
Stream: Internet Engineering Task Force (IETF)
RFC: [9699](#)
Category: Informational
Published: December 2024
ISSN: 2070-1721
Authors: R. Krishna A. Rahman
Ericsson

RFC 9699

Use Case for an Extended Reality Application on Edge Computing Infrastructure

Abstract

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

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9699>.

Copyright Notice

Copyright (c) 2024 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Use Case	4
2.1. Processing of Scenes	4
2.2. Generation of Images	5
3. Technical Challenges and Solutions	5
4. XR Network Traffic	6
4.1. Traffic Workload	6
4.2. Traffic Performance Metrics	8
5. Conclusion	9
6. IANA Considerations	9
7. Security Considerations	10
8. Informative References	10
Acknowledgements	15
Authors' Addresses	15

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 along a continuum that connects a 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 have several requirements for the network and the mobile devices running these applications. Some XR applications (such as AR applications) require real-time processing of video streams to recognize specific objects. This processing is then used to overlay information on the video being displayed to the user. In addition, other XR applications (such as AR and VR applications) 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 following can be deployed:

Disaggregation: Using Software-Defined Networking (SDN) to break vertically integrated systems into independent components. These components can have open interfaces that are standard, well documented, and non-proprietary.

Virtualization: Being able to run multiple independent copies of those components, such as SDN Controller applications and 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 that require 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 a 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. For the purpose of this document, a network operator is any organization or individual that manages or operates the computing resources or storage in close network proximity to a mobile device or sensor. Examples of form factors include the following: 1) head-mounted displays (HMDs), such as optical see-through HMDs and video see-through HMDs, 2) hand-held displays, and 3) smartphones with video cameras and location-sensing capabilities using systems such as a global navigation satellite system (GNSS). These devices have limited battery capacity and dissipate heat when running. Also, 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 Bitrate (ABR) 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 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 examines these issues with the use case presented in the following section.

2. Use Case

This use case involves an XR application running on a mobile device. Consider a group of tourists who are taking 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 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 are discussed in greater detail below.

2.1. Processing of Scenes

The task of processing a scene can be broken down into a pipeline of three consecutive subtasks: tracking, acquisition of a model of the real world, and 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 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 a model based on an annotated point cloud (a set of points in space that are annotated with descriptors) that is then stored in a database. To ensure that this database can be scaled up, techniques such as combining client-side simultaneous tracking and mapping with 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 a 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 onto an object's surface. A third modeling 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 a 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; 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 real objects [OCCL_1] [OCCL_2]. The XR application also applies photometric registration [PHOTO_REG] by aligning 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 are also required to deal with artifacts generated by lens distortion [LENS_DIST], blur [BLUR], noise [NOISE], etc.

2.2. Generation of Images

The XR application must generate a high-quality video that has the properties described above and overlay the video on the XR device's display. This step is 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 that are computationally intensive, such as real-time generation and processing of high-quality video content. This section discusses the challenges such applications can face as a consequence and offers some solutions.

As a result of performing computationally intensive tasks on XR devices such as XR glasses, excessive heat is generated by the chipsets 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 problem of heat dissipation and battery drainage 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 that this delay can be at most 20 ms and preferably between 7-15 ms in order to avoid motion sickness [PER_SENSE] [XR] [OCCL_3]. Out of these 20 ms, display techniques including the refresh rate of write displays and pixel switching take 12-13 ms [OCCL_3] [CLOUD]. This leaves 7-8 ms 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 suggests an Ultra-Reliable Low Latency of 0.1 to 1 ms 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. For example, a sudden surge in demand from a large group of tourists can overwhelm the 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 ABR algorithms that choose bitrate policies tailored in a fine-grained manner to the resource demands and play back the videos with appropriate QoE metrics as the user moves around with the group of tourists.

However, the heavy-tailed nature of several operational parameters (e.g., buffer occupancy, throughput, client-server latency, and variable transmission times) makes prediction-based adaptation by ABR algorithms sub-optimal [ABR_2]. This is because with such distributions, the law of large numbers (how long it takes for the sample mean to stabilize) works too slowly [HEAVY_TAIL_2] and the mean of sample does not equal the mean of distribution [HEAVY_TAIL_2]; 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] (the longer the wait for an event, the longer a further need to wait) and the mismatch between the size and count of events [HEAVY_TAIL_1]. These issues make designing an algorithm for adaptation error-prone and challenging. 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 in Sections 1 and 3, 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 invocations, 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 is also 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, the table indicates that it will require 200 to 1000 Mbps of bandwidth. Also, the table shows that XR applications, such as the one in our use case, transmit a larger amount of data per unit time as compared to regular video applications. As a result, issues arising from heavy-tailed parameters, such as long-range dependent traffic [METRICS_2] and self-similar traffic [METRICS_3], would be experienced at timescales of milliseconds and microseconds rather than hours or seconds. Additionally, burstiness at the timescale of tens of milliseconds due to the 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 times can show correlation [HEAVY_TAIL_1]. Self-similar traffic contains bursts at a wide range of timescales [HEAVY_TAIL_1]. Multi-fractal spectrum bursts for traffic summarize the statistical distribution of local scaling exponents found in a traffic trace [HEAVY_TAIL_1]. The operational consequence 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 milliseconds 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 support for edge servers in techniques such as those described in [RFC8939], [RFC9023], and [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 network operators could be to deploy equipment that supports differentiated services [RFC2475] or per-connection Quality-of-Service (QoS) guarantees using RSVP [RFC2210].

Thus, the provisioning of edge servers (in terms of the number of servers, the topology, the placement of servers, the assignment of link capacity, CPUs, and Graphics Processing Units (GPUs)) should be performed with the above factors in mind.

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

Application	Throughput Required
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
6DoF video or point cloud	200 to 1000 Mbps

Table 1: Throughput Requirements for Streaming Video Applications

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 discussed in this section.

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 timescale of an RTT propagation, and this must be considered when operationalizing a network.

Table 2 shows a taxonomy of applications with their associated required response times and bandwidths (this data is from Table V in [METRICS_6]). 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 RTT 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 for the sum of the time taken by computational tasks (such as processing of scenes and generation of images) and the RTT. This response time depends only on the 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 20 ms at most and preferably between 7-15 ms, as discussed earlier. This requirement for response time is similar to the first two entries in Table 2. Additionally, the required bandwidth for our use case is 200 to 1000 Mbps (see Section 4.1). Since

our use case envisages multiple users running the XR application on their devices and connecting to the 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	Guidance in theme parks, shopping malls, archaeological sites, and museums
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 computing resources or storage to the mobile devices running the XR application. This is required to ensure a response time of 20 ms at most and preferably between 7-15 ms. Additionally, a bandwidth of 200 to 1000 Mbps 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 those described in [DIST], [NIST1], [CWE], and [NIST2]. This document does not introduce any new security issues.

8. Informative References

- [ABR_1] Mao, H., Netravali, R., and M. Alizadeh, "Neural Adaptive Video Streaming with Pensieve", SIGCOMM '17: Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 197-210, DOI 10.1145/3098822.3098843, 2017, <<https://dl.acm.org/doi/10.1145/3098822.3098843>>.
- [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", 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI '20), pp. 495-511, February 2020, <<https://www.usenix.org/conference/nsdi20/presentation/yan>>.
- [AUGMENTED] Schmalstieg, D. and T. Höllerer, "Augmented Reality: Principles and Practice", Addison-Wesley Professional, 2016, <<https://www.oreilly.com/library/view/augmented-reality-principles/9780133153217/>>.
- [AUGMENTED_2] Azuma, R.T., "A Survey of Augmented Reality", Presence: Teleoperators and Virtual Environments, vol. 6, no. 4, pp. 355-385, DOI 10.1162/pres.1997.6.4.355, August 1997, <<https://direct.mit.edu/pvar/article-abstract/6/4/355/18336/A-Survey-of-Augmented-Reality?redirectedFrom=fulltext>>.
- [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", IEEE Communication Surveys and Tutorials, vol. 19 no. 4, pp. 2573-2620, DOI 10.1109/COMST.2017.2731979, 2017, <<https://ieeexplore.ieee.org/document/7993011>>.
- [BLUR] Kan, P. and H. Kaufmann, "Physically-Based Depth of Field in Augmented Reality", Eurographics 2012 - Short Papers, pp. 89-92, DOI 10.2312/conf/EG2012/short/089-092, 2012, <<https://diglib.org/items/6954bf7e-5852-44cf-8155-4ba269dc4cee>>.
- [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", WWW '21: Proceedings of the Web Conference 2021, pp. 295-304, DOI 10.1145/3442381.3449854, 2021, <<https://dl.acm.org/doi/10.1145/3442381.3449854>>.
- [CWE] SANS Institute, "CWE/SANS TOP 25 Most Dangerous Software Errors", <<https://www.sans.org/top25-software-errors/>>.

- [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", APSys '14: 5th Asia-Pacific Workshop on Systems, pp. 1-7, DOI 10.1145/2637166.2637230, 2014, <<https://dl.acm.org/doi/10.1145/2637166.2637230>>.
- [DEV_HEAT_2]** Matsushashi, K., Kanamoto, T., and A. Kurokawa, "Thermal Model and Countermeasures for Future Smart Glasses", Sensors, vol. 20, no. 5, p. 1446, DOI 10.3390/s20051446, 2020, <<https://www.mdpi.com/1424-8220/20/5/1446>>.
- [DIST]** Coulouris, G., Dollimore, J., Kindberg, T., and G. Blair, "Distributed Systems: Concepts and Design", Addison-Wesley, 2011, <<https://dl.acm.org/doi/10.5555/2029110>>.
- [EDGE_1]** Satyanarayanan, M., "The Emergence of Edge Computing", Computer, vol. 50, no. 1, pp. 30-39, DOI 10.1109/MC.2017.9, 2017, <<https://ieeexplore.ieee.org/document/7807196>>.
- [EDGE_2]** Satyanarayanan, M., Klas, G., Silva, M., and S. Mangiante, "The Seminal Role of Edge-Native Applications", 2019 IEEE International Conference on Edge Computing (EDGE), pp. 33-40, DOI 10.1109/EDGE.2019.00022, 2019, <<https://ieeexplore.ieee.org/document/8812200>>.
- [EDGE_3]** Peterson, L. and O. Sunay, "5G Mobile Networks: A Systems Approach", Synthesis Lectures on Network Systems, DOI 10.1007/978-3-031-79733-0, 2020, <<https://link.springer.com/book/10.1007/978-3-031-79733-0>>.
- [GLB_ILLUM_1]** Kan, P. and H. Kaufmann, "Differential Irradiance Caching for fast high-quality light transport between virtual and real worlds", 2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 133-141, DOI 10.1109/ISMAR.2013.6671773, 2013, <<https://ieeexplore.ieee.org/document/6671773>>.
- [GLB_ILLUM_2]** Franke, T., "Delta Voxel Cone Tracing", 2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 39-44, DOI 10.1109/ISMAR.2014.6948407, 2014, <<https://ieeexplore.ieee.org/document/6948407>>.
- [HEAVY_TAIL_1]** Crovella, M. and B. Krishnamurthy, "Internet Measurement: Infrastructure, Traffic and Applications", John Wiley and Sons, 2006, <<https://www.wiley.com/en-us/Internet+Measurement%3A+Infrastructure%2C+Traffic+and+Applications-p-9780470014615>>.
- [HEAVY_TAIL_2]** Taleb, N., "Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications", Revised Edition, STEM Academic Press, 2022, <<https://arxiv.org/pdf/2001.10488>>.
- [HEAVY_TAIL_3]** Ehrenberg, A., "A Primer in Data Reduction: An Introductory Statistics Textbook", John Wiley and Sons, 2007, <<https://www.wiley.com/en-us/A+Primer+in+Data+Reduction%3A+An+Introductory+Statistics+Textbook-p-9780471101352>>.

-
- [LENS_DIST]** Fuhrmann, A., Schmalstieg, D., and W. Purgathofer, "Practical Calibration Procedures for Augmented Reality", *Virtual Environments 2000*, pp. 3-12, DOI 10.1007/978-3-7091-6785-4_2, 2000, <https://link.springer.com/chapter/10.1007/978-3-7091-6785-4_2>.
- [METRICS_1]** ABI Research, "Augmented and Virtual Reality: The first Wave of Killer Apps: Qualcomm - ABI Research", April 2017, <<https://gsacom.com/paper/augmented-virtual-reality-first-wave-5g-killer-apps-qualcomm-abi-research/>>.
- [METRICS_2]** Paxson, V. and S. Floyd, "Wide area traffic: the failure of Poisson modeling", *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226-244, DOI 10.1109/90.392383, June 1995, <<https://ieeexplore.ieee.org/document/392383>>.
- [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", *IEEE/ACM Transactions on Networking*, vol. 5, no. 1, pp. 71-86, DOI 10.1109/90.554723, February 1997, <<https://ieeexplore.ieee.org/abstract/document/554723>>.
- [METRICS_4]** Gilbert, A.C., "Multiscale Analysis and Data Networks", *Applied and Computational Harmonic Analysis*, vol. 10, no. 3, pp. 185-202, DOI 10.1006/acha.2000.0342, May 2001, <<https://www.sciencedirect.com/science/article/pii/S1063520300903427>>.
- [METRICS_5]** Beyer, B., Ed., Jones, C., Ed., Petoff, J., Ed., and N.R. Murphy, Ed., "Site Reliability Engineering: How Google Runs Production Systems", O'Reilly Media, Inc., 2016, <<https://research.google/pubs/site-reliability-engineering-how-google-runs-production-systems/>>.
- [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", *IEEE Communications Surveys and Tutorials*, vol. 23, no. 2, pp. 1160-1192, DOI 10.1109/COMST.2021.3061981, 2021, <<https://ieeexplore.ieee.org/document/9363323>>.
- [NIST1]** NIST, "Cloud Computing Synopsis and Recommendations", NIST SP 800-146, DOI 10.6028/NIST.SP.800-146, May 2012, <<https://csrc.nist.gov/pubs/sp/800/146/final>>.
- [NIST2]** NIST, "Guide to General Server Security", NIST SP 800-123, DOI 10.6028/NIST.SP.800-123, July 2008, <<https://csrc.nist.gov/pubs/sp/800/123/final>>.
- [NOISE]** Fischer, J., Bartz, D., and W. Strasser, "Enhanced visual realism by incorporating camera image effects", 2006 IEEE/ACM International Symposium on Mixed and Augmented Reality, pp. 205-208, DOI 10.1109/ISMAR.2006.297815, 2006, <<https://ieeexplore.ieee.org/document/4079277>>.
- [OCCL_1]** Breen, D.E., Whitaker, R.T., Rose, E., and M. Tuceryan, "Interactive Occlusion and Automatic Object Placement for Augmented Reality", *Computer Graphics Forum*, vol. 15, no. 3, pp. 11-22, DOI 10.1111/1467-8659.1530011, August 1996, <<https://onlinelibrary.wiley.com/doi/10.1111/1467-8659.1530011>>.

-
- [OCCL_2]** Zheng, F., Schmalstieg, D., and G. Welch, "Pixel-wise closed-loop registration in video-based augmented reality", 2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 135-143, DOI 10.1109/ISMAR.2014.6948419, 2014, <<https://ieeexplore.ieee.org/document/6948419>>.
- [OCCL_3]** Lang, B., "Oculus Shares 5 Key Ingredients for Presence in Virtual Reality", Road to VR, 24 September 2014, <<https://www.roadtovr.com/oculus-shares-5-key-ingredients-for-presence-in-virtual-reality/>>.
- [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.", APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization, pp. 39-47, DOI 10.1145/1012551.1012559, 2004, <<https://dl.acm.org/doi/10.1145/1012551.1012559>>.
- [PHOTO_REG]** Liu, Y. and X. Granier, "Online Tracking of Outdoor Lighting Variations for Augmented Reality with Moving Cameras", IEEE Transactions on Visualization and Computer Graphics, vol. 18, no. 4, pp. 573-580, DOI 10.1109/TVCG.2012.53, 2012, <<https://ieeexplore.ieee.org/document/6165138>>.
- [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", Human Factors, vol. 54, no. 2, pp. 235-249, DOI 10.1177/0018720811428734, April 2012, <<https://pubmed.ncbi.nlm.nih.gov/22624290/>>.
- [REG]** Holloway, R.L., "Registration Error Analysis for Augmented Reality", Presence: Teleoperators and Virtual Environments, vol. 6, no. 4, pp. 413-432, DOI 10.1162/pres.1997.6.4.413, August 1997, <<https://direct.mit.edu/pvar/article-abstract/6/4/413/18334/Registration-Error-Analysis-for-Augmented-Reality?redirectedFrom=fulltext>>.
- [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, C.J., 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", 2014 IEEE Conference on Computer Vision and Pattern Recognition, pp. 422-429, DOI 10.1109/CVPR.2014.61, 2014, <<https://ieeexplore.ieee.org/document/6909455>>.
- [SLAM_2]** Sweeny, C., Fragoso, V., Höllerer, T., and M. Turk, "gDLS: A Scalable Solution to the Generalized Pose and Scale Problem", Computer Vision - ECCV 2014, pp. 16-31, DOI 10.1007/978-3-319-10593-2_2, 2014, <https://link.springer.com/chapter/10.1007/978-3-319-10593-2_2>.
- [SLAM_3]** Gauglitz, S., Sweeney, C., Ventura, J., Turk, M., and T. Höllerer, "Model Estimation and Selection towards Unconstrained Real-Time Tracking and Mapping", IEEE Transactions on Visualization and Computer Graphics, vol. 20, no. 6, pp. 825-838, DOI 10.1109/TVCG.2013.243, 2014, <<https://ieeexplore.ieee.org/document/6636302>>.
- [SLAM_4]** Pirchheim, C., Schmalstieg, D., and G. Reitmayr, "Handling pure camera rotation in keyframe-based SLAM", 2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pp. 229-238, DOI 10.1109/ISMAR.2013.6671783, 2013, <<https://ieeexplore.ieee.org/document/6671783>>.
- [UBICOMP]** Bardram, J. and A. Friday, "Ubiquitous Computing Systems", Ubiquitous Computing Fundamentals, 1st Edition, Chapman and Hall/CRC Press, pp. 37-94, 2009, <<https://www.taylorfrancis.com/chapters/edit/10.1201/9781420093612-6/ubiquitous-computing-systems-jakob-bardram-adrian-friday>>.
- [URLLC]** 3GPP, "Study on enhancement of Ultra-Reliable Low-Latency Communication (URLLC) support in the 5G Core network (5GC)", 3GPP TR 23.725, 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3453>>.
- [VIS_INTERFERE]** Kalkofen, D., Mendez, E., and D. Schmalstieg, "Interactive Focus and Context Visualization for Augmented Reality", 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, pp. 191-201, DOI 10.1109/ISMAR.2007.4538846, 2007, <<https://ieeexplore.ieee.org/document/4538846>>.
- [XR]** 3GPP, "Extended Reality (XR) in 5G", 3GPP TR 26.928, 2020, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3534>>.
- [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", 2020 17th Annual IEEE

International Conference on Sensing, Communication, and Networking (SECON), pp. 1-9, DOI 10.1109/SECON48991.2020.9158434, 2020, <<https://ieeexplore.ieee.org/document/9158434>>.

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 helpful feedback, suggestions, and comments.

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