

# Performance Evaluation of Transport Protocols and Roadmap to a High-Performance Transport Design for Immersive Applications

Inayat Ali

*Intelligent Network Research Section*

*Electronics and Telecommunications Research Institute (ETRI)*

Daejeon, South Korea

inayat@etri.re.kr

Seungwoo Hong

*Intelligent Network Research Section*

Daejeon, South Korea

swhong@etri.re.kr

Pyung-koo Park

*Intelligent Network Research Section*

*Electronics and Telecommunications Research Institute (ETRI)*

Daejeon, South Korea

parkpk@etri.re.kr

Tae Yeon Kim

*Intelligent Network Research Section*

Daejeon, South Korea

tykim@etri.re.kr

**Abstract**—Immersive technologies such as virtual reality (VR), augmented reality (AR), and holograms will change users' digital experience. These immersive technologies have a multitude of applications, including telesurgeries, teleconferencing, Internet shopping, computer games, etc. Holographic-type communication (HTC) is a type of augmented reality media that provides an immersive experience to Internet users. However, HTC has different characteristics and network requirements, and the existing network architecture and transport protocols may not be able to cope with the stringent network requirements of HTC. Therefore, in this paper, we provide an in-depth and critical study of the transport protocols for HTC. We also discuss the characteristics and the network requirements for HTC. Based on the performance evaluation of the existing transport protocols, we propose a roadmap to design new high-performance transport protocols for immersive applications.

**Index Terms**—Holographic-type communication, Transport protocols, Immersive applications, AR, VR

## I. INTRODUCTION

Virtual reality (VR), augmented reality (AR), mixed reality (MR), and holographic technologies are examples of immersive technologies. These technologies have gained much attention from industry and academia in recent years. Applications of immersive technologies include but are not limited to holographic teleportation, tactile robotic surgery, virtual orchestra, E-commerce, and video games. Holographic-type communication is an emerging type of communication that involves the transmission and interaction of volumetric holograms over the Internet. For example, holographic teleportation will facilitate real-time interaction with the remote participants of a meeting by projecting them into the meeting room in the form of

a holographic object. Similarly, holographic telesurgery will allow surgeons to perform surgery with more ease, and clarity [3]. HTC and immersive applications are used interchangeably in this article since both are immersive technologies and generate volumetric data. The need to handle this kind of immersive applications and the volumetric data generated by these has become vital as big Internet giants like Facebook [1] and Roblox kick-started their immersive services. Roblox's monthly active users are over 150 million [2]; imagine if most services on their platform go immersive, this will put a heavy burden on network resources. Particularly, we focus on holographic-type communication (HTC). A plethora of Internet applications, e.g., telesurgery, virtual orchestra, and tele-conferencing, are on their way to adopting holographic technologies. Moreover, these applications will put significant demands on network infrastructures, such as high bandwidth, ultra-low latency, and synchronization of multiple data streams, which the current network standards, especially the transport protocols, do not support.

Novel network architecture and transport protocols are needed to combat the requirement of HTC. The current standard of transport protocols may not be able to handle this kind of volumetric data while maintaining other application-specific requirements such as reliability, application-specific latency, and on-time delivery of specific data. At the transport layer, video streaming services either use TCP for quality or UDP for low latency and high throughput. TCP-based solutions include HTTP with adaptive streaming (HAS), and real-time media protocol (RTMP) [4]. UDP-based solutions, which are adopted for low latency, include the real-time protocol (RTP) combined with RTCP. Novel approaches are mostly built on top of UDP, such as QUIC [5] and WebRTC [6]. For designing a novel transport layer solution for HTC, it is necessary to first

This work is supported by Electronics and Telecommunications Research Institute (ETRI) grant funded by ICT R&D program of MSIT/IITP[2021-0-00715, Development of End-to-End Ultra-high Precision Network Technologies].

deeply study the working and performance of current transport layer standards for HTC. Until now, these standard protocols have not been evaluated for holographic-type communication. QUIC, particularly is a new protocol and has never been critically studied in comparison with the existing protocol for HTC.

The contribution of this paper is threefold:

- (1). This paper provides a detailed evaluation and comparison of TCP, UDP, and QUIC protocols for HTC.
- (2). Discuss the challenges in the existing transport protocol and network requirements of HTC and other immersive network applications.
- (3). Present a roadmap to the design of novel high-performance transport protocols for immersive network applications that will guide our future research.

The rest of the paper is organized as follows. We discuss the challenges and network requirements of holographic-type communication in section 2. In section 3, we give an overview of the standard transport protocols. We provide critical discussion and performance analysis of the UDP, TCP, and QUIC under different congestion variants in section 4. In section 5, a roadmap to designing novel transport protocols for HTC will be presented, and then we will conclude the work in section 6.

## II. CHALLENGES AND NETWORK REQUIREMENTS OF IMMERSIVE APPLICATIONS

Holograms are volumetric objects and require high bandwidth even after applying high-quality compression techniques. HTC, therefore, have very stringent network requirements. HTC objects require continuous user interactivity, whether streamed in real-time or dealing with offline content, as the user will still interact with the object by changing the viewing angle and position. The portion of the HTC object visible from the current view angle must be retrieved first compared to the data corresponding to the portion not visible from the current angle. This user interactivity is difficult to achieve as predicting the current view angle is challenging. This nature of HTC content in immersive applications demands ultra-high bandwidth and ultra-low latency.

### A. High Bandwidth

The bandwidth demand of HTC will increase by several orders of magnitude over 3D VR content. For example, a Microsoft Kinect for Windows v2 generates nearly 70 Mbits of raw data per frame or 2 Gbits of raw data per second at 30 FPS. After applying high-quality compression techniques, this data can be reduced to approximately 40 Mbits per second per single camera sensor. The quality of a hologram depends on the number of sensors used to capture a hologram. A minimum of 8 sensors will generate over 300 Mbits of data per second. The amount increases with the increased quality of images captured, the number of camera sensors, and the number of viewpoints. Hence very high bandwidth will be required to stream HTC.

### B. Ultra-Low Latency

Ultra-low latency will be required for a truly immersive experience. However, ultra-low latency for such volumetric data with high bandwidth requirements is hard to achieve. The latency requirement becomes very challenging in the case of HTC. Because not only real-time HTC streaming requires ultra-low latency, but also static HTC content needs ultra-low latency as the user may change the view angle continuously, which will require the application to retrieve part of the HTC content visible from that view angle. So this continuous user interactivity with HTC objects makes the latency requirement of HTC more challenging.

### C. Application Requirements

Different network applications have different requirements from the network. Some applications are delay sensitive, and some are loss sensitive. These requirements are common in traditional network applications, but in immersive applications, along with in-time delivery, on-time delivery of data may also be vital, for example, in telesurgery. Data spread over multiple streams must be received on time to avoid inconsistency and ensure synchronization. These strict application requirements must be addressed while designing standards for holographic-type communication.

## III. TRANSPORT PROTOCOLS OVERVIEW

Until recently, only UDP and TCP were the dominant transport layer protocols. Different applications adopt either UDP or TCP based on the application requirements. Though fast, UDP is a connection-less protocol used by applications that do not require guaranteed delivery like most video streaming and broadcasting applications. Applications that require guaranteed delivery of data and quality use TCP. Many transport protocols, such as stream control transmission protocol (SCTP) [9], were developed with time. However, due to the design of the Internet, these protocols could not be used on the public scale because the middleboxes usually drop packets other than UDP and TCP. QUIC [11], developed by google and deployed in past years, is now becoming the most dominant transport protocol since it is integrated into google and YouTube platforms. QUIC runs on top of UDP, and most of the QUIC header is encrypted, enabling it to traverse the Internet without being dropped by the middleboxes. QUIC overcomes many problems in TCP, e.g., Head of line blocking (HoL), long handshake delay, packet reordering issues, etc. A single QUIC connection is composed of multiple streams recognized by stream id. These streams can be mapped to HTTP/2 streams. This feature is very effective as the loss of packets in one stream does not block other streams in the same connection hence resolving the HoL issue. Furthermore, QUIC connection establishment requires one RTT compared to TCP/TLS 1.3, which requires 2 RTTs. QUIC also introduced 0 RTT connection establishment for recently closed connections. QUIC also resolves the retransmission ambiguity in TCP, where the TCP can not recognize whether the ACK is for the first transmission or retransmission. QUIC uses per stream

buffer at sender and receiver that ameliorates HoL issues among streams.

#### IV. PERFORMANCE EVALUATION OF TRANSPORT PROTOCOLS FOR IMMERSIVE APPLICATIONS

We perform a critical analysis of the standard transport protocols, including QUIC and congestion variants like YeAH [7] and Vegas [8]. We use immersive applications like volumetric data in different network scenarios to check the performance of the transport protocols. We compare the performance of the standard transport protocols based on end-to-end delay, throughput, jitter, delivery ratio, and retrieval time of the HTC objects. We also analyze the protocols by looking at the RTT variations, congestion window, and queue occupancy.

##### A. Simulation Environment

We use ns-3, an open-source network simulator, for creating the simulation environment. We created two network scenarios, i.e., a simple LAN scenario to see the performance of the protocols at full throttle and a high latency WAN scenario to analyze the performance of transport protocols in a dynamic high latency WAN as shown in figure 1.

##### B. Results and Discussion

We are particularly interested in the performance of TCP and QUIC as being the widely used protocols by applications that require guaranteed delivery of data. We use Vegas and YeAH as congestion variants in TCP and QUIC for a fair comparison. We first recorded the retrieval time of the HTC object, which is 400 Mbits each by UDP, TCP, and QUIC under different congestion variants. The results in figure 2 show that UDP takes less time to retrieve the HTC object as it does not have any flow and congestion control. However, UDP can not be adopted by applications that require guaranteed delivery since UDP encounters huge packet loss and high end-to-end delay, especially when the network is congested. Among TCP and QUIC, QUIC-YeAh and Vegas outperform TCP-YeAH and Vegas, respectively. This is because of the multi-streaming features of the QUIC and the per-stream buffer at the sender and receiver. We also look into the throughput achieved by these protocols in figure 3. As expected, UDP achieves the highest throughput, 374mbps, near the link capacity (400mbps). However, UDP also encounters packet losses, which is not suitable for most immersive applications. QUIC-YeAH achieves 316 Mbps of

throughput compared to 288 Mbps achieved by TCP-YeAH, as shown in figure 3a. Similarly, QUIC-Vegas achieves higher throughput than TCP-vegas though the number is lower than what is achieved by YeAH. We also compared the throughput performance in the WAN scenario as shown in figure 3b. The average throughput achieved with respect to varying send/receive buffer size is shown in figure 3b. The throughput drops heavily in connection-oriented protocols compared to unreliable connection-less UDP. This heavy drop is because of the high WAN latency and the flow control measures of the connection-oriented protocols. A critical observation from this figure is that unlike TCP, throughput achieved by QUIC is not affected by send/receive buffer size; this is because QUIC maintains per stream buffer, which avoids memory contention at send/receive buffer. Also, we see that the throughput in TCP increases with send/receive buffer size, and the line becomes straight after the buffer size is more than 2 megabytes. This also attests to our previous observation that at smaller buffer sizes, the throughput decreases because of the memory contention at the buffer that increases the RTT, which activates the flow control to avoid congestion.

We now investigate further into the protocols by analyzing the end-to-end delay. In LAN scenario figure 4a, QUIC-YeAH has a high end-to-end delay compared to TCP-YeAH. We investigated the reasons behind this and concluded that this high end-to-end delay is because QUIC-YeAH keeps the pipe full to achieve high throughput, which causes queue latency at nodes that add up to the average end-to-end delay. Figure 6 shows queue occupancy. In figure 6, UDP keeps the queue full all the time, which puts the network in a congestion state and causes packet loss and very high end-to-end delay (247ms). The queue size in the experiment is 100 packets. QUIC-YeAH keeps the queue nearly full compared to TCP-YeAH. Once the queue becomes full, QUIC flow control slows down the ingestion of packets to avoid congestion. The reasons for the average end-to-end delay trend of the protocols become more clear from the results of queue occupancy. This behavior of QUIC also affects the average jitter, as shown in figure 5. The jitter in QUIC is high compared to TCP as QUIC keeps

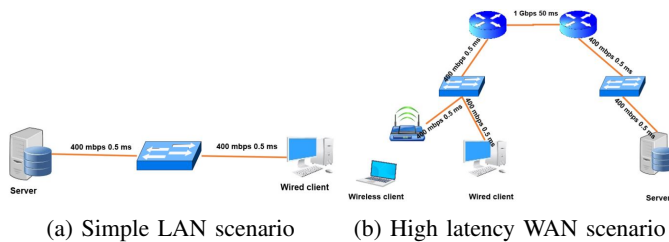


Fig. 1: Network scenario for performance evaluation

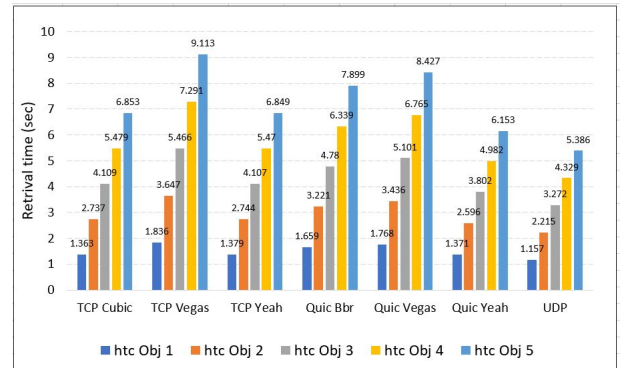


Fig. 2: Retrieval time of HTC object by different transport protocols and congestion variants

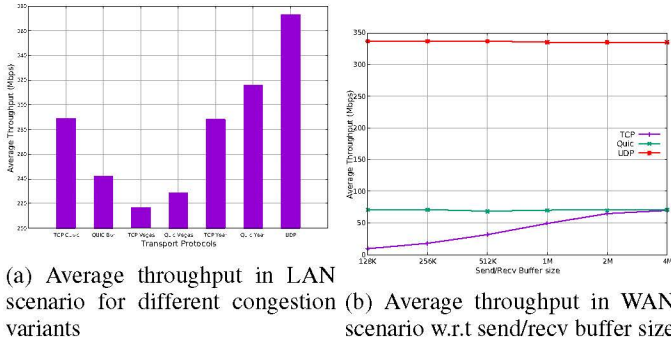


Fig. 3: Average throughput of transport protocols

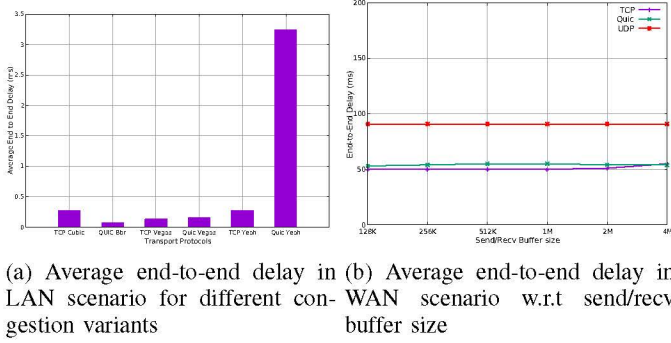


Fig. 4: Average end-to-end delay in transport protocols

the queue near its full capacity and then abruptly controls the flow to avoid congestion, unlike TCP, which keeps the queue capacity at nearly half. This causes variation in packet arrival at the receiver. On the other hand, the jitter in UDP is very low for obvious reasons. The average jitter results in the WAN scenario figure 5b are rather more interesting. Our observations from the figure are that jitter increases with the increase in throughput; as we can see from the throughput results (figure 3b), QUIC achieves high throughput than TCP when the send/receive buffer size varies from 128K to 2M. A similar trend can be seen in average jitter. QUIC faces high jitter than TCP when the send/receive buffer size varies from 128K to 2M (figure 5b). At 4M send/receive buffer size, the average throughput of TCP equates to QUIC throughput, and the average jitter of TCP shows an abrupt increase and equates to QUIC's average jitter. On further investigation, this pattern is caused by the similar queue occupancy of TCP and QUIC at 4 MB of send/receive buffer size.

We also show the congestion window change in figure 7 to make it easy to understand the congestion control behavior of YeAH. The congestion window increases exponentially and steeply drops when congestion is detected. Connection-oriented reliable transport protocols perform worst in the wireless environment due to heavy losses of packets at the wireless channel. Figure 8 shows the average throughput in a wireless environment where only the client is connected through WiFi. The throughput of all three protocols drops heavily because of the losses due to the wireless interface. UDP throughput,

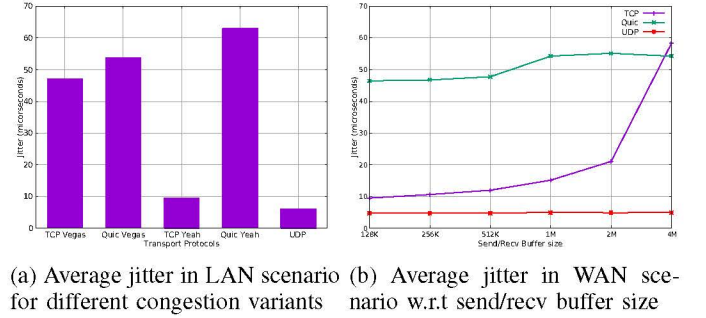


Fig. 5: Average jitter in transport protocols

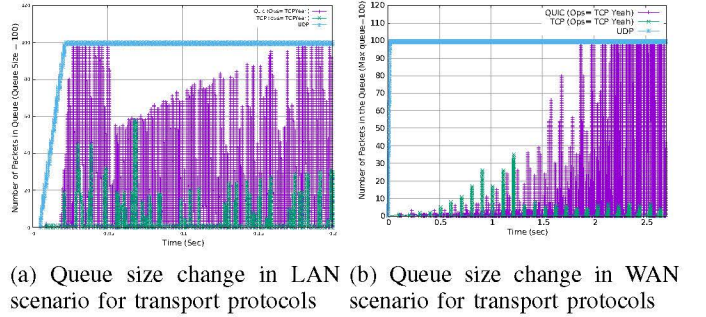


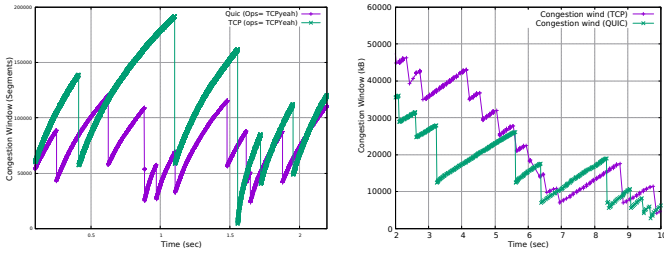
Fig. 6: Queue size change in transport protocols (Congestion option is "Yeah")

which was 340 Mbps, drops to 76 Mbps when retrieving data using WiFi. Similarly, QUIC and TCP throughput drops heavily, as shown in figure 8a. The trend is similar to that in the wired environment, except that TCP at a larger send/receive buffer size achieves slightly higher throughput. The end-to-end Delay of UDP is very high because of the frequent losses at the wireless interface and re-transmission at the mac layer. The end-to-end delay of QUIC and TCP shows a similar trend as in a wired environment; however, the average end-to-end delay increases by approximately 15 ms, which is the effect of losses at wireless interfaces and latency added by mac layer re-transmission. Finally, figure 9 shows the Delivery ratio of wired and wireless clients in the WAN scenario. Since TCP and QUIC guarantee packet delivery gives a 100% delivery ratio at the cost of higher end-to-end delay. On the other hand, the delivery ratio of UDP is worst in the wireless environment, as shown in figure 9a. The delivery ratio of the wired client is shown in figure 9b, which is as expected 100% for TCP and QUIC while 94% for UDP, keeping in mind that the Linux traffic control layer was enabled, resulting in a high delivery ratio for UDP.

## V. ROADMAP FOR DESIGNING NOVEL TRANSPORT PROTOCOLS FOR IMMERSIVE APPLICATIONS

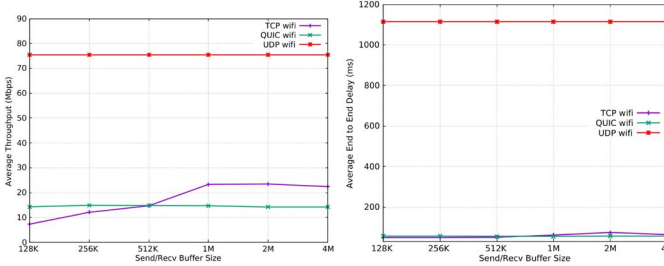
Network applications and their requirements have changed rapidly in the past decade resulting in new challenges for the existing transport protocols. Soon most of the Internet traffic will be generated by immersive applications, including holographic-type communication. Therefore, a new high-





(a) Congestion window in LAN scenario for transport protocols (b) Congestion window in WAN scenario for transport protocols

Fig. 7: Congestion window in transport protocols (Congestion option is "Yeah")



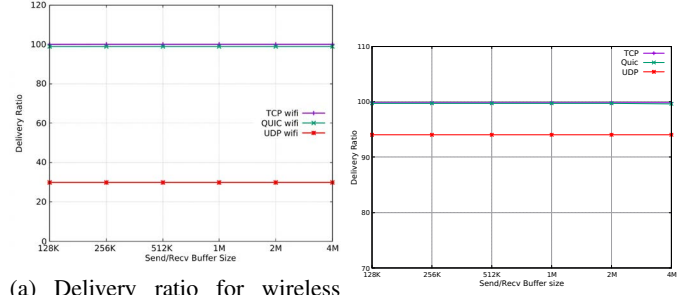
(a) Average throughput for wireless client (b) End-to-end delay for wireless client

Fig. 8: Throughput and delay for wireless client

performance transport design is needed from scratch to tackle the challenges of modern immersive applications and their possible future demands.

Congestion control algorithms are the key to transport protocols. However, their performance is affected by erroneous congestion detection. For example, RTT variation caused by MAC layer re-transmission and packet loss caused by a lossy wireless link sometimes triggers the congestion control algorithm. Artificial intelligence may be a valuable alternative to detect congestion in the network. Similarly, out-of-order packet delivery triggers false congestion. A new robust mechanism for congestion detection and out-of-order packet delivery must be a feature of the transport protocols. The congestion control algorithm must consider the challenges caused by the wireless environment. The fairness concept in TCP must be revisited in the design of new transport protocols. A significant amount of Internet traffic is caused by software updates and some other mice flow, which are highly delay tolerant and can be assigned fewer network resources than other connections. Furthermore, in the case of multi-streaming, per stream buffer can further enhance the performance.

Holographic objects can be viewed from different angles, and not all areas of the HTC object are visible to the user based on the user's viewing position. Transport protocols should be able to adapt themselves to decide which part of the content to stream first and at what quality. However, this requires the protocols to predict the user's viewpoint, and movement [12]. A few essential features of high-performance transport



(a) Delivery ratio for wireless client (b) Delivery ratio for wired client

Fig. 9: Delivery ratio for wireless and wired client

protocols are discussed below.

#### A. Application-specific Requirements

Different network applications may have different requirements, and fixed rules of protocols will not be able to do justice to the application requirements. For example, some applications may require high reliability but may not require ultra-low latency. Some applications may be loss tolerant but may require high bandwidth and low latency, such as online games and OTT services. Similarly, some applications, such as telesurgery, may require ultra-low latency and strict deadlines. Therefore, new high-performance transport protocols must consider application requirements in their design and try to meet the application requirements. A viable solution is a cross-layer design to meet the application-specific requirements. The network layer can provide explicit knowledge of the network condition, such as congestion state, open paths, multi-path etc. The transport layer can easily exploit this knowledge to ensure application-specific requirements.

#### B. Multicast support

Most connection-oriented reliable protocols do not support multicast communication. In traditional transport protocols, if there is one sender and  $n$  receivers, a single packet will trigger  $n$  acknowledgment packets and there will be  $n$  total flows as shown in figure 10a. The design of new high-performance transport protocols must include multicast support as a huge percentage of communication on the Internet is based on multicast. Multicast saves bandwidth as the server sends only one content packet, which is multicasted by the network to a multicast group as shown in figure 10b. A new mechanism of ack and retransmission of lost packets should be implemented. For example, the gateway router can receive individual acks and sends one ack back to the server. Moreover, a small buffer could be used at the multicast edge router to retransmit the lost packet. This will save massive bandwidth.

#### C. Segment level reliability

The existing transport protocols transmit data without knowing its importance in a connection. For example, some content may go unnoticed by the users and may be sent at lower quality or even be ignored if that particular segment is lost. A considerable amount of bandwidth can be saved if the protocol deals

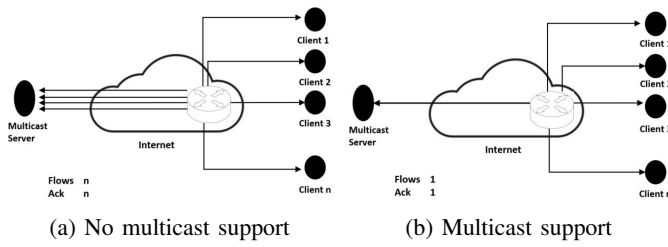


Fig. 10: Multicast support

with these less critical segments differently. This feature will impact the performance of transport protocols for immersive applications as a considerable amount of volumetric data can be sent at lower quality and can be ignored if a loss occurs, saving considerable bandwidth.

#### D. Multi-streaming and connection migration

Multi-streaming is sending data of a single flow on multiple streams. This not only achieves high throughput but also avoid HoL blocking issue. Moreover, high-priority data in a flow could also benefit from multi-streaming by assigning it to a single high-priority stream. For example, critical data in a flow which is loss sensitive must be assigned to a specific stream with loss-sensitive requirements from the application. Similarly, connection migration is also a must-needed feature in any new transport protocol because of the widespread of the Internet in the form of WiFi and LTE. With the increasing popularity of mobile devices with multiple Internet connections simultaneously, connection migration from one connection over to a more stable connection is vital for reliability, as explained in figure 11. Traditional transport protocols do not have multi-streaming and connection migration features. However, the recent release of QUIC protocols has both of the mentioned features. These features are vital for modern high-performance transport protocols for immersive applications. Application-level protocols like HTTP/2 supports multi-streaming to avoid HoL blocking issue; however, these multiple streams are multiplexed on a single TCP stream at the transport layer limiting the benefit of HTTP/2 multi-streaming. Connection migration and keeping recent connection states to avoid handshake delay when the connection times out due to a congested wired link or lossy wireless channel must be included in the design of the transport protocols.

## VI. CONCLUSION AND FUTURE WORK

Holographic-type communication is going to be vital for future hologram-based networked applications. Holograms and extended reality are emerging media types that will replace the current 2D and 3D volumetric content. However, this kind of media requires new network architecture and protocols for guaranteed user experience. In this paper, we have discussed the network requirements of the HTC, then we have provided an in-depth analysis of the standard transport protocols, i.e., UDP, TCP, and QUIC. We compared their performance and provided a critical discussion on the results. Based on the

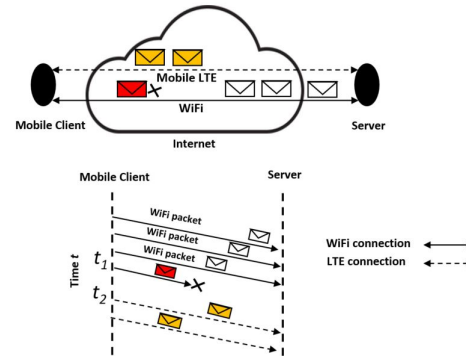


Fig. 11: Connection migration in transport protocol

results and discussion, we presented a roadmap for the design of a novel high-performance transport protocol. This detailed study will help the researchers to get a deep insight into the current transport protocols for immersive applications and takeaways for designing a novel, efficient transport protocol.

In future work, we are going to design and develop a new high-performance transport protocol for immersive applications.

## ACKNOWLEDGMENT

This work was supported by Electronics and Telecommunications Research Institute (ETRI) grant funded by ICT R&D program of MSIT/IITP[2021-0-00715, Development of End-to-End Ultra-high Precision Network Technologies].

## REFERENCES

- [1] M. Zuckerberg, "Founder's Letter, 2021," Accessed: December, 2022, [Online]. Available: <https://about.fb.com/news/2021/10/founders-letter/>
- [2] R. U. Long, "Roblox and effect on education," M.S. thesis, Dept. Educ., Drury Univ., Springfield, MO, USA, 2019
- [3] R. Christophe, Laurent, Dirk K. Ysebaert, Eddy Vanhoecke, Robrecht H. Van Hee, *Interactive Technology and the New Paradigm for Healthcare* Ebook. IOS Press pages 187-190, vol 18.
- [4] V. Swaminathan, "Are We in the Middle of a Video Streaming Revolution?", ACM TOMM. vol. 9, no. 1s, pp. 40:1–40:6, Oct. 2013.
- [5] Langley, A., Riddoch, A., Wilk, A., Vicente, A., Krasic, C., Zhang, D., Yang, F., Kouranov, F., Swett, I., Iyengar, J. and Bailey, J., "The quic transport protocol: Design and internet-scale deployment", In Proceedings of the conference of the ACM special interest group on data communication. pp. 183-196, 2017, Los Angeles, USA.
- [6] H. Alvestrand, "Overview: Real Time Protocols for Browser-based Applications draft-ietf-rtcweb-overview-19", Network Working Group, Internet-draft, Nov 2017.
- [7] A. Baiocchi, A. P. Castellani, F. Vacirca., "YeAH-TCP: Yet Another Highspeed TCP", In Proc. PFLDnet, vol. 7, pp. 37-42, Feb 2007.
- [8] Brakmo, L. S., Peterson, L. L., "TCP Vegas: End to end congestion avoidance on a global Internet", IEEE Journal on selected Areas in communications, vol. 13, no. 8, pp. 1465-1480.
- [9] R. Stewart, "Stream control transport protocol.", IETF, RFC 6582, Sep. 2007. [Online]. Available: <https://rfc-editor.org/rfc/rfc6582.txt>
- [10] P. Srisuresh, J. Kuthan, J. Rosenberg, A. Molitor, and A. Rayhan, "Middlebox communication architecture and framework.", IETF, RFC 3303, Aug. 2002. [Online]. Available: <https://rfc-editor.org/rfc/rfc3303.txt>
- [11] J. Iyengar, Ed. Fastly, M. Thomson, Ed. Mozilla, "QUIC: A UDP-Based Multiplexed and Secure Transport.", IETF, In RFC 9000. 2021 [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-34>
- [12] A. Clemm, M. Torres Vega, H. K. Ravuri, T. Wauters and F. De Turck, "Towards Truly Immersive Holographic-Type Communications: Challenges and Solutions.", IEEE Communications Magazine vol. 58, no. 1, pp. 93-99. 2020.