Enabling Wearables Towards the Metaverse: Trade-offs involving Blockage and Heat

# Enabling Wearables towards the Metaverse: Trade-Offs Involving Blockage and Heat

Olga Chukhno, Francesco Malandrino, Alessandro Catania, Antonella Molinaro, and Carla Fabiana Chiasserini


## Abstract

Immersive and extended reality (XR) represents one of the most exciting applications of beyond-5G networks, paving the way towards the metaverse. Wearable devices (*wearables* for short) are the core component of XR and, being small and lightweight, face several unique challenges. Most notably, they must handle large amounts of complex computations despite their limited processing power, which can lead to device overheating. This issue becomes more pronounced when heavy computational tasks cannot be offloaded to the edge due to connectivity issues like blockage. In this work, we aim at identifying the best approach for delivering XR services, considering that *both* local and edge-based XR processing can be combined to best cope with challenging operational conditions, most importantly, channel blockage and device overheating. Through numerical experiments based on real-world wearables, we demonstrate that balancing the trade-offs between heat management and blockage mitigation leads to the best performance and quality of experience for XR applications.


## Introduction

Extended reality (XR) applications are a crucial building block towards the metaverse and represent one of the most exciting — and challenging — applications of next-generation mobile networks. They are unique in the multifaceted roles of end-user devices (a.k.a. {*wearables*) as sensors, displays, and computing units. For example, virtual/augmented/mixed reality (VR/AR/MR) gaming applications use a wearable as both a sensor that transmits real-time data and a terminal to display multimedia content to users [1]. XR functionality encompasses a wide range of tasks, including simultaneous localization and mapping, hand gesture tracking, pose estimation, machine learning (ML) inference (e.g., object detection), and multimedia processing [2].

In spite of the significant ensuing requirements, wearables are limited by several factors. Today's wearables do not have any cable connection and rely solely on embedded batteries and wireless connectivity. It follows that performing any computation on a wearable has to account for the available processing power, memory, storage, connectivity, and battery life, which are far more limited in wearables than in other end-user devices like laptops. An additional issue is represented by the heat generation resulting from intensive tasks: indeed, excessive heat from a wearable poses serious risks to user comfort and safety, as it may cause not only skin irritation or burns but also dizziness and disorientation [3].

As a result, performing all required operations on the wearable is not always feasible, and XR can greatly benefit from a distributed processing approach [1], where the processing is *split* between edge servers and local devices. On the negative side, as offloading XR applications requires *substantial throughput and ultra-low latency*, XR systems also face connectivity hurdles. The network bandwidth to support immersive experiences necessitates new network technologies and higher frequencies, which, in turn, pose challenges such as intermittent connectivity and blockage [4].

In summary, edge computing can mitigate battery consumption and heat generation in XR devices, but its reliability in high-frequency bands is often compromised by blockages. Performing computations locally on the wearables sidesteps connectivity issue; however, it must be balanced against the wearables' tendency to become hot with usage, causing discomfort, thus disrupting the XR experience itself. This makes *the extent* of the split between local computation and edge offloading one of the crucial decisions to make for XR.

In this study, we focus on a scenario like the one depicted in Fig. 1, where multiple radio technologies are available, and a user device may run the XR application (or part of it) locally — as long as overheating does not occur — or offload it to the edge. Task offloading can be performed using different radio technologies, which entail different speed/reliability trade-offs. In this context, it is critical to assess how to best exploit the wearable's limited capability for local processing, jointly accounting for the thermal management of wearables and potential connectivity issues due to blockage, which hinders edge offloading. These two challenges call for opposite actions: thermal management would benefit from reducing on-device processing, while blockage issues can be tackled by increasing it.

Our main contributions can thus be summarized as follows:
- We identify the main challenges associated with enabling XR applications on wearables, including the risk of overheating and the potential for blockage that prevents task offloading to the edge;

*Olga Chukhno and Antonella Molinaro are with Mediterranea University of Reggio Calabria and CNIT, Italy; Francesco Malandrino is with CNR-IEIIT, Italy; Alessandro Catania is with University of Pisa, Italy; Carla Fabiana Chiasserini is with Politecnico di Torino, Italy, CNIT, Italy and Chalmers University of Technology, Sweden.*

ISSN: 0163-6804
Digital Object Identifier: 10.1109/MCOM.001.2400591

The data flow begins at the wearable, which transmits data from its embedded sensors, and ends therein, with the wearable displaying the complete XR video stream to the user.

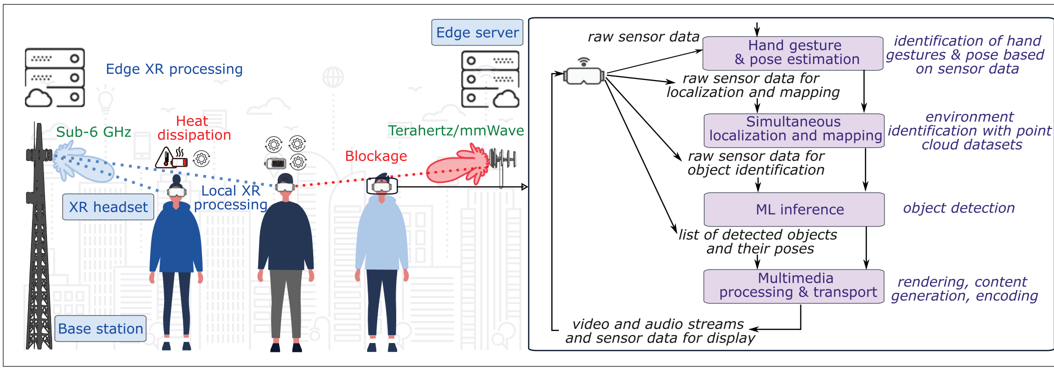


FIGURE 1. Our reference scenario, including the key entities — XR devices, base stations, and edge servers — (blue boxes), their roles in XR processing (violet text), main radio technologies (green text), and the faced issues (red text).

- We review the current approaches to blockage prediction and heat/power characterization and point out the need to explore their mutual interaction to optimize offloading strategies;
- We conduct the thermal performance analysis of a real-world wearable device, considering both local computation and offloading, using it as a case study. Our findings confirm that accounting for the trade-off between blockage prediction and thermal performance estimation is essential for optimizing system performance;
- Finally, we outline the road ahead, highlighting open issues that need to be addressed, including optimizing network infrastructure, advancing hardware materials, and predicting and managing mobility-related blockages for enhanced XR performance.

## Context and Motivation

In this section, we review the latest trends in XR technology and highlight how they are intertwined with the challenges of device overheating and link blockage.

**XR technology evolution.** XR devices have become widely available and are continuously evolving. This evolution can be traced back to the early days of virtual reality, from innovations such as Google Cardboard to recent developments from companies like Oculus and Apple, which are already offering highly immersive and accessible experiences. Beyond hardware and software advancements, wearables themselves are evolving rapidly, and key areas of development include increased display resolution, higher tracking accuracy, more realistic haptic feedback, eye tracking, improved rendering, and advancements in wireless connectivity [5]. These developments not only collectively contribute to a more immersive, realistic, and accessible XR experience, but also create significant computational demands, often exceeding the capabilities of modern wearables.

As an example, Fig. 1 illustrates the typical traffic flow of an XR application and the various XR functionalities it requires. The data flow begins at the wearable, which transmits data from its embedded sensors, and ends therein, with the wearable displaying the complete XR video stream to the user. As a consequence of the wearables' limited capabilities, substantial research is focused on how best to *offload* computing tasks from wearables to the edge.

In this context, Ericsson envisions several *levels* of offloading based upon the task to perform and the local conditions [2]. Options range from a low-offload architecture, where most processing occurs on the device, and only heavy tasks like compression and transmission of point cloud data are delegated to the edge, to high-offload setups, where only essential sensor data, such as the camera feed, are transmitted to the edge, with the bulk of processing handled locally on the end-user device.

Offloading requires ultra-low latency data transfer to and from the edge, which in turn necessitates high-speed wireless connectivity.High-frequency technologies can provide the necessary speed, however, they are vulnerable to blockage. While local computation can address this issue, it may lead to the same overheating problems that offloading was originally meant to mitigate.

**Challenge 1: blockage.** Offloading XR workloads to edge servers requires significant data transmission over the network. To provide the required capacity, 5G and beyond systems can operate at higher frequencies using directional antennas. However, harsh propagation conditions at high-frequency environments make them susceptible to channel intermittency and blockage caused by both fixed structures (e.g., buildings) and human movement.

The dynamic nature of human blockage in XR scenarios, including wearables with state changes occurring within milliseconds, further complicates the problem. The changing position of the body relative to the mobile device or head movement may be an additional cause of rapid drops in signal strength. Specifically, dynamic blockages can cause additional attenuation (of approximately 15 — 40 dB) [6], which negatively impacts signal quality, making the network unreliable for time-sensitive applications like XR. Moreover, such events trigger frequent handovers, leading to increased data transmission delays, lower throughput, and higher power consumption for user equipment, ultimately degrading the user experience.

**Challenge 2: heat.** The second challenge is unique to wearables and stems from their compact design, which is necessary for their physical use. The dense integration of components within such a limited space exacerbates heat dissipation issues, making traditional convective cooling solutions, such as fans, impractical. This limitation restricts heat dissipation capacity to a few watts. At the same time, wearables have tighter temperature limits than other mobile devices; hence, accounting for the *instantaneous* power consumption becomes necessary to avoid overheating and the ensuing damages like low-temperature burns [7].

Another power-related issue stems from the high computational demands associated with high-resolution displays, complex 3D rendering, and real-time tracking. Significant amounts of computation must occur directly on the wearable, meaning that both hardware and software techniques are needed to keep energy consumption as low as possible, thus maximizing battery duration.

Therefore, it is essential to develop optimal management strategies that address the intertwined challenges of blockage and overheating by maintaining temperature and battery levels within safe ranges and dynamically adjusting computational load.

## Existing Approaches

In this section, we examine the latest approaches to address the challenges of blockage and heating in XR applications and highlight the gap in current research that we seek to fill. Note that the task of *detecting* blockages is distinct from their mitigation and is typically performed using timeouts.

**Blockage mitigation.** To address connectivity challenges, particularly signal blockage, prior work has proposed using multiple network technologies concurrently. However, these solutions can be complex to implement at high frequencies, especially in power-limited devices. Beyond the fundamental requirements of power and timing tracking, both the network and the user device must constantly monitor the transmission direction of each potential link. Specifically, this necessitates real-time estimation and tracking of the angles of departure/arrival for each beam. This means that devices must dynamically reconfigure or maintain multiple beams simultaneously to avoid losing connectivity. Without proper motion and blockage prediction models, which have yet to appear, this process incurs substantial overhead, increasing computational complexity for beam selection, tracking, and management, as well as handover procedures.

To help dealing with this issue, computer vision and neural networks have been exploited to enhance reliability. For example, [8] leverages computer vision to detect obstacles, user location, and speed, and a neural network model to mitigate beam blockage and reduce handover frequency. While computer vision focuses on real-time detection, blockage prediction aims to anticipate potential obstructions, allowing for proactive actions. These predictors typically utilize ML models that consider communication link metrics and external factors such as visual and location data. However, one major challenge with ML predictors is their reliance on large deployment-specific training datasets, an issue that can be addressed using meta-learning techniques [9].

Furthermore, analytically tractable models have been developed to handle blockages. An extensive survey of such models is presented in [6], covering essential components for modeling scenarios that include deployment, propagation, antenna characteristics, blockages, micromobility, beam searching, traffic patterns, and service models under various system and environmental conditions. These models provide a foundation for the development of blockage prediction models that are essential in the complex environment of future-generation networks and applications like XR. However, effective prediction of link blockages in XR requires integrating user-side factors, such as application-dependent behavior [4], as usage patterns directly influence motion patterns (e.g., movement speed and direction), which in turn impact blockage probability.

**Thermal and battery management.** To tackle the challenge of heat dissipation, the existing research has explored various strategies to optimize thermal management and power consumption in XR devices. A recent thermal analysis model [10] breaks down XR glasses into individual components, creating an equivalent thermal resistance circuit for each part. This simplifies heat-generating elements, incorporating power consumption as the heat source, and provides heat transfer coefficients for natural convection in air.

The influence of material selection on temperature regulation in XR glasses has also been explored, revealing that the choice of frame materials significantly impacts temperature control. For example, while aluminum can decrease overall temperature, it may cause low-temperature burns near the ear due to its high thermal conductivity. Conversely, cellulose acetate plastic reduces ear temperature but increases the surface temperature of the device body. A combination of cellulose acetate plastic for the temples and temple tips together with aluminum for the rims offers a balanced solution for temperature regulation across the device. Recent studies on other innovative materials, such as advanced nanopolyhybrids [11], suggest significant performance improvements in XR glasses. Beyond thermal management, these nanopolyhybrids contribute to the miniaturization of electronic systems. These materials may offer efficient heat management, extending the usage time of wearables and enhancing user comfort.

Another area of focus in current research is optimizing the arrangement of electronic components to effectively mitigate thermal challenges. For example, [12] employs a genetic algorithm and a thermal resistance circuit to minimize temperature around the ears and frequently touched areas of XR glasses, enhancing user comfort and safety. However, the time that XR glasses can be worn in active operation mode before they become hot can be influenced by patterns of use [13]. Given their close proximity to the body and prolonged use, maintaining safe operating temperatures is crucial to prevent tissue damage. The thermal models developed so far allow for an estimation of the steady state temperature of the system, given the heat sources, the geometrical constraints, and the heat transfer coefficients. An enhanced model for transient simulations would be beneficial to predict maximum temperatures during intense computational periods, thus ensuring user safety and comfort throughout the experience.

In summary, predictive modeling holds notable significance in anticipating when XR devices are prone to thermal issues and begin to overheat. These models must consider various aspects, such as device characteristics, usage patterns, environmental factors, and thermal properties of materials, to facilitate preemptive thermal and battery management. This capability would greatly help XR developers and network operators by reliably predicting on-device processing limitations before thermal issues disrupt the user experience.

**Research gap.** Existing research addresses blockage mitigation through multi-connectivity, computer vision, and analytical modeling, and thermal man-

Significant amounts of computation must occur directly on the wearable, meaning that both hardware and software techniques are needed to keep energy consumption as low as possible, thus maximizing battery duration.

Making optimal offloading decisions — determining where to run computations — depends on multiple factors, including traditional ones like network load, channel conditions, and device capabilities, which are commonly considered in existing offloading strategies.

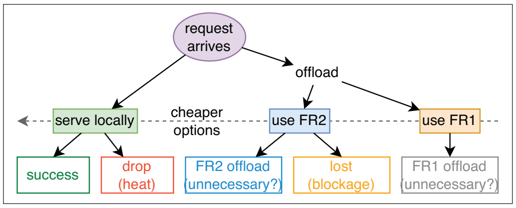


FIGURE 2. Decisions made by wearables about how to serve an XR request. First, they choose between local execution and edge offloading; if they opt for offloading, then they decide between FR1 and FR2. Locally-served requests may be dropped due to overheating; while FR2-offloaded requests may be lost due to blockage. In terms of cost, local execution is the most efficient, followed by FR2, and then FR1. If a cheaper option is viable, offloading is deemed "unnecessary."

agement through material selection, component optimization, and steady-state modeling. However, these solutions treat blockage and thermal challenges separately. We identify a need to integrate these factors within the specific context of XR.

Making optimal offloading decisions — determining where to run computations — depends on multiple factors, including traditional ones like network load, channel conditions, and device capabilities, which are commonly considered in existing offloading strategies. In this work, we seek to broaden the focus to include novel, XR- and wearable-specific factors, most notably, the interaction between network blockage prediction and thermal management. These two aspects are intertwined, as both are critical for ensuring the overall performance and quality of XR experience. Accurately predicting the duration of a blockage is essential for deciding whether local computation is viable. At the same time, precise models of wearable thermal behavior are needed to determine how much local computation can be performed before overheating occurs. Neglecting either factor could lead to suboptimal decisions — either excessive or insufficient local computation — resulting in high costs or overheating issues.

## Mutual Impact of Task Offloading, Link Blockage, and Wearable Temperature

In this section, we characterize the relationship among task offloading decisions, network conditions, and wearable temperature. To this end, we first present our reference system, and then the strategies we compare and their performance.

### Reference System

In the following, we first introduce the options available to wearables to serve XR requests, their possible outcomes, and the device thermal model we consider.

**Decisions, outcomes, and objectives.** We consider a system similar to the one in Fig. 1, where a wearable executes an XR application by combining local computation and task offloading. For simplicity and without loss of generality, we model the application load as a series of discrete *requests*. As summarized in Fig. 2, for each request, the wearable can:

- Serve the request locally (the cheapest option);
- Offload the request to the edge using Frequency Range 2 (FR2) technology (e.g., mmWave), which provides plentiful, cheap bandwidth but is prone to signal blockages;
- Offload the request to the edge using Frequency Range 1 (FR1) technology (e.g., 5G NR in the C-band), which is more robust to blockage but incurs higher cost.

Serving a request locally impacts three key aspects of the wearable behavior, on different time scales:

- It increases the instantaneous power consumption, which shall not exceed the wearable's thermal design power;
- It increases the wearable temperature, which shall not exceed 43°C to avoid user discomfort;
- It depletes the wearable battery, whose level shall not drop to zero.

It is worth mentioning that repeated overheating cycles also accelerate battery aging, leading to reduced capacity and increased internal resistance over time [14]. Although this long-term degradation process is beyond the scope of the present study, it represents a key consideration for future work on sustainable thermal management in XR wearables. Among the three mentioned aspects, power and battery constraints are common to most mobile devices; in contrast, temperature is especially relevant and challenging to manage in XR wearables. For this reason, our focus is on thermal constraints. Accordingly, as shown at the bottom of Fig. 2, locally-served requests may be dropped due to overheating. For offloading, requests offloaded via FR2 might be lost due to blockages, while FR1 is assumed to be always reliable.

1. If the wearable estimates that the request can be served locally without overheating, it opts for local execution;
2. Otherwise, if the wearable predicts an unblocked transmission path, it offloads the request using FR2;
3. Otherwise, it offloads the request using FR1.

The decisions (and their outcome) depend critically upon two main factors:

1. The accuracy of the blockage prediction,
2. The precision in estimating heat increase due to local computation.

**Thermal behavior.** We simulate the thermal profile of a wearable device matching a real-world scenario [10], with a Ball Grid-Array (BGA)-packaged chip similar to the OMAP4430 system-on-chip (SoC) by Texas Instruments, which was designed for image/video processing and used in the Google Glass Explorer. We perform a time-dependent thermal analysis using the "Heat Transfer" module within COMSOL Multiphysics, a finite element analyzer, solver, and simulator. In the simulation, time is divided into 60-second slots. During local computation, the chip dissipates 0.6 W, corresponding to the thermal design power of the OMAP4430. When computation is offloaded to the edge, power dissipation drops to 0.1 W.

Figure 3 highlights how power dissipation influences the chip's temperature. As expected, temperature rises during periods of on-device processing and falls during offloading. The purple horizontal line in the figure marks the 43°C safety threshold, identified in [7] as the limit beyond which there is a risk of low-temperature burns. Critically, when slots with on-device processing occur in close succession, the device temperature can exceed this threshold, as demonstrated in the figure. This highlights the importance of understanding the device's thermal behavior to prevent overheating while taking full advantage of its local computational capability.

Importantly, in practical scenarios, wearables may not have full knowledge of their internal state, therefore, our model assumes that temperature estimates are affected by jitter, simulating the uncertainty in the wearable's awareness capabilities.

### Benchmark Approaches and Results

Throughout our performance evaluation, we compare the following approaches:

- Our **prediction**-based approach, leveraging existing blockage-prediction strategies to inform decisions;
- A **history-based** approach, where a blockage duration is assumed as the average of the last 10 observed durations;
- Two benchmark strategies called **FR1 only** and **FR2 only**, using (respectively) only FR1 or only FR2 for offloading.

Importantly, we do not propose a new blockage prediction strategy, rather we exploit existing ones. We do, however, account for different levels of prediction accuracy by introducing a jitter over the blockage duration, ranging from 0% (ideal prediction) to 30% (rough, yet serviceable, prediction). Similarly, we consider different values of jitter in estimation of the heat increase due to locally served requests. Intuitively, a larger jitter means a higher probability of hitting the temperature threshold, hence, dropping a request.

Figure 4 presents the outcome of the service requests across the different approaches. Each bar in the figure represents one approach, and all bars are of equal height, corresponding to the total number of requests. The colored segments within each bar indicate the possible outcomes. Of particular interest is the combined size of the orange and red segments, corresponding to failed requests, respectively, lost and dropped requests. The prediction-based approach always achieves a very low failure rate, which increases gradually with the prediction jitter. Notably, even with a 30% jitter, it outperforms the state-of-the-art history-based benchmark. Looking at the last two bars, corresponding to the "only FR1" and "only FR2" approaches, we highlight their contrasting behaviors. The "only FR1" strategy, in particular, results in a very low failure rate, matching that yielded by perfect prediction. The "only FR2" approach instead disregards blockage and results in the highest number of lost requests across all strategies.

The hatched areas in all bars correspond to unnecessary offloads, e.g., requests that are offloaded via FR1 even though FR2 or local computation would have been viable options. We observe that larger jitter values always result in more unnecessary offloads. However, even with a 30% jitter, our prediction-based approach still outperforms the history-based method. The "only FR1" strategy, in particular, results in a very large number of unnecessary offloads, as it altogether disregards the option of cheaper, FR2-based offloading. Although the lower failure rate of this approach may justify its higher costs in specific scenarios, in most practical cases, our prediction-based approach offers the best trade-off between reducing costs by exploiting all local service and offloading opportunities and maximizing the probability of successful request completion.

We now increase the jitter in heat estimation from 5% to 30%, and present the resulting outcomes in Fig. 5. We can immediately observe larger red areas, corresponding to higher dropped requests; this highlights the critical role that accurate thermal behavior estimation plays in any XR˜offloading scheme. Despite this increase in jitter, the relative performance of the different strategies remains consistent with the results shown in Fig. 4, which confirms the robustness of our prediction-based strategy under varying conditions.

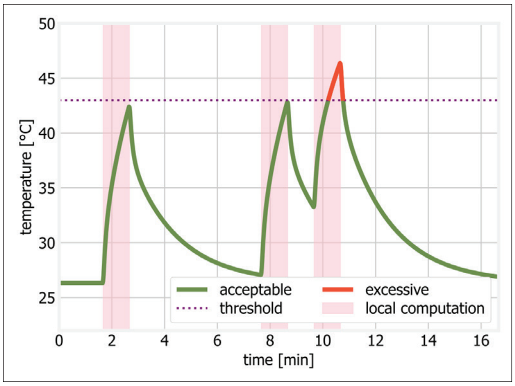


**FIGURE 3.** Thermal simulation of a device with an alternation of local computation (high dissipated power, pink areas in the figure) and task offloading to the edge (low dissipated power), with the duration of time slots fixed at 60 s.

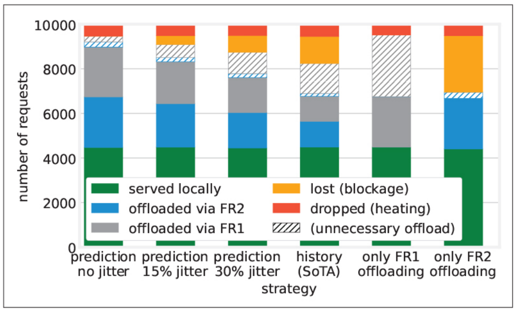


**FIGURE 4.** Request outcomes for all prediction strategies under a 5% jitter in the heat increase estimation.

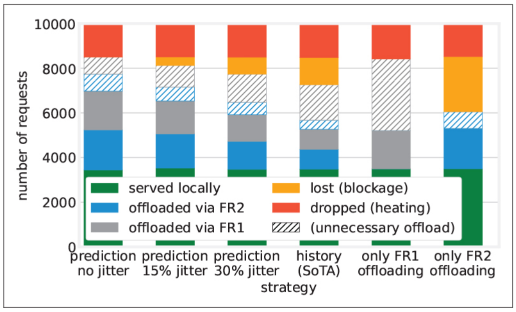


**FIGURE 5.** Request outcomes for all prediction strategies under a 5% jitter in the heat increase estimation.

However, ensuring seamless XR experience still presents challenges from both network and device perspectives, not only due to the demand for high throughput and ultra-low latency, but also because of issues related to link blockage and XR devices' thermal management.

### Open Issues and Envisioned Solutions

Further research is required to optimize network infrastructure and to fully utilize available resources from both networks and devices. Moreover, there is a need for research in transportation and neuroscience to effectively predict and handle blockages related to macro- and micro-mobility in XR environments. Advancements in electronics are essential to improve hardware materials and component distributions for better heat dissipation in XR devices. Another aspect is thermal management, which is essential for ensuring human safety and preventing thermal discomfort during

XR experiences [3]. The microclimate temperature between the user and the XR device, along with factors like battery life and weight, play a significant role in enhancing immersion quality, especially during prolonged use. Consequently, research must focus on thermal-aware solutions to mitigate the risk of overheating.

To enhance XR performance, several strategies could be considered, such as optimizing network efficiency through selective radio resource management, measurement skipping, and intra-device multi-modality. These approaches improve scheduling and support diverse QoS flows, tackling network congestion and device diversity. Latency and reliability issues can be addressed through delay-aware logical channel priority and radio link control in acknowledged modes. In addition, improving application information provision and conditional mobility can further enhance performance, by directly addressing the complexities of user movement prediction and obstacle management. Research on localized services, AI/ML integration, and improved XR positioning should also be considered to improve user perception and interaction [15]. This should be explored along with novel materials with superior thermal conductivity and innovative cooling mechanisms while testing new component distribution strategies to prevent overheating and prolonging devices' lifespan.

These solutions, among others, will contribute to the realization of future XR applications, ensuring their mainstream adoption across various industries, from gaming to education, ultimately making XR an integral part of human daily life.

## Conclusions

In recent years, wireless XR setups have revolutionized immersive experiences by allowing users to move freely and interact with their surroundings. However, ensuring seamless XR experience still presents challenges from both network and device perspectives, not only due to the demand for high throughput and ultra-low latency, but also because of issues related to link blockage and XR devices' thermal management. In this work, we addressed these challenges, identifying network outage as one of the most prominent issues. Furthermore, we demonstrated that effective blockage prediction and thermal management are key methods to successfully overcome network outages. Of paramount importance is the simultaneous consideration of these methods, which requires novel metrics and traffic management strategies that account for their mutual effect.

## Acknowledgments

This work was supported by SNS JU under the European Union's HE research and innovation programme under Grant Agreements No. 101095363 and No. 101192521, and by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU (PE00000001 — program "RESTART").

## References

[1] B. Sébire, WI Rapporteur RAN2 (Nokia) and Huilin Xu, WI Rapporteur RAN1 (Qualcomm), "eXtended Reality for NR," accessed July 21, 2025; https://www.3gpp.org/technologies, 2023.
[2] F. Alriksson *et al.*, "XR and 5G: Extended Reality at Scale with Time-Critical Communication," accessed July 21, 2025; https://www.ericsson.com/en/reports-and-papers, 2021.
[3] M. A. Rupp, "Is It Getting Hot in Here? the Effects of VR Headset Microclimate Temperature on Perceived Thermal Discomfort, VR Sickness, and Skin Temperature," *Applied Ergonomics*, vol. 114, 2024, p. 104,128.
[4] O. Chukhno *et al.*, "Interplay of User Behavior, Communication, and Computing in Immersive Reality 6G Applications," *IEEE Commun. Mag.*, vol. 60, no. 12, 2022, pp. 28–34.
[5] Güven Çelikkaya, "Top 10 XR Technology Trends in 2024," accessed July 21, 2025; https://www.linkedin.com/pulse, 2024.
[6] D. Moltchanov *et al.*, "A Tutorial on Mathematical Modeling of 5G/6G Millimeter Wave and Terahertz Cellular Systems," *IEEE Commun. Surveys & Tutorials*, vol. 24, no. 2, 2022, pp. 1072–1116.
[7] A. Q. Han, "Thermal Management and Safety Regulation of Smart Watches," *2016 15th IEEE Intersociety Conf. Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2016, pp. 939–44.
[8] M. Al-Quraan *et al.*, "Intelligent Beam Blockage Prediction for Seamless Connectivity in Vision-Aided Next-Generation Wireless Networks," *IEEE Trans. Network and Service Management*, 2022.
[9] A. E. Kalør, O. Simeone, and P. Popovski, "Prediction of mmWave/THz Link Blockages Through Meta-Learning and Recurrent Neural Networks," *IEEE Wireless Commun. Letters*, vol. 10, no. 12, 2021, pp. 2815–19.
[10] K. Matsuhashi, T. Kanamoto, and A. Kurokawa, "Thermal Model and Countermeasures for Future Smart Glasses," *Sensors*, vol. 20, no. 5, 2020, p. 1446.
[11] D. G. Atinafu *et al.*, "Nanopolyhybrids: Materials, Engineering Designs, and Advances in Thermal Management," *Small Methods*, vol. 7, no. 6, 2023, p. 2,201,515.
[12] K. Kusumi *et al.*, "Electronic Component Placement Optimization for Heat Measures of Smartglasses," *Ieice Electronics Express*, vol. 20, no. 6, 2023, p. 20,230,011.
[13] R. Bahru, A. A. Hamzah, and M. A. Mohamed, "Thermal Management of Wearable and Implantable Electronic Healthcare Devices: Perspective and Measurement Approach," *Int'l. J. Energy Research*, vol. 45, no. 2, 2021, pp. 1517–34.
[14] Y. Che *et al.*, "Battery Aging Behavior Evaluation Under Variable and Constant Temperatures with Real Loading Profiles," *2023 IEEE Applied Power Electronics Conference and Exposition (APEC)*, 2023, pp. 2979–83.
[15] M. Gapeyenko *et al.*, "Overview of NR Enhancements for Extended Reality (XR) in 3GPP 5G-Advanced," *IEEE Access*, 2024.

## Biographies

Olga Chukhno (olga.chukhno@unirc.it) obtained her double Ph.D. degree (2022) as an MSCA fellow from Tampere University, Finland and Mediterranea University of Reggio Calabria, Italy, where she is currently an Assistant Professor. Her main research interests include wireless communications and mathematical modeling.

Francesco Malandrino earned his Ph.D. degree from Politecnico di Torino in 2012 and is now a senior researcher at the National Research Council of Italy (CNR-IEIIT). His research interests include distributed ML as well as wireless, cellular, and vehicular networks.

Alessandro Catania obtained his Ph.D. degree from the University of Pisa, Italy, in 2020, where he is currently an Assistant Professor. His research interests include mixed-signal microelectronic design for high-temperature environments and implantable medical devices.

Antonella Molinaro received her Ph.D. degree from the University of Calabria in 1996. She is a Full Professor of telecommunications at the University Mediterranea of Reggio Calabria, Italy, with double affiliation at CentraleSupelec, University of Paris-Saclay, France. Her research mainly focuses on wireless and mobile networking, vehicular networks, and future Internet.

Carla Fabiana Chiasserini received her Ph.D. from Politecnico di Torino in 2000. She is currently a Full Professor with the Department of Electronic Engineering and Telecommunications at Politecnico di Torino, Italy. She is also an Associate Researcher at CNR-IEIIT and CNIT, and a WASP Guest Professor at Chalmers University of Technology, Sweden.