MOPS
Internet-Draft
Intended status: Informational
Expires: 21 December 2024

R. Krishna

A. Rahman
Ericsson

19 June 2024

Media Operations Use Case for an Extended Reality Application on Edge Computing Infrastructure

draft-ietf-mops-ar-use-case-18

Abstract

This document explores the issues involved in the use of Edge Computing resources to operationalize media use cases that involve Extended Reality (XR) applications. In particular, this document discusses those applications that run on devices having different form factors (such as different physical sizes and shapes) and need Edge computing resources to mitigate the effect of problems such as a need to support interactive communication requiring low latency, limited battery power, and heat dissipation from those devices. The intended audience for this document are network operators who are interested in providing edge computing resources to operationalize the requirements of such applications. This document discusses the expected behavior of XR applications which can be used to manage the traffic. In addition, the document discusses the service requirements of XR applications to be able to run on the network.

Status of This Memo

This Internet-Draft will expire on 21 December 2024.

Copyright Notice

Table of Contents

1. Introduction

Extended Reality (XR) is a term that includes Augmented Reality (AR), Virtual Reality (VR) and Mixed Reality (MR) [XR]. AR combines the real and virtual, is interactive and is aligned to the physical world of the user [AUGMENTED_2]. On the other hand, VR places the user inside a virtual environment generated by a computer [AUGMENTED].MR merges the real and virtual world along a continuum that connects completely real environment at one end to a completely virtual environment at the other end. In this continuum, all combinations of the real and virtual are captured [AUGMENTED].

XR applications will bring several requirements for the network and the mobile devices running these applications. Some XR applications such as AR require a real-time processing of video streams to recognize specific objects. This is then used to overlay information on the video being displayed to the user. In addition, XR applications such as AR and VR will also require generation of new video frames to be played to the user. Both the real-time processing of video streams and the generation of overlay information are computationally intensive tasks that generate heat [DEV_HEAT_1], [DEV_HEAT_2] and drain battery power [BATT_DRAIN] on the mobile device running the XR application. Consequently, in order to run applications with XR characteristics on mobile devices, computationally intensive tasks need to be offloaded to resources provided by Edge Computing.

Edge Computing is an emerging paradigm where for the purpose of this document, computing resources and storage are made available in close network proximity at the edge of the Internet to mobile devices and sensors [EDGE_1], [EDGE_2]. A computing resource or storage is in close network proximity to a mobile device or sensor if there is a short and high-capacity network path to it such that the latency and bandwidth requirements of applications running on those mobile devices or sensors can be met. These edge computing devices use cloud technologies that enable them to support offloaded XR applications. In particular, cloud implementation techniques [EDGE_3] such as the follows can be deployed:

* Disaggregation (using SDN to break vertically integrated systems into independent components- these components can have open interfaces which are standard, well documented and not proprietary),

* Virtualization (being able to run multiple independent copies of those components such as SDN Controller apps, Virtual Network Functions on a common hardware platform).

* Commoditization (being able to elastically scale those virtual components across commodity hardware as the workload dictates).

Such techniques enable XR applications requiring low-latency and high bandwidth to be delivered by proximate edge devices. This is because the disaggregated components can run on proximate edge devices rather than on remote cloud several hops away and deliver low latency, high bandwidth service to offloaded applications [EDGE_2].

This document discusses the issues involved when edge computing resources are offered by network operators to operationalize the requirements of XR applications running on devices with various form factors. A network operator for the purposes of this document is any organization or individual that manages or operates the compute resources or storage in close network proximity to a mobile device or sensors. Examples of form factors include Head Mounted Displays (HMD) such as Optical-see through HMDs and video-see-through HMDs and

Hand-held displays. Smart phones with video cameras and location sensing capabilities using systems such as a global navigation satellite system (GNSS) are another example of such devices. These devices have limited battery capacity and dissipate heat when running. Besides as the user of these devices moves around as they run the XR application, the wireless latency and bandwidth available to the devices fluctuates and the communication link itself might fail. As a result, algorithms such as those based on adaptive-bit-rate techniques that base their policy on heuristics or models of deployment perform sub-optimally in such dynamic environments [ABR_1]. In addition, network operators can expect that the parameters that characterize the expected behavior of XR applications are heavy-tailed. Heaviness of tails is defined as the difference from the normal distribution in the proportion of the values that fall a long way from the mean [HEAVY_TAIL_3]. Such workloads require appropriate resource management policies to be used on the Edge. The service requirements of XR applications are also challenging when compared to the current video applications. In particular several Quality of Experience (QoE) factors such as motion sickness are unique to XR applications and must be considered when operationalizing a network. This document motivates these issues with a use-case that is presented in the following sections.

2. Use Case

A use case is now described that involves an application with XR systems' characteristics. Consider a group of tourists who are being conducted in a tour around the historical site of the Tower of London. As they move around the site and within the historical buildings, they can watch and listen to historical scenes in 3D that are generated by the XR application and then overlaid by their XR headsets onto their real-world view. The headset then continuously updates their view as they move around.

The XR application first processes the scene that the walking tourist is watching in real-time and identifies objects that will be targeted for overlay of high-resolution videos. It then generates high-resolution 3D images of historical scenes related to the perspective of the tourist in real-time. These generated video images are then overlaid on the view of the real-world as seen by the tourist.

This processing of scenes and generation of high-resolution images is now discussed in greater detail.

2.1. Processing of Scenes

The task of processing a scene can be broken down into a pipeline of three consecutive subtasks namely tracking, followed by an acquisition of a model of the real world, and finally registration [AUGMENTED].

Tracking: The XR application that runs on the mobile device needs to track the six-dimensional pose (translational in the three perpendicular axes and rotational about those three axes) of the user’s head, eyes and the objects that are in view [AUGMENTED]. This requires tracking natural features (for example points or edges of objects) that are then used in the next stage of the pipeline.

Acquisition of a model of the real world: The tracked natural features are used to develop a model of the real world. One of the ways this is done is to develop an annotated point cloud (a set of points in space that are annotated with descriptors) based model that is then stored in a database. To ensure that this database can be scaled up, techniques such as combining a client-side simultaneous tracking and mapping and a server-side localization are used to construct a model of the real world [SLAM_1], [SLAM_2], [SLAM_3], [SLAM_4]. Another model that can be built is based on polygon mesh and texture mapping technique. The polygon mesh encodes a 3D object’s shape which is expressed as a collection of small flat surfaces that are polygons. In texture mapping, color patterns are mapped on to an object’s surface. A third modelling technique uses a 2D lightfield that describes the intensity or color of the light rays arriving at a single point from arbitrary directions. Such a 2D lightfield is stored as a two-dimensional table. Assuming distant light sources, the single point is approximately valid for small scenes. For larger scenes, many 3D positions are additionally stored making the table 5D. A set of all such points (either 2D or 5D lightfield) can then be used to construct a model of the real world [AUGMENTED].

Registration: The coordinate systems, brightness, and color of virtual and real objects need to be aligned with each other and this process is called registration [REG]. Once the natural features are tracked as discussed above, virtual objects are geometrically aligned with those features by geometric registration. This is followed by resolving occlusion that can occur between virtual and the real objects [OCCL_1], [OCCL_2]. The XR application also applies photometric registration [PHOTO_REG] by aligning the brightness and color between the virtual and real objects. Additionally, algorithms that calculate global illumination of both the virtual and real objects [GLB_ILLUM_1], [GLB_ILLUM_2] are executed. Various algorithms to deal with artifacts generated by lens distortion [LENS_DIST], blur [BLUR], noise [NOISE] etc. are also required.

2.2. Generation of Images

The XR application must generate a high-quality video that has the properties described in the previous step and overlay the video on the XR device’s display- a step called situated visualization. A situated visualization is a visualization in which the virtual objects that need to be seen by the XR user are overlaid correctly on the real world. This entails dealing with registration errors that may arise, ensuring that there is no visual interference [VIS_INTERFERE], and finally maintaining temporal coherence by adapting to the movement of user’s eyes and head.

3. Technical Challenges and Solutions

As discussed in section 2, the components of XR applications perform tasks such as real-time generation and processing of high-quality video content that are computationally intensive. This section will discuss the challenges such applications can face as a consequence.

As a result of performing computationally intensive tasks on XR devices such as XR glasses, excessive heat is generated by the chip-sets that are involved in the computation [DEV_HEAT_1], [DEV_HEAT_2]. Additionally, the battery on such devices discharges quickly when running such applications [BATT_DRAIN].

A solution to the heat dissipation and battery drainage problem is to offload the processing and video generation tasks to the remote cloud. However, running such tasks on the cloud is not feasible as the end-to-end delays must be within the order of a few milliseconds. Additionally, such applications require high bandwidth and low jitter to provide a high QoE to the user. In order to achieve such hard timing constraints, computationally intensive tasks can be offloaded to Edge devices.

Another requirement for our use case and similar applications such as 360-degree streaming (streaming of video that represents a view in every direction in 3D space) is that the display on the XR device should synchronize the visual input with the way the user is moving their head. This synchronization is necessary to avoid motion sickness that results from a time-lag between when the user moves their head and when the appropriate video scene is rendered. This time lag is often called "motion-to-photon" delay. Studies have shown [PER_SENSE], [XR], [OCCL_3] that this delay can be at most 20ms and preferably between 7-15ms in order to avoid the motion sickness problem. Out of these 20ms, display techniques including the refresh rate of write displays and pixel switching take 12-13ms [OCCL_3], [CLOUD]. This leaves 7-8ms for the processing of motion sensor inputs, graphic rendering, and round-trip-time (RTT) between the XR device and the Edge. The use of predictive techniques to mask latencies has been considered as a mitigating strategy to reduce motion sickness [PREDICT]. In addition, Edge Devices that are proximate to the user might be used to offload these computationally intensive tasks. Towards this end, a 3GPP study indicates an Ultra Reliable Low Latency of 0.1ms to 1ms for communication between an Edge server and User Equipment (UE) [URLLC].

Note that the Edge device providing the computation and storage is itself limited in such resources compared to the Cloud. So, for example, a sudden surge in demand from a large group of tourists can overwhelm that device. This will result in a degraded user experience as their XR device experiences delays in receiving the video frames. In order to deal with this problem, the client XR applications will need to use Adaptive Bit Rate (ABR) algorithms that choose bit-rates policies tailored in a fine-grained manner to the resource demands and playback the videos with appropriate QoE metrics as the user moves around with the group of tourists.

However, heavy-tailed nature of several operational parameters makes prediction-based adaptation by ABR algorithms sub-optimal [ABR_2]. This is because with such distributions, law of large numbers (how long does it take for sample mean to stabilize) works too slowly [HEAVY_TAIL_2], the mean of sample does not equal the mean of distribution [HEAVY_TAIL_2], and as a result standard deviation and variance are unsuitable as metrics for such operational parameters [HEAVY_TAIL_1]. Other subtle issues with these distributions include the "expectation paradox" [HEAVY_TAIL_1] where the longer the wait for an event, the longer a further need to wait and the issue of mismatch between the size and count of events [HEAVY_TAIL_1]. This makes designing an algorithm for adaptation error-prone and challenging. Such operational parameters include but are not limited to buffer occupancy, throughput, client-server latency, and variable transmission times. In addition, edge devices and communication links may fail and logical communication relationships between various software components change frequently as the user moves around with their XR device [UBICOMP].

4. XR Network Traffic

4.1. Traffic Workload

As discussed earlier, the parameters that capture the characteristics of XR application behavior are heavy-tailed. Examples of such parameters include the distribution of arrival times between XR application invocation, the amount of data transferred, and the inter-arrival times of packets within a session. As a result, any traffic model based on such parameters are themselves heavy-tailed. Using these models to predict performance under alternative resource allocations by the network operator is challenging. For example, both uplink and downlink traffic to a user device has parameters such as volume of XR data, burst time, and idle time that are heavy-tailed.

Table 1 below shows various streaming video applications and their associated throughput requirements [METRICS_1]. Since our use case envisages a 6 degrees of freedom (6DoF) video or point cloud, it can be seen from the table that it will require 200 to 1000Mbps of bandwidth. As seen from the table, the XR application such as our use case transmit a larger amount of data per unit time as compared to traditional video applications. As a result, issues arising out of heavy-tailed parameters such as long-range dependent traffic [METRICS_2], self-similar traffic [METRICS_3], would be experienced at time scales of milliseconds and microseconds rather than hours or seconds. Additionally, burstiness at the time scale of tens of milliseconds due to multi-fractal spectrum of traffic will be experienced [METRICS_4]. Long-range dependent traffic can have long bursts and various traffic parameters from widely separated time can show correlation [HEAVY_TAIL_1]. Self-similar traffic contains bursts at a wide range of time scales [HEAVY_TAIL_1]. Multi-fractal spectrum bursts for traffic summarizes the statistical distribution of local scaling exponents found in a traffic trace [HEAVY_TAIL_1]. The operational consequences of XR traffic having characteristics such as long-range dependency, and self-similarity is that the edge servers to which multiple XR devices are connected wirelessly could face long bursts of traffic [METRICS_2], [METRICS_3]. In addition, multi-fractal spectrum burstiness at the scale of milli-seconds could induce jitter contributing to motion sickness [METRICS_4]. This is because bursty traffic combined with variable queueing delays leads to large delay jitter [METRICS_4]. The operators of edge servers will need to run a 'managed edge cloud service' [METRICS_5] to deal with the above problems. Functionalities that such a managed edge cloud service could operationally provide include dynamic placement of XR servers, mobility support and energy management [METRICS_6]. Providing Edge server support for the techniques being developed at the DETNET Working Group at the IETF [RFC8939], [RFC9023], [RFC9450] could guarantee performance of XR applications. For example, these techniques could be used for the link between the XR device and the edge as well as within the managed edge cloud service. Another option for the network operators could be to deploy equipment that supports differentiated services [RFC2475] or per-connection quality-of-service guarantees [RFC2210].

Application	Throughput Required
Real-world objects annotated with text and images for workflow assistance (e.g. repair)	1 Mbps
Video Conferencing	2 Mbps
3D Model and Data Visualization	2 to 20 Mbps
Two-way 3D Telepresence	5 to 25 Mbps
Current-Gen 360-degree video (4K)	10 to 50 Mbps
Next-Gen 360-degree video (8K, 90+ Frames-per-second, High Dynamic Range, Stereoscopic)	50 to 200 Mbps
6 Degree of Freedom Video or Point Cloud	200 to 1000 Mbps

Table 1: Throughput requirement for streaming video applications

Thus, the provisioning of edge servers in terms of the number of servers, the topology, where to place them, the assignment of link capacity, CPUs and GPUs should keep the above factors in mind.

4.2. Traffic Performance Metrics

The performance requirements for XR traffic have characteristics that need to be considered when operationalizing a network. These characteristics are now discussed.

The bandwidth requirements of XR applications are substantially higher than those of video-based applications.

The latency requirements of XR applications have been studied recently [XR_TRAFFIC]. The following characteristics were identified.:

* The uploading of data from an XR device to a remote server for processing dominates the end-to-end latency.

* A lack of visual features in the grid environment can cause increased latencies as the XR device uploads additional visual data for processing to the remote server.

* XR applications tend to have large bursts that are separated by significant time gaps.

Additionally, XR applications interact with each other on a time scale of a round-trip-time propagation, and this must be considered when operationalizing a network.

The following Table 2 [METRICS_6] shows a taxonomy of applications with their associated required response times and bandwidths. Response times can be defined as the time interval between the end of a request submission and the end of the corresponding response from a system. If the XR device offloads a task to an edge server, the response time of the server is the round-trip time from when a data packet is sent from the XR device until a response is received. Note that the required response time provides an upper bound on the sum of the time taken by computational tasks such as processing of scenes, generation of images and the round-trip time. This response time depends only on the Quality of Service (QOS) required by an application. The response time is therefore independent of the underlying technology of the network and the time taken by the computational tasks.

Our use case requires a response time of 20ms at most and preferably between 7-15ms as discussed earlier. This requirement for response time is similar to the first two entries of Table 2 below. Additionally, the required bandwidth for our use case as discussed in section 5.1, Table 1, is 200Mbps-1000Mbps. Since our use case envisages multiple users running the XR applications on their devices, and connected to an edge server that is closest to them, these latency and bandwidth connections will grow linearly with the number of users. The operators should match the network provisioning to the maximum number of tourists that can be supported by a link to an edge server.

Application	Required Response Time	Expected Data Capacity	Possible Implementations/ Examples
Mobile XR based remote assistance with uncompressed 4K (1920x1080 pixels) 120 fps HDR 10-bit real-time video stream	Less than 10 milliseconds	Greater than 7.5 Gbps	Assisting maintenance technicians, Industry 4.0 remote maintenance, remote assistance in robotics industry
Indoor and localized outdoor navigation	Less than 20 milliseconds	50 to 200 Mbps	Theme Parks, Shopping Malls, Archaeological Sites, Museum guidance
Cloud-based Mobile XR applications	Less than 50 milliseconds	50 to 100 Mbps	Google Live View, XR-enhanced Google Translate

Table 2: Traffic Performance Metrics of Selected XR Applications

5. Conclusion

In order to operationalize a use case such as the one presented in this document, a network operator could dimension their network to provide a short and high-capacity network path from the edge compute resources or storage to the mobile devices running the XR application. This is required to ensure a response time of 20ms at most and preferably between 7-15ms. Additionally, a bandwidth of 200 to 1000Mbps is required by such applications. To deal with the characteristics of XR traffic as discussed in this document, network operators could deploy a managed edge cloud service that operationally provides dynamic placement of XR servers, mobility support and energy management. Although the use case is technically feasible, economic viability is an important factor that must be considered.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

The security issues for the presented use case are similar to other streaming applications [DIST], [NIST1], [CWE], [NIST2]. This document itself introduces no new security issues.

8. Acknowledgements

Many Thanks to Spencer Dawkins, Rohit Abhishek, Jake Holland, Kiran Makhijani, Ali Begen, Cullen Jennings, Stephan Wenger, Eric Vyncke, Wesley Eddy, Paul Kyzivat, Jim Guichard, Roman Danyliw, Warren Kumari, and Zaheduzzaman Sarker for providing very helpful feedback, suggestions and comments.

9. Informative References

[ABR_1] Mao, H., Netravali, R., and M. Alizadeh, "Neural Adaptive Video Streaming with Pensieve", In Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 197-210, 2017.

[ABR_2] Yan, F., Ayers, H., Zhu, C., Fouladi, S., Hong, J., Zhang, K., Levis, P., and K. Winstein, "Learning in situ: a randomized experiment in video streaming", In 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20), pp. 495-511, 2020.

[AUGMENTED] Schmalstieg, D. S. and T.H. Hollerer, "Augmented Reality", Addison Wesley, 2016.

[AUGMENTED_2] Azuma, R. T., "A Survey of Augmented Reality.", Presence:Teleoperators and Virtual Environments 6.4, pp. 355-385., 1997.

[BATT_DRAIN] Seneviratne, S., Hu, Y., Nguyen, T., Lan, G., Khalifa, S., Thilakarathna, K., Hassan, M., and A. Seneviratne, "A survey of wearable devices and challenges.", In IEEE Communication Surveys and Tutorials, 19(4), p.2573-2620., 2017.

[BLUR] Kan, P. and H. Kaufmann, "Physically-Based Depth of Field in Augmented Reality.", In Eurographics (Short Papers), pp. 89-92., 2012.

[CLOUD] Corneo, L., Eder, M., Mohan, N., Zavodovski, A., Bayhan, S., Wong, W., Gunningberg, P., Kangasharju, J., and J. Ott, "Surrounded by the Clouds: A Comprehensive Cloud Reachability Study.", In Proceedings of the Web Conference 2021, pp. 295-304, 2021.

[CWE] "CWE/SANS TOP 25 Most Dangerous Software Errorss", Common Weakness Enumeration, SANS Institute, 2012.

[DEV_HEAT_1] LiKamWa, R., Wang, Z., Carroll, A., Lin, F., and L. Zhong, "Draining our Glass: An Energy and Heat characterization of Google Glass", In Proceedings of 5th Asia-Pacific Workshop on Systems pp. 1-7, 2013.

[DEV_HEAT_2] Matsuhashi, K., Kanamoto, T., and A. Kurokawa, "Thermal model and countermeasures for future smart glasses.", In Sensors, 20(5), p.1446., 2020.

[DIST] Coulouris, G., Dollimore, J., Kindberg, T., and G. Blair, "Distributed Systems: Concepts and Design", Addison Wesley, 2011.

[EDGE_1] Satyanarayanan, M., "The Emergence of Edge Computing", In Computer 50(1) pp. 30-39, 2017.

[EDGE_2] Satyanarayanan, M., Klas, G., Silva, M., and S. Mangiante, "The Seminal Role of Edge-Native Applications", In IEEE International Conference on Edge Computing (EDGE) pp. 33-40, 2019.

[EDGE_3] Peterson, L. and O. Sunay, "5G mobile networks: A systems approach.", In Synthesis Lectures on Network Systems., 2020.

[GLB_ILLUM_1] Kan, P. and H. Kaufmann, "Differential irradiance caching for fast high-quality light transport between virtual and real worlds.", In IEEE International Symposium on Mixed and Augmented Reality (ISMAR),pp. 133-141, 2013.

[GLB_ILLUM_2] Franke, T., "Delta voxel cone tracing.", In IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 39-44, 2014.

[HEAVY_TAIL_1]
Crovella, M. and B. Krishnamurthy, "Internet measurement: infrastructure, traffic and applications", John Wiley and Sons Inc., 2006.

[HEAVY_TAIL_2]
Taleb, N., "The Statistical Consequences of Fat Tails", STEM Academic Press, 2020.

[HEAVY_TAIL_3]
Ehrenberg, A., "A Primer in Data Reduction.", John Wiley, London, 1982.

[LENS_DIST]
Fuhrmann, A. and D. Schmalstieg, "Practical calibration procedures for augmented reality.", In Virtual Environments 2000, pp. 3-12. Springer, Vienna, 2000.

[METRICS_1]
ABI Research, "Augmented and Virtual Reality: The first Wave of Killer Apps.", https://gsacom.com/paper/augmented-virtual-reality-first-wave-5g-killer-apps-qualcomm-abi-research/, 2017.

[METRICS_2]
Paxon, V. and S. Floyd, "Wide Area Traffic: The Failure of Poisson Modelling.", In IEEE/ACM Transactions on Networking, pp. 226-244., 1995.

[METRICS_3]
Willinger, W., Taqqu, M.S., Sherman, R., and D.V. Wilson, "Self-Similarity Through High Variability: Statistical Analysis and Ethernet LAN Traffic at Source Level.", In IEEE/ACM Transactions on Networking, pp. 71-86., 1997.

[METRICS_4]
Gilbert, A.C., "Multiscale Analysis and Data Networks.", In Applied and Computational Harmonic Analysis, pp. 185-202., 2001.

[METRICS_5]
Beyer, B., Jones, C., Petoff, J., and N.R. Murphy, "Site Reliability Engineering: How Google Runs Production Systems.", O’Reilly Media, Inc., 2016.

[METRICS_6]
Siriwardhana, Y., Porambage, P., Liyanage, M., and M. Ylianttila, "A survey on mobile augmented reality with 5G

mobile edge computing: architectures, applications, and technical aspects.", In IEEE Communications Surveys and Tutorials, Vol 23, No. 2, 2021.

[NIST1] "NIST SP 800-146: Cloud Computing Synopsis and Recommendations", National Institute of Standards and Technology, US Department of Commerce, 2012.

[NIST2] "NIST SP 800-123: Guide to General Server Security", National Institute of Standards and Technology, US Department of Commerce, 2008.

[NOISE] Fischer, J., Bartz, D., and W. StraÃ\237er, "Enhanced visual realism by incorporating camera image effects.", In IEEE/ACM International Symposium on Mixed and Augmented Reality, pp. 205-208., 2006.

[OCCL_1] Breen, D.E., Whitaker, R.T., and M. Tuceryan, "Interactive Occlusion and automatic object placementfor augmented reality", In Computer Graphics Forum, vol. 15, no. 3 , pp. 229-238,Edinburgh, UK: Blackwell Science Ltd, 1996.

[OCCL_2] Zheng, F., Schmalstieg, D., and G. Welch, "Pixel-wise closed-loop registration in video-based augmented reality", In IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 135-143, 2014.

[OCCL_3] Lang, B., "Oculus Shares 5 Key Ingredients for Presence in Virtual Reality.", https://www.roadtovr.com/oculus-shares-5-key-ingredients-for-presence-in-virtual-reality/, 2014.

[PER_SENSE] Mania, K., Adelstein, B.D., Ellis, S.R., and M.I. Hill, "Perceptual sensitivity to head tracking latency in virtual environments with varying degrees of scene complexity.", In Proceedings of the 1st Symposium on Applied perception in graphics and visualization pp. 39-47., 2004.

[PHOTO_REG] Liu, Y. and X. Granier, "Online tracking of outdoor lighting variations for augmented reality with moving cameras", In IEEE Transactions on visualization and computer graphics, 18(4), pp.573-580, 2012.

[PREDICT] Buker, T. J., Vincenzi, D.A., and J.E. Deaton, "The effect of apparent latency on simulator sickness while using a see-through helmet-mounted display: Reducing apparent latency with predictive compensation..", In Human factors 54.2, pp. 235-249., 2012.

[REG] Holloway, R. L., "Registration error analysis for augmented reality.", In Presence:Teleoperators and Virtual Environments 6.4, pp. 413-432., 1997.

[RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, DOI 10.17487/RFC2210, September 1997, <https://www.rfc-editor.org/info/rfc2210>.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <https://www.rfc-editor.org/info/rfc2475>.

[RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <https://www.rfc-editor.org/info/rfc8939>.

[RFC9023] Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP over IEEE 802.1 Time-Sensitive Networking (TSN)", RFC 9023, DOI 10.17487/RFC9023, June 2021, <https://www.rfc-editor.org/info/rfc9023>.

[RFC9450] Bernardos, CJ., Ed., Papadopoulos, G., Thubert, P., and F. Theoleyre, "Reliable and Available Wireless (RAW) Use Cases", RFC 9450, DOI 10.17487/RFC9450, August 2023, <https://www.rfc-editor.org/info/rfc9450>.

[SLAM_1] Ventura, J., Arth, C., Reitmayr, G., and D. Schmalstieg, "A minimal solution to the generalized pose-and-scale problem", In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 422-429, 2014.

[SLAM_2] Sweeny, C., Fragoso, V., Hollerer, T., and M. Turk, "A scalable solution to the generalized pose and scale problem", In European Conference on Computer Vision, pp. 16-31, 2014.

[SLAM_3] Gauglitz, S., Sweeny, C., Ventura, J., Turk, M., and T. Hollerer, "Model estimation and selection towards unconstrained real-time tracking and mapping", In IEEE transactions on visualization and computer graphics, 20(6), pp. 825-838, 2013.

[SLAM_4] Pirchheim, C., Schmalstieg, D., and G. Reitmayr, "Handling pure camera rotation in keyframe-based SLAM", In 2013 IEEE international symposium on mixed and augmented reality (ISMAR), pp. 229-238, 2013.

[UBICOMP] Bardram, J. and A. Friday, "Ubiquitous Computing Systems", In Ubiquitous Computing Fundamentals pp. 37-94. CRC Press, 2009.

[URLLC] 3GPP, "3GPP TR 23.725: Study on enhancement of Ultra-Reliable Low-Latency Communication (URLLC) support in the 5G Core network (5GC).", https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3453, 2019.

[VIS_INTERFERE] Kalkofen, D., Mendez, E., and D. Schmalstieg, "Interactive focus and context visualization for augmented reality.", In 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, pp. 191-201., 2007.

[XR] 3GPP, "3GPP TR 26.928: Extended Reality (XR) in 5G.", https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3534, 2020.

[XR_TRAFFIC] Apicharttrisorn, K., Balasubramanian, B., Chen, J., Sivaraj, R., Tsai, Y., Jana, R., Krishnamurthy, S., Tran, T., and Y. Zhou, "Characterization of Multi-User Augmented Reality over Cellular Networks", In 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON), pp. 1-9. IEEE, 2020.

Authors’ Addresses

Renan Krishna
United Kingdom
Email: renan.krishna@gmail.com

Akbar Rahman
Ericsson
349 Terry Fox Drive
Ottawa Ontario K2K 2V6
Canada
Email: Akbar.Rahman@ericsson.com