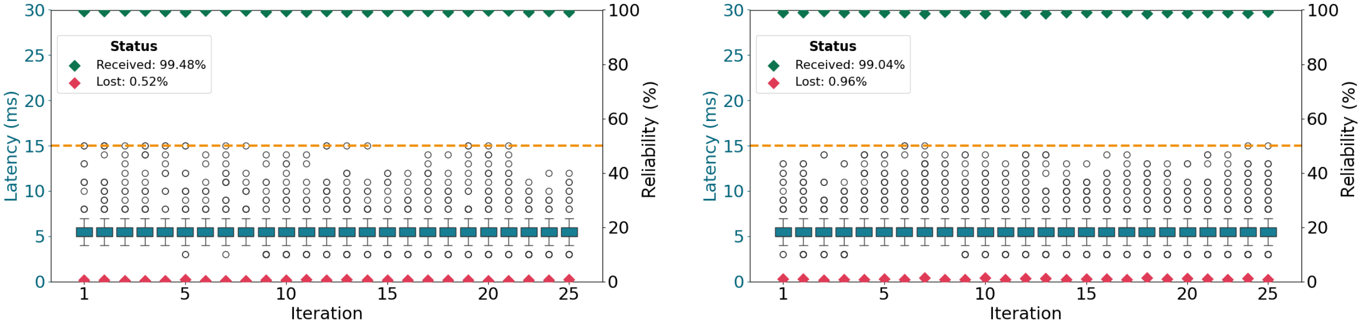

(c) MTIP-A.

(d) MTIP-B.

Fig. 5: Latency and reliability of scenario d) Congestion incidents.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

maintain the required 15 ms latency deadline. This duplication approach allows the protocol to respond to fluctuations in network performance, as continuously leveraging both paths enhances reliability while also reducing the likelihood of latency spikes or packet loss. However, when losses occur in both paths simultaneously, it does not apply recovering mechanisms like MPTCP, resulting in a high but not full reliability with 99.48% of received orders. Finally, MTIP-B also manages to keep packet delivery within the 15 ms constraint, although it faces occasional challenges. In its effort to minimize resource usage, MTIP-B may reduce the number of active paths, which can sometimes result in unrecoverable losses that might affect stricter use cases. However, on average, the protocol can meet the 99% reliability constraint of the use case. The use of paths by each protocol will be further examined in Section V.

V. DISCUSSION

In this section, we explore the reasons behind the behavior of the protocols and evaluate their feasibility for time-sensitive applications. Specifically, we will focus on three aspects: latency, reliability, and resource usage.

A. Latency and reliability evaluation

With regard to latency and reliability, we observe that all protocols perform reasonably well in a stable scenario. It is particularly noteworthy how MPTCP maintains a very stable end-to-end delay, despite this not being its primary objective. However, when network impairments or link failures are introduced, both MPTCP and MPQUIC experience significant degradation latency-wise. This is expected with MPTCP, as it prioritizes reliability at the expense of latency, which is undesirable for time-sensitive applications and should be avoided if impairments are anticipated. In the case of MPQUIC, the official implementation shows poor adaptation to link changes, over-prioritizing reliability over latency, and at times failing to recover from link failures. Finally, MTIP showcases the advantages of its flexibility, specifically designed for time-sensitive applications. It prioritizes latency while maintaining a reasonable level of reliability, even in adverse scenarios. The MTIP-A configuration highlights the benefits of redundancy, providing the best possible functioning based on available paths. On the other hand, MTIP-B demonstrates how the protocol can be configured to achieve the desired behavior without constantly using both paths simultaneously. While it includes mechanisms to re-evaluate paths, the figures show that the impact on latency is minimal, maintaining performance within the required thresholds for the use case.

B. Resource utilization

For the analysis of resource utilization, Figure 6 illustrates the use of the different paths, along with an additional line representing their combined total usage. The examples shown correspond to protocol executions in scenario c) and are accessible in Github[4]. This scenario and the selected images were chosen because they effectively represent the typical patterns and trends observed during all scenarios under evaluation. This discussion is divided into path usage and the amount of data required to transmit the same information.

In terms of path usage, MPTCP and MPQUIC primarily rely on a single path, resulting in the lowest possible combined path utilization as they do not usually use redundant paths. However, this impacts their operation, as discussed in the previous subsection. Regarding MTIP, we note that MTIP-A continuously utilizes all available paths, leading to potential resource waste. In contrast, MTIP-B adapts to use the fewest

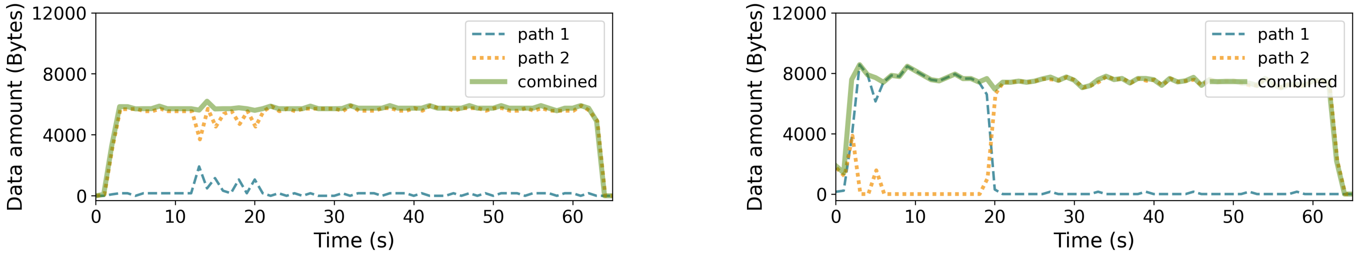

(a) MPTCP (b) MPQUIC

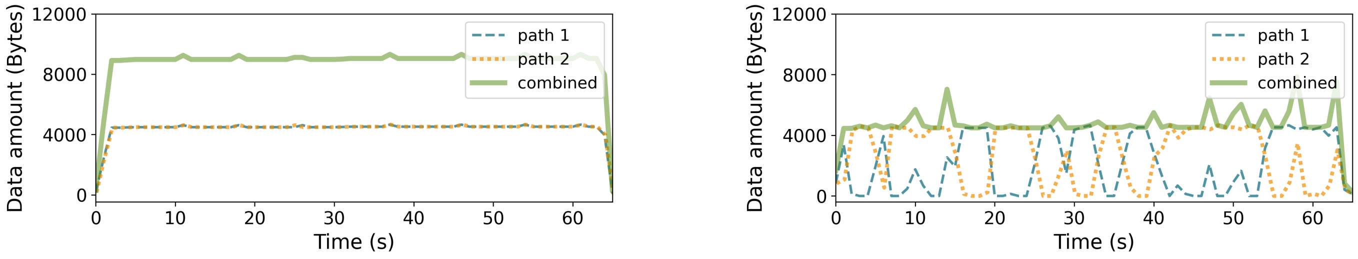

(c) MTIP-A (d) MTIP-B

Fig. 6: Examples of the observed behaviour of the protocols in the use of paths.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This article has been accepted for publication in IEEE Network. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MNET.2025.3583261

necessary paths, constantly reelecting the primary path, while still using some duplication to maintain expected behavior. Consequently, MTIP exhibits higher resource utilization in terms of path usage compared to MPTCP and MPQUIC, as it sometimes uses redundant paths even when a single path might be sufficient. However, this can be considered acceptable for time-sensitive applications, where some degree of resource waste is justified to prioritize performance. The decision between MTIP-A and MTIP-B should be made based on the acceptable level of resource usage versus the tolerance to losses in the specific use case.

There is an additional discussion in the scope of resource utilization, which is the amount of data used to transmit at transport level the same number of application orders. Even if MTIP is the protocol with higher resource utilization in terms of path usage, it is also the one with lower header overhead. This is shown in the images with a low data amount line per path at any given time. Focusing on just one path, MTIP uses less amount of data than MPTCP or MPQUIC to transmit the application orders. The higher data amount is presented by MPQUIC, due to the fact that it not only sends the data, but also performs encryption and security mechanisms.

VI. CONCLUSION

In this paper, we have analyzed the effectiveness of Multipath TCP, Multipath QUIC, and the Multi-connection Tactile Internet Protocol for time-sensitive applications, in particular in the transport of remote control traffic of a quadruped robot. We have created a setup that uses the different transport protocols to forward the remote control traffic of the Ghost Vision 60 UGV, over two differentiated 5G network paths. We have designed multiple scenarios, starting with a 5G baseline, and then introduced various impairments to analyze how the protocols react to link failures, latency variations and congestion incidents.

MPTCP and MPQUIC can support this traffic under favorable network conditions, but they lack specific mechanisms to meet the requirements of such applications. As a result, the application is responsible for key functions, such as discarding traffic that exceeds the latency threshold. MPTCP focuses on 100% reliability that compromises latency in adverse scenarios, whereas the official implementation of MPQUIC lacks quick adaptability, especially to link failures. MTIP demonstrates better operation in terms of latency constraints, particularly in adverse network conditions, due to its context-aware path management. The evaluation results highlight the advantages of MTIP in maintaining low latency and high reliability even under challenging network conditions, although it incurs some losses if resource waste wants to be kept low.

In conclusion, MTIP stands out for time-sensitive applications in the context of future 6G networks, as it leverages advanced network programmability to dynamically adapt to changing conditions, maintaining low latency without compromising the expected reliability. Future work will focus on enhancing the context-aware adaptation of MTIP, incorporating intelligent solutions such as automata learning [14] to automatically reconfigure MTIP internal parameters and refine its decision-making capabilities to further optimize performance in complex network scenarios toward 6G, as the unification of satellite (Non-Terrestrial Networks) and the terrestrial cellular network.

ACKNOWLEDGMENT

Work supported by the UNICO I+D projects 5G+TACTILE_1 (grant agreement TSI-063000-2021-11) and Victoria Satellite (grant agreement TSI-064100-2023-5). As well as the the SNS-JU projects 6G-SANDBOX (grant agreement No 101096328) and 6G-VERSUS (grant agreement No 101192633).

REFERENCES

[1] O. Holland, E. Steinbach, R. V. Prasad, Q. Liu, Z. Dawy, A. Aijaz, N. Pappas, K. Chandra, V. S. Rao, S. Oteafy, M. Eid, M. Luden, A. Bhardwaj, X. Liu, J. Sachs, and J. Araújo, "The IEEE 1918.1 "Tactile Internet" Standards Working Group and its Standards," *Proceedings of the IEEE*, vol. 107, no. 2, pp. 256–279, 2019.

[2] D. Rico and P. Merino, "A Survey of End-to-End Solutions for Reliable Low-Latency Communications in 5G Networks," *IEEE Access*, vol. 8, pp. 192 808–192 834, 2020.

[3] 3rd Generation Partnership Project (3GPP), 3GPP System architecture for the 5G system (5GS), Technical Specification (TS) 23.501, v18.7.0 Sept. 2024, [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/23_series/23.501/23501-i70.zip.

[4] Q. Peng, A. Walid, J. Hwang, and S. H. Low, "Multipath TCP: Analysis, Design, and Implementation," *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 596–609, 2016.

[5] T. Viernickel, A. Froemmgen, A. Rizk, B. Koldehofe, and R. Steinmetz, "Multipath QUIC: A Deployable Multipath Transport Protocol," in *2018 IEEE International Conference on Communications (ICC)*, 2018, pp. 1–7.

[6] D. Rico, K. Grinemmo, A. Brunstrom, and P. Merino, "Performance analysis of the multi-connection tactile internet protocol over 5g," *Journal of Network and Systems Management*, vol. 31, no. 49, 2023. [Online]. Available: https://doi.org/10.1007/s10922-023-09737-0

[7] M. Polese, F. Chiariotti, E. Bonetto, F. Rigotto, A. Zanella, and M. Zorzi, "A Survey on Recent Advances in Transport Layer Protocols," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3584–3608, 2019.

[8] R. Poorzare and A. C. Augé, "Challenges on the Way of Implementing TCP Over 5G Networks," *IEEE Access*, vol. 8, pp. 176 393–176 415, 2020.

[9] B. Han, C. Xu, Y. Li, X. Wang, and P. Xun, "Mpr-quic: Multi-path partially reliable transmission for priority and deadline-aware video streaming," *Journal of Systems Architecture*, vol. 153, p. 103195, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1383762124001322

[10] D. Rico, M. Gallardo, and P. Merino, "Verification of a multi-connectivity protocol for tactile internet applications," *Comput. Commun.*, vol. 212, pp. 390–406, 2023. [Online]. Available: https://doi.org/10.1016/j.comcom.2023.10.013

[11] G. Sarwar, R. Boreli, E. Lochin, and A. Mifdaoui, "Performance evaluation of multipath transport protocol in heterogeneous network environments," in *2012 International Symposium on Communications and Information Technologies (ISCIT)*, 2012, pp. 985–990.

[12] R. Mavropodi and C. Douligeris, "Multipath Routing Protocols for Mobile Ad Hoc Networks: Security Issues and Performance Evaluation," in *Autonomic Communication*, I. Stavrakakis and M. Smirnov, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 165–176.

[13] "Victoria Network: Research in Advanced 5G & 6G," https://www.victoria-network.eu/, accessed: 2024-11-11.

[14] F. Luque-Schempp, L. Panizo, M.-D.-M. Gallardo, and P. Merino, "Automadapt: Zero touch configuration of 5g qos flows extended for time-sensitive networking," *IEEE Access*, vol. 11, pp. 82 960–82 977, 2023.

[15] E. Aumayr, G. Caso, A.-M. Bosneag, A. D. Zayas, Özgü Alay, B. Garcia, K. Kousias, A. Brünstrom, P. M. Gomez, and H. Koumaras, "Service-based analytics for 5g open experimentation platforms," *Computer Networks*, vol. 205, p. 108740, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128621005892

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/