XR and 5G: Extended reality at scale with time-critical communication

The value of 5G extends far beyond the enhanced mobile broadband that more than 1 billion people already have access to through upwards of 150 communication service providers around the world [1]. The time-critical communication capabilities in 5G networks will enable major breakthroughs in a wide range of application areas, including extended reality (XR).

FREDRIK ALRIKSSON, DU HO KANG, CHRIS PHILLIPS, JOSE LUIS PRADAS, ALI ZAIDI

In contrast to enhanced mobile broadband (MBB), the majority of the emerging 5G value-adding applications are time critical in nature with demanding requirements on reliable low latency. Time-critical use cases can be broken down into four categories: real-time media, remote control, industrial control and mobility automation [1, 2].

■ The real-time media category includes innovative extended reality (XR) applications that are of growing interest for consumers, enterprises and public institutions alike. The emergence of time-critical communications over 5G networks makes it possible to offload parts of the XR processing and functionality to the edge cloud and enhance the user experience with lightweight and cost-efficient head-mounted displays (HMDs).

XR use cases

XR is an umbrella term that covers immersive technologies ranging from virtual reality (VR) to mixed reality (MR) and augmented reality (AR). In VR, users are totally immersed in a simulated digital environment or a digital replica of reality. MR includes all variants where virtual and real

environments are mixed. AR is one such variant, where digital information is overlaid on images of reality viewed through a device. The level of augmentation can vary from a simple information display to the addition of virtual objects and even complete augmentation of the real world. MR can also include variants where real objects are included in the virtual world.

XR is expected to improve productivity and convenience for consumers, enterprises and public institutions in a wide variety of application areas such as entertainment, training, education, remote support, remote control, communications and virtual meetings. It can be used in virtually all industry segments, including health care, real estate, shopping, transportation and manufacturing. VR is already used for gaming both at home and at dedicated venues, for virtual tours in the context of real estate, for education and training purposes and for remote participation at live events such as concerts and sports.

While VR holds great promise, AR and MR use cases have even greater transformational potential. In VR, the headsets cut users off from their physical surroundings and restrict mobility [3]. With AR, users are present in reality and free to move even when using HMDs. Many smartphone users have already experienced basic forms of AR, through games like Pokémon Go and apps that enable shoppers to visualize new furniture in their homes before making a purchase. AR technology becomes much more powerful, though, when it is used with HMDs. By freeing up the user's hands, AR HMDs transform the user interaction. The ability to have

●● WHILE VR HOLDS GREAT PROMISE, AR AND MR USE CASES HAVE EVEN GREATER TRANSFORMATIONAL POTENTIAL ●●

information overlaid on the real world while simultaneously having your hands free has been shown to increase worker efficiency dramatically [4].

XR edge-processing architectures

VR HMDs are already available at scale on the market today, but they are rapidly evolving. Because VR applications are often processing-intensive, enabling high-end VR requires connecting the HMDs to high-end processing units – typically powerful PCs or gaming consoles. VR HMDs with local processing capabilities are also starting to become available, but these are relatively large and heavy and cannot provide the same experience as when off-device processing is used.

AR HMDs are also available on the market today, predominantly for enterprise use [5]. Mass-market adoption will require further progress on ease of use, attractive appearance and content availability [6]. Building fashionable, small form factor HMDs to meet the demands for XR is challenging due to limited processing power, storage, battery life and heat dissipation. We believe that the best way to address these challenges is by offloading parts of XR processing to the mobile network edge.

Terms and abbreviations

5GC – 5G Core | **AAS** – Advanced Antenna System | **AR** – Augmented Reality | **CG** – Configured Grant | **CoMP** – Coordinated Multi-Point | **DAPS** – Dual Active Protocol Stack | **DC** – Data Center | **DL** – Downlink | **E2E** – End-to-End | **L4S** – Low Latency, Low Loss, Scalable Throughput | **HMD** – Head-Mounted Display | **MBB** – Mobile Broadband | **MR** – Mixed Reality | **SPS** – Semi-Persistent Scheduling | **TRP** – Transmission Reception Point | **TTI** – Time Transmission Interval | **UE** – User Equipment | **UL** – Uplink | **VR** – Virtual Reality | **XR** – Extended Reality
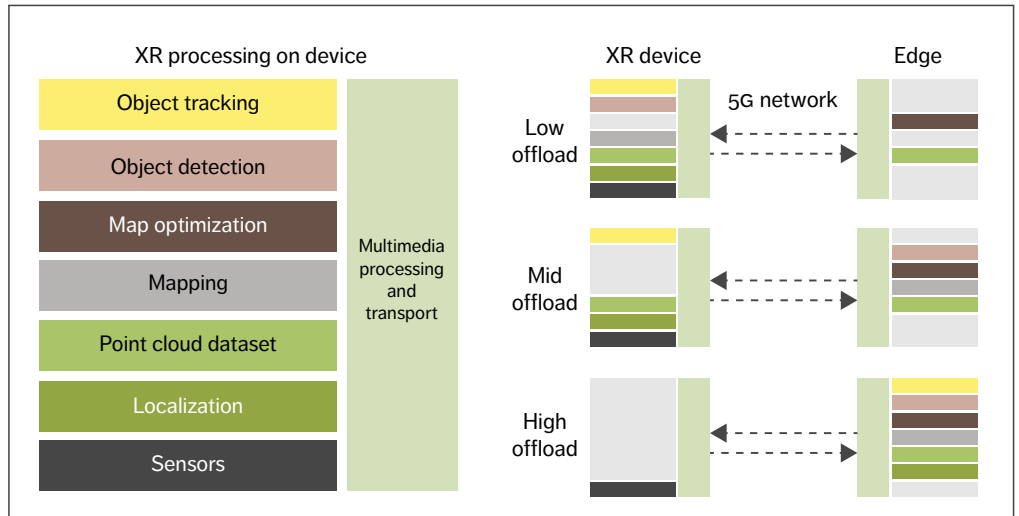
Figure 1 Split architecture options with 5G connectivity

Figure 1 shows the eight main types of XR functionality and how they can be split up between an XR device and the network edge. The major components in XR processing include SLAM (simultaneous localization, mapping and map optimization) [7] with point cloud datasets [8], hand gesture and pose estimation [9], object detection and tracking [10] and multimedia processing and transport. Examples of multimedia processing and transport are rendering, asynchronous time warp [11] and video, audio and sensor encoding.

Figure 1 also illustrates three architectures for splitting the XR processing: low offload, mid offload and high offload. Our internal studies indicate that low offload reduces device energy consumption by

❝❝ OUR INTERNAL STUDIES INDICATE THAT HIGH OFFLOAD REDUCES DEVICE ENERGY CONSUMPTION BY MORE THAN SEVENFOLD ❞❞

threefold while mid offload reduces it by fourfold. High offload reduces device energy consumption by more than sevenfold.

In the low-offload architecture, almost all processing is done on the device. The point cloud dataset, spatial map generation and localization occur on the device. As portions of the point cloud and spatial map are built, the point cloud is compressed through multimedia processing and transmitted over the 5G network to the edge, where uploaded spatial map data is created and merged into existing global spatial map data. The object detection and tracking are performed on the device. The rendering can occur locally on the device or at the network edge.

In the mid-offload case, the localization and object tracking functions are performed on the device. The spatial map and point cloud datasets are generated in the network. Key image frames are compressed using a video codec and sent to the network edge to generate the spatial map and point cloud datasets, and to perform object detection. At the network edge, the point cloud datasets and spatial map are created and merged with the global point cloud

●● THE XR CONNECTIVITY REQUIREMENTS DEPEND ON THE LEVEL OF SPLIT ARCHITECTURE AND THE TARGETED QoE ●●

datasets and spatial map. The overlay rendering occurs at the network edge. The edge-rendered video and audio content is encoded into video and audio streams and transmitted to the device along with the rendering mesh data.

In the high-offload case, only sensor data is sent over the uplink. Sensor data includes camera data. Many AR/VR devices have multiple cameras, including infrared and RGB (red, green and blue) ones. Sensor data also covers sensors such as LiDAR (light detection and ranging) and IMU (inertial measurement unit). The image data is encoded using video compression such as Motion Pictures Experts Group (MPEG) High Efficiency Video Coding (HEVC) or Versatile Video Coding (VVC).

Several compression techniques have been proposed for IMU data, such as delta encoding, linear extrapolation, second- to fifth-order polynomial regression and spline extrapolation [12]. There are numerous techniques for compressing 3D point cloud generated by LiDAR sensors. ISO/MPEG currently has two tracks for point cloud compression standardization under development.

These are Video-based Point Cloud Compression (V-PCC) and Geometry Based Point Cloud Compression (G-PCC). LiDAR is covered in the MPEG G-PCC track [13]. The edge-rendered video and audio content is encoded into video and audio streams and transmitted to the device along with the rendering mesh data.

XR traffic characteristics and connectivity requirements

XR traffic is characterized by a mixture of pose and video from/to the same XR device, varying video frame size over time and quasi-periodic packet arrival with application jitter after IP segmentation. Traffic arrival time to the RAN is periodic with non-negligible jitter due to application-processing-time uncertainty. Video frame sizes are an order of magnitude larger and, at the same time, not fixed over time compared with packets in voice or industrial control communication. The segmentation of each frame is expected, which implies that packets arrive in bursts that must be handled together to meet stringent bounded latency requirements.

XR connectivity requirements depend on the level of split architecture and the targeted QoE, leading to a wide range of bit rates and bounded latency requirements. *Figure 2* presents the 5G connectivity requirements for AR, VR and cloud gaming based on the developments in the ecosystem, including the 3GPP [14]. The requirements assume local processing techniques in the split architecture to

Use cases	DL bitrates (Mbps)	UL bitrates (Mbps)	One-way latency (ms)	Frame realibility (%)
Cloud gaming	8-30	~0.3	10-30	≥99
VR	30-100	< 2	5-20	≥99
AR	2-60	2-20	5-50	≥99

Figure 2 Use-case requirements for 5G networks

mitigate consumer latency requirements [15]. Note that latency and reliability requirements are on a video frame (or file) level excluding application error and delay.

For downlink (DL) video traffic, VR typically needs higher bit rates than cloud gaming to support retinal resolution when using HMD and lower compression efficiency due to low-latency encoding. Some AR applications for conversational services can have DL video traffic as VR. However, they potentially have lower resolution to render a video on only part of display, leading to a lower bit rate than VR DL video. In addition, they can have uplink (UL) video streaming for the object detection and tracking, but the bit rate requirements can be lower compared with the DL video. All cloud gaming, AR and VR include pose traffic in the UL, which has much lower bit rates than video traffic, but AR and VR will have higher bit rate requirements than cloud gaming to convey more pose information, such as six degrees of freedom.

AR and VR require more stringent end-to-end (E2E) bounded latencies than cloud gaming since a human is more sensitive to the discrepancy in 3D virtual environments. For instance, it is widely accepted that rendering motion to photon latency greater than 20ms starts to cause nausea when a human wears a VR HMD. Processing techniques such as asynchronous time warp [11] relax the latency requirement to some extent, making VR feasible over 5G networks. For AR, object detection can be performed at the network edge, and object tracking can be done on the device as shown in the low- and mid-offload cases in Figure 1. By performing object detection in the network and tracking on the device, the bounded latency requirement for a 5G network can be relaxed up to 50ms.

Network architecture for time-critical communications

Time-critical communications is an emerging 5G concept for enabling services with reliable low latency requirements such as XR [2]. The aim is to secure data delivery within specific latency bounds (X ms) with the desired reliability level (Y percent).

Depending on the user requirements, X ranges from tens of milliseconds to 1 millisecond latency and Y ranges from 99 percent to 99.999 percent reliability. To ensure bounded latencies, the system may have to compromise on capacity, throughput, energy efficiency or coverage.

The 5G RAN, the 5G Core (5GC) and the transport network together with the device contribute to the E2E reliability and latency. E2E latency is the sum of individual latency contributions from every component. E2E reliability cannot be better than the reliability of the weakest link.

Edge deployment of 5GC and applications is key to reducing the transport latency between the application and the RAN. If an application is hosted in a central national data center (DC), the transport network round-trip latency can be in the order of 10-40ms, depending on the distance to the DC and how well the transport network is built out. The transport latency can be reduced to 5-20ms by moving applications to a regional DC or even to 1-5ms for edge sites. For local network deployments with networking functions and applications hosted on-premises, transport latencies become negligible [2].

Achievable RAN latency/reliability performance depends on general deployment factors (such as frequency band, bandwidth, inter-site distances, numerology, duplexing schemes and TDD configuration), RAN and user equipment (UE) capabilities (in terms of hardware and software features) and traffic characteristics (such as data rate and packet size).

5G toolbox to achieve time-critical communications

To ensure that XR applications work well over 5G networks, it is important to separate the XR traffic from best-effort MBB traffic using the 5G QoS framework with optimized QoS flows, as illustrated at the top of *Figure 3*. This enables optimized treatment of XR throughout the mobile network and specifically in the RAN to mitigate the different sources of delay.

The provision of bounded latency for XR and

Figure 3 5G toolbox to realize time-critical communications – the tools that are important for XR are marked in green

other services requires the prevention of delays and interruptions. The bottom half of Figure 3 presents the comprehensive 5G toolbox for realizing time-critical communications that makes it possible to overcome the following five main causes of delays and interruptions:

1. Congestion
2. Dynamic radio environment
3. Standards/protocols
4. Mobility
5. Device power saving.

Congestion

There are several ways to mitigate the congestion-related delays that can occur when end hosts transmit at a higher bitrate than the network can sustain. Fast application rate adaptation is key to avoid congestion for rate-adaptive traffic such as XR and cloud gaming. Low Latency, Low Loss, Scalable throughput (L4S) is an existing method that provides fast indication of congestion from networks that can be applied in the RAN, enabling RAN rate recommendation and forming the basis for a common framework for rate adaptation over 5G [16]. Network slicing and radio resource partitioning together with admission control and latency-optimized scheduling are important tools for reserving network resources for time-critical services and avoiding congestion-related delays to provide a minimum guaranteed bit rate, for example. Network slicing is also beneficial for protecting MBB and other services from resource-hungry, time-critical services.

Dynamic radio environment

Advanced scheduling together with robust link adaptation and signal transmission formats are important tools to combat delays related to the

●● THE 5G QoS FRAMEWORK MAKES IT POSSIBLE TO ESTABLISH QoS FLOWS THAT PROVIDE OPTIMIZED NETWORK TREATMENT FOR SPECIFIC FLOWS ●●

dynamic radio environment such as fading/blocking and interference. High-rate XR traffic places new demands on these tools to deliver very high spectral efficiency, while ensuring reliable and timely video frame delivery. Advanced antenna systems (AASs) have tremendous potential to improve the link budget, reduce interference, and increase spatial multiplexing, ultimately leading to more radio system capacity for XR.

Standards/protocols

Features that minimize delays associated with standards/protocols include tools like prescheduling, UL configured grants (CGs) and DL semi-persistent scheduling (SPS). As an example, a combination of periodic CGs and dynamic grants can be used to reduce latency for UL AR traffic consisting of periodic frames of varying size. XR traffic arrival time characteristics call for further optimizations of UL CGs and DL SPS to avoid unnecessary waiting times. Protocol enhancements throughout the different protocol layers, from the SDAP (Service Data Adaptation Protocol) to the physical layer, including the PDCP (Packet Data Convergence Protocol), the RLC (Radio Link Control) and the MAC (Medium Access Control) protocols, can be important for improving the capacity of XR applications as well. For example, control signaling efficiency could be improved to provide grants for multiple radio resource allocations needed for a large XR video frame.

Mobility

Time-critical services place much more stringent requirements on mobility performance than MBB services. For XR services, mobility interruptions need to be well below the inter-frame arrival time (which is typically between 15-50ms) to pass unnoticed. This will require smarter and faster network algorithms as well as stricter processing requirements at the device. 5G New Radio provides a few options for supporting seamless and more robust mobility, including multiple transmission reception points (multi-TRPs), dual active protocol stack (DAPS) handover and conditional handover.

Device power saving

Device power saving is important to support low-power XR devices, and XR traffic arrival time characteristics also call for further optimizations of discontinuous reception.

Traffic awareness in the RAN

One interesting area of future standardization work in the 3GPP is to further optimize 5G performance and capacity by improving XR traffic awareness in the RAN, especially for very short-term traffic variation that may be difficult for an application layer to handle. If, for example, the RAN had the ability to know which IP packets are associated with the same application frame, it could potentially use that knowledge to optimize radio resource allocation, scheduling, link adaptation, packet discarding and other features to increase capacity.

5G QoS approaches for extended reality

The 5G QoS framework makes it possible to establish QoS flows that provide optimized network treatment for specific traffic flows, in addition to the default QoS flow used for MBB. Such additional QoS flows can be established either using 5GC QoS-exposure application programming interfaces to communicate service requirements, or by traffic detection together with pre-provisioned service requirements, such as relying on standardized 5G QoS identifier characteristics. Two complementary 5G QoS approaches can be implemented for XR: a generic latency-optimized QoS flow and service-optimized QoS flows.

RAN deployment	Frequency allocation	Min. bit rate	Max. latency	Frame reliability
Wide area	Mid-band FDD 2x20MHz@2GHz	DL: 8-30Mbps UL: 2-10Mbps	DL: 10-30ms UL: 10-30ms	99%
	Mid-band TDD 100MHz@3.5GHz			
Indoor	Mid-band TDD 100MHz@3.5GHz	DL: 30-60Mbps UL: 10-20Mbps	DL: 10-30ms UL: 10-30ms	99%
	mmWave 800MHz@30GHz		DL: 5-10ms UL: 5-10ms	

Figure 4 Simulation assumptions for 5G RAN

Generic latency-optimized QoS flow

A generic latency-optimized QoS flow that can be implemented in any network is an important tool to enable a large ecosystem of high-rate, rate-adaptive, time-critical applications and services including XR to emerge and hopefully flourish in the same way as MBB and smartphones have done [16]. Implementing such a generic QoS flow enables networks to provide L4S-based early RAN congestion detection and fast rate recommendation to applications as well as packet treatment optimized for low latency and jitter rather than throughput. By not relying on knowledge of specific service requirements it avoids a tight coupling between application and network.

Service-optimized QoS flows

Service-optimized QoS flows can be established for specific XR services with known service requirements in terms of packet delay budgets, packet error rates, minimum guaranteed bit rates and so on, in cases where there is a need to increase coverage or capacity while still providing a good minimum QoE, for example. This requires a tighter coupling between application and network, which adds complexity to the ecosystem but provides QoE beyond what the generic latency-optimized QoS flow can enable and may be required for the most demanding XR applications.

Deployment strategy

To develop insights regarding the 5G network deployment strategy for various XR applications, we have carried out simulation studies for a wide-area deployment and an indoor enterprise deployment. For the wide-area scenario, we evaluated capacity for low- to mid-range XR applications in terms of requirements. For the indoor deployment, we considered high-end applications as well. The simulation assumptions are summarized in *Figure 4*. The 99 percent frame reliability implies less than 0.1 percent packet error rate if each frame is segmented into more than 10 IP packets.

The wide-area scenario is based on a macro deployment in central London with an inter-site distance of approximately 450m. For the mid-band, wide-area deployments, we include an AAS with 32 elements and 16 elements per polarization for 3.5GHz and 2GHz, respectively. The indoor deployment also assumes eight elements and 32 elements per polarization of an AAS for 3.5GHz and 30GHz, respectively. Devices with four receiver
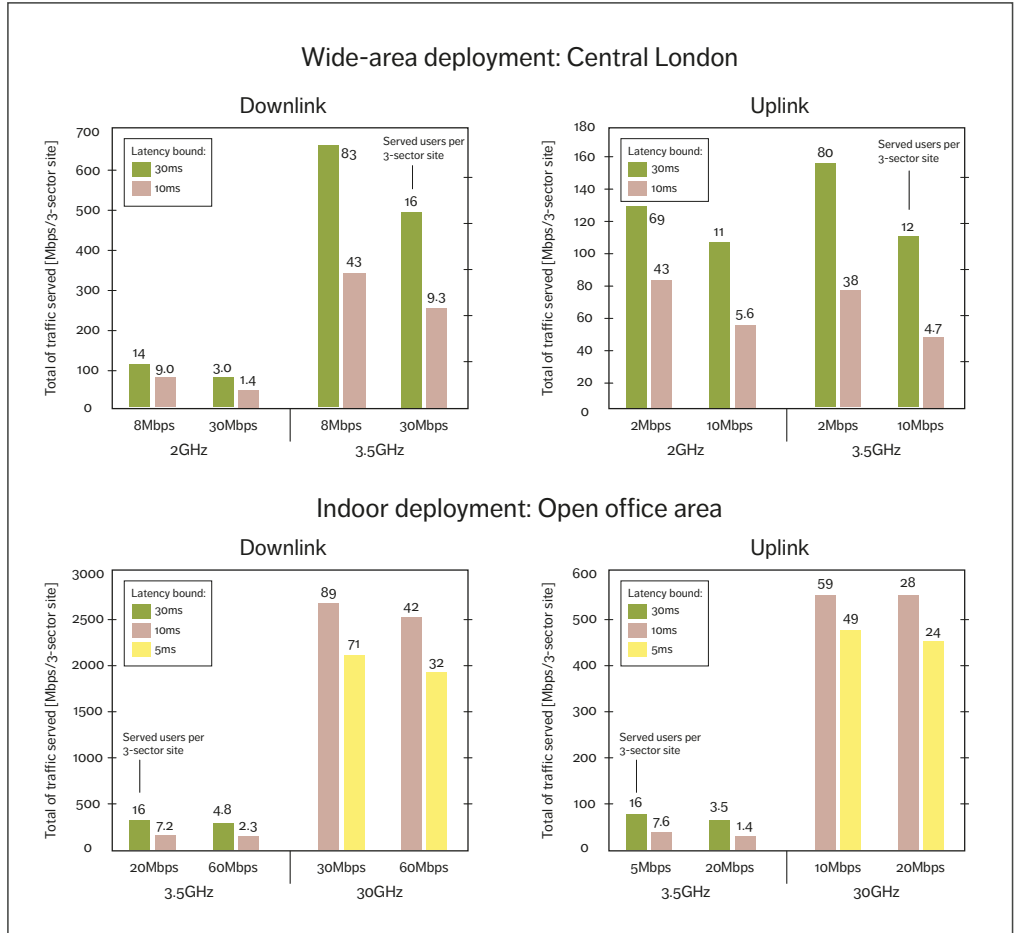
Figure 5 5G RAN-served traffic for various requirements

branches were used in the evaluation. The indoor deployment assumes a 120m by 50m open office area, where four pico sites with three sectors each are placed on the ceiling. For the TDD bands, we have assumed 4:1 DL and UL configuration.

Figure 5 shows the average served traffic per cell (with different combinations of the minimum bit rate and maximum latency requirements per user) for the wide-area and the indoor deployment scenarios assuming 90 percent and 95 percent coverage availability, respectively. From the cell throughput and the minimum bit rate per user, it is possible to derive the maximum number of users served simultaneously per cell, which is shown next to each bar in Figure 5. When the cell capacity is underutilized, the applications can increase the bit rate for superior QoE by rate adaptation.

We can observe that the cell capacity in the traffic

served decreases by increasing the minimum user bit rate requirements or by reducing the maximum latency target. Increasing the user bit rate requirement at a given latency budget creates more interference and resource utilization, leading to a smaller amount of total traffic being served. Similarly, when reducing the latency requirement for a given bit rate requirement, a more stringent latency target implies that the resources are available on a shorter time scale, which leads to fewer users being served.

These results show that communication service providers can start to address cloud gaming and low-end AR use cases (for consumers and enterprises) in a wide area with the ongoing 5G network rollout utilizing mid bands, and gradually evolve capabilities and densify deployments to address greater coverage and more demanding requirements in terms of throughput and bounded latency. For users with coverage needs in a small geographic area such

as a theme park, industry campus or office, there is an opportunity to support high-end XR with indoor deployments and address more stringent latency and higher bit rate requirements with a local breakout of the core network.

Conclusion

Although extended reality (XR) has the potential to be transformational for both business and society, widespread adoption has previously been hindered by issues such as heat generation and the limited processing power, storage and battery life of small form factor head-mounted devices. The time-critical communication capabilities in 5G make it possible to overcome these challenges by offloading XR processing to the mobile network edge. By capitalizing on their ongoing 5G rollouts, mobile network operators are in an excellent position to enable the realization of XR on a large scale.

Further reading

» **Switch on a better 5G network, available at:** *https://www.ericsson.com/en/5g/5g-networks*

» **5G by Ericsson, available at:** *https://www.ericsson.com/en/5g*

» **How 5G and Edge Computing can enhance virtual reality, available at:**
  https://www.ericsson.com/en/blog/2020/4/how-5g-and-edge-computing-can-enhance-virtual-reality

» **A technical overview of time-critical communication with 5G NR, available at:**
  https://www.ericsson.com/en/blog/2021/2/time-critical-communication--5g-nr

References

1. **Ericsson Mobility Report, November 2020, available at:**
   https://www.ericsson.com/4adc87/assets/local/mobility-report/documents/2020/november-2020-ericsson-mobility-report.pdf

2. **Critical IoT connectivity: Ideal for time-critical communications, Ericsson Technology Review, June 2, 2020, Alriksson, F; Boström, L; Sachs, J; Eric Wang, Y.-P; Zaidi, A, available at:**
   https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/critical-iot-connectivity

3. **Merged Reality: Understanding how virtual and augmented realities could transform everyday reality, Ericsson ConsumerLab, June 2017, available at:**
   https://www.ericsson.com/en/reports-and-papers/consumerlab/reports/merged-reality

4. **Augmented Reality Is Already Improving Worker Performance, Harvard Business Review, March 13, 2017, Abraham, M; Annunziata, M, available at:**
   https://hbr.org/2017/03/augmented-reality-is-already-improving-worker-performance

5. **Augmented Reality Headsets Market, Industry Report, 2019-2025, Grand View Research, November 2019, available at:** *https://www.grandviewresearch.com/industry-analysis/augmented-reality-ar-headsets-market*

6. **Is 2021 finally the year for smart glasses? Here's why some experts still say no, CNBC Evolve, January 23, 2021, Subin, S, available at:**
   https://www.cnbc.com/2021/01/23/why-experts-dont-expect-smart-glasses-to-surge-in-2021.html

7. **AR Products Enhance its Accuracy Through SLAM Technology, PS Marketing Intelligence, February 24, 2021, Shrivastava, D, available at:**
   https://psmarketing.home.blog/2021/02/24/ar-products-enhance-its-accuracy-through-slam-technology/

8. **Point clouds and VR: The future of point cloud visualisation, Vercator blog, April 21, 2020, Thomson, C, available at:** *https://info.vercator.com/blog/point-clouds-and-vr-the-future-of-point-cloud-visualisation*

9. **Pose Estimation Guide, Fritz AI, 2021, available at:** *https://www.fritz.ai/pose-estimation/*

10. **Object Detection and Pose Tracking for Augmented Reality: Recent Approaches, HAL, 18th Korea-Japan Joint Workshop on Frontiers of Computer Vision (FCV), February 2012, Uchiyama, H; Marchand, E, available at:** *https://hal.inria.fr/hal-00751704/document*

11. **The asynchronous time warp for virtual reality on consumer hardware, Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology (pp. 37-46), 2016, van Waveren, J.M.P, available at:** *https://dl.acm.org/doi/pdf/10.1145/2993369.2993375*

12. **Lossless Compression of Human Movement IMU Signals, Sensors (Basel), October 2020, Chiasson, D; Xu, J; Shull, P, available at:** *https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7590134/*

13. **3D Point Cloud Compression: A Survey, Web3D '19, Los Angeles, CA, USA, July 26–28, 2019, Chao Cao, C; Preda, M; Zaharia, T, available at:** *https://dl.acm.org/doi/pdf/10.1145/3329714.3338130*

14. **Study on XR Evaluations for NR (RP-193241), 3GPP TSG RAN Meeting #86, 2019, available at:**
   https://www.3gpp.org/ftp/TSG_RAN/TSG_RAN/TSGR_86/Docs/RP-193241.zip

15. **Extended Reality (XR) in 5G (Release 16) TR 26.928, V16.1.0, 3GPP Technical Specification Group Services and System Aspects, December 23, 2020, available at:**
   https://www.3gpp.org/ftp/Specs/archive/26_series/26.928/26928-g10.zip

16. **Enabling time-critical applications over 5G with rate adaptation, Ericsson white paper, 2021, available at:**
   https://www.ericsson.com/en/reports-and-papers/white-papers/enabling-time-critical-applications-over-5g-with-rate-adaptation

THE AUTHORS

Fredrik Alriksson
◆ is a researcher at Development Unit Networks, where he leads strategic technology and concept development within IoT & New Industries. He joined Ericsson in 1999 and has worked in R&D with architecture evolution, covering a broad set of technology areas including RAN, Core, IMS and VoLTE. Alriksson holds an M.Sc. in electrical engineering from KTH Royal Institute of Technology in Stockholm, Sweden.

Du Ho Kang
◆ is a senior specialist at Ericsson Research who joined the company in 2014. His expertise is in 5G-and-beyond concept developments and performance evaluation toward diverse international standardization and spectrum regulation bodies including 3GPP RAN, ETSI BRAN (European Telecommunications Standards Institute Broadband Radio Access Networks) and the ITU-R (International Telecommunication Union – Radiocommunication Sector). His current interest is developing future RAN concepts for emerging services. Kang holds a Ph.D. in wireless infrastructure and deployment from KTH Royal Institute of Technology.

Chris Phillips
◆ is a master researcher at Ericsson Research and the technical lead for the company's internal XR research project. He has been with Ericsson since 2007. His primary area of expertise is video processing and transport optimization. His latest focus has been in foveated remote rendering for VR, 360 video, cloud gaming and processing/transport optimization for distributed spatial mapping with point cloud datasets. He is also active in the 3GPP SA4, VRIF, SVA and OpenXR organizations. Phillips holds an M.Sc. in computer science from the University of Georgia, USA.

Jose Luis Pradas
◆ is a master researcher at Ericsson Research whose current focus is on RAN enhancements to support XR services in 5G networks. He joined Ericsson in 2007 and has worked in research, performing concept development in architecture and RAN protocols as well as driving 3GPP standardization. Pradas hold an M.Sc. in telecommunication from Universitat Politècnica de València, Spain, as well as an M.Sc. in communications from Helsinki University of Technology, Finland.

Ali Zaidi
◆ is a strategic product manager for Cellular IoT at Ericsson and also serves as the company's head of IoT Competence. Since joining Ericsson in 2014, Zaidi has been working with technology and business development of 4G and 5G radio access. He is currently responsible for Ericsson's radio products for time-critical communication, industrial automation, XR and automotive. Zaidi holds an M.Sc. in innovation management and a Ph.D. in telecommunications from KTH Royal Institute of Technology.