

# Scoping Sustainable Collaborative Mixed Reality

Yasra Chandio<sup>†</sup>, Noman Bashir<sup>‡</sup>, Tian Guo<sup>†</sup>, Elsa Olivetti<sup>‡</sup>, Fatima M. Anwar<sup>†</sup>

<sup>†</sup>University of Massachusetts Amherst, <sup>‡</sup>MIT, <sup>†</sup>Worcester Polytechnic Institute

**Abstract**—Mixed Reality (MR) is becoming ubiquitous as it finds its applications in education, healthcare, and other sectors beyond leisure. While MR end devices, such as headsets, have low energy intensity, the total number of devices and resource requirements of the entire MR ecosystem, which includes cloud and edge endpoints, can be significant. The resulting operational and embodied carbon footprint of MR has led to concerns about its environmental implications. Recent research has explored reducing the carbon footprint of MR devices by exploring hardware design space or network optimizations. However, many additional avenues for enhancing MR’s sustainability remain open, including energy savings in non-processor components and carbon-aware optimizations in collaborative MR ecosystems. In this paper, we aim to identify key challenges, existing solutions, and promising research directions for improving MR sustainability. We explore adjacent fields of embedded and mobile computing systems for insights and outline MR-specific problems requiring new solutions. We identify the challenges that must be tackled to enable researchers, developers, and users to avail themselves of these opportunities in collaborative MR systems.

## I. INTRODUCTION

Mixed Reality (MR) is an emerging technology increasingly used for leisure and safety-critical collaborative applications [1, 2]. From 2020 to 2021, 33 million AR/VR headsets were sold, a trend poised to accelerate with the release of Apple Vision Pro [3] and Generative AI fostering new applications. If used daily for 2 hours, the 33 million headsets could generate  $2.6 \times 10^5$  metric tons of greenhouse gas (GHG) emissions annually, based on a 50Wh daily usage and the average global carbon intensity of  $432\text{gCO}_2\text{eq/kWh}$  [4]. The carbon footprint of the broader MR ecosystem, including edge computing systems and cloud datacenters, will be much higher than low-power MR headsets.

Recent efforts have explored opportunities to improve the sustainability of MR [5, 6]. Zhang et al. [5] explore optimizing the energy efficiencies of networking components to enable sustainable development in Metaverse. Elgamal et al. [6] investigate MR hardware design space optimizations to reduce the lifecycle emissions of a single headset. Prior work on the energy efficiency of MR headsets has explored energy-efficient video processing, optimizing display power, and reducing the power used for tracking, among other applications [7]. There is also work on estimating and optimizing the energy consumption of gaming, which may apply to MR headsets [8]. While prior work takes essential initial steps, significant avenues for improving the sustainability of *collaborative MR* remain open, including carbon-aware spatiotemporal workload optimizations, reducing the carbon footprint of non-processing components, and leveraging prior work in the adjacent fields of embedded and mobile computing systems.

While prior work can inform sustainability efforts in collaborative MR, reducing MR’s energy and carbon footprint involves additional challenges due to the nature of computing tasks in MR. For instance, real-time visual processing for immersive environments is an especially demanding task [9]. Additionally, continuous user tracking through interaction modalities like hand and eye tracking presents unique challenges not encountered in other domains [10]. MR devices’ need for portability and wearability brings specific design and operational constraints, such as highly effective thermal management. It may also necessitate task offloading to edge and cloud systems for applications that could otherwise run efficiently on smartphones. These challenges are in addition to the usual issues faced by interactive, battery-powered, and mobile devices, such as balancing battery life with performance and ensuring reliable wireless connectivity.

This paper identifies the potential opportunities for energy- and carbon-aware optimizations in collaborative MR. In doing so, we make the following contributions.

- 1) Outline the ecosystem of MR applications and analyze the major energy/carbon footprint sources in MR pipelines.
- 2) Identify the opportunities for reducing the energy/carbon footprint and highlight the tradeoffs that must be navigated.
- 3) Outline research directions for researchers, application developers, and end users to enhance MR’s sustainability.

## II. LANDSCAPE OF COLLABORATIVE MR

This section overviews the collaborative mixed reality (MR) components and sources of energy consumption and carbon emissions. As illustrated in Figure 1, the MR ecosystem consists of MR devices connected to a local network, an edge, or a cloud in various configurations. Across these tiers, the MR sustainability implications include *embodied carbon emissions* in the hardware supply chain and *software operational emissions* due to electricity use. Next, we briefly describe its hardware and software components.

**HARDWARE COMPONENTS** include the headset and physical components in the network, at the edge, or in the cloud.

**1. Headsets** are a user’s primary interface to the collaborative MR ecosystem. They include sensors, a display, processors, networking components, and a battery. Similar to mobile phones, embodied emissions dominate the lifecycle carbon footprint of MR headsets [6], which end users and application developers cannot change. However, operational emissions are significant and will increase as advances in battery technology or wireless power transfer extend daily usage time.

**2. Network infrastructure** connects headsets to the edge or cloud servers, and its energy consumption depends on the

data transfer requirements and the distance between endpoints. While the network’s carbon emissions can be significant, prior work has not been done to quantify and reduce them.

**3. Edge computing** is pivotal to enabling real-time MR applications by providing high processing power of a dozen to hundreds of powerful servers close to the end user. While the sustainability implications of edge computing vary depending on the energy source, such as diesel generator-powered edge vs. solar-powered edge, it is primarily used to enable latency-critical applications. The use of edge infrastructure also reduces energy consumption in the network.

**4. Cloud** servers often handle the most resource-intensive MR tasks, and their data processing and storage emissions can be substantial. Despite their significant power needs, the cloud’s system-level efficiency is often higher than the edge but lower than the headsets that employ energy-efficient embedded processors, such as ARM-based. However, using the cloud may be inevitable for some applications as smaller processors cannot fit the bigger artificial intelligence (AI) models.

**SOFTWARE COMPONENTS** in collaborative MR are numerous; we outline essential tasks and the related work in Table I.

**1. Data collection and pre-processing** includes key tasks such as offline sensor calibration [11] and synchronization [12] and online data filtering [13] before it is used to capture user interactions and render MR experiences. The energy intensity of these tasks depends on the complexity of the environment, the types of sensors used, and application data requirements.

**2. System services** in MR systems provide fundamental services that maintain the operational efficiency of the device, such as display brightness adjustment. Managing power-aware system states [14] and optimizing idle states conserve energy [15], enhance performance, and extend battery life.

**3. User interaction** management includes gesture recognition, object manipulation and rendering, and display control.

Gesture recognition technologies such as eye-gaze tracking [10, 16], hand gesture recognition [9], and voice commands [17] require continuous tracking of users to enable interactions with the virtual environment aligned with human behaviors and expectations. The rendering of 3D objects [18] enhances user interaction. However, enabling immersive experiences requires many tasks, such as direct manipulation of virtual objects [19], simulating real-world physics to enable realistic interactions, proper alignment and anchoring of virtual objects in the real world, and occlusion handling to ensure virtual objects are consistent with the physical world [20].

Display management tasks such as resolution management [21], color calibration [22], adaptive brightness [23], and foveated rendering [24] ensure that the visual output is optimized for device capabilities and tailored to the user’s viewing comfort and environmental conditions. These settings are crucial for maintaining clarity, color fidelity, and overall visual comfort to prolong engagement in MR environments.

However, the energy aspects of these interactions are ignored to provide a safe, immersive, and secure MR experience.

**4. World understanding** tasks refine the system’s ability to interpret the surroundings. Depth estimation measures distances




Fig. 1. An overview of the collaborative mixed reality landscape.

and relationships between objects, which helps place virtual items accurately in a real space [25]. Object segmentation [26] and detection [27] identify and categorize different environmental elements, allowing the system to interact intelligently. These capabilities make MR more useful for practical applications, blending digital content seamlessly with the physical world [28] and enhancing user interaction by accurate head and pose tracking [29]. However, it is worth noting that most of the improvements for these services leverage compute-intensive deep learning approaches.

**5. Collaboration** is crucial in enabling immersive user experiences through effective content sharing across devices and with cloud/edge requiring services such as content delivery [30] and content caching [31]. It also requires managing computational loads efficiently by processing them at the edge [32] or cloud [33]. The shared remote experiences [34] require blending co-located users with remote participants [35]. In shared experiences, collaboration spans the entire ecosystem and is highly energy-intensive, with a further energy consumption challenge across multiple locations.

### III. SUSTAINABLE COLLABORATIVE MR

MR sustainability can be improved by reducing the energy and carbon footprint or exploiting potential tradeoffs between energy, carbon, and performance. We expand Table I to outline energy efficiency improvement (EEI), carbon efficiency improvement (CEI), energy performance tradeoff (EPT), carbon-energy tradeoff (CET), and carbon-performance tradeoff (CPT) opportunities in relevant tasks.

**IMPROVING EFFICIENCY** refers to the opportunities for reducing energy consumption and carbon footprint of MR without impacting performance, potentially at higher cost.

While the energy efficiency of computing has significantly improved, there are further optimizations possible in computing hardware’s energy efficiency, software’s algorithmic efficiency, and hardware-software co-design [6]. As outlined in Table I, improving energy efficiency may be difficult for the tasks requiring significant performance gains, as these tasks are likely to use more computationally intensive methods. Additionally, headset vendors must give application developers more control over the hardware to enable application-specific and context-aware optimizations.

The operational carbon footprint of MR depends on the carbon intensity of electricity used to power the headset, edge, and cloud. The carbon intensity depends on the mix of generation resources used to generate the electricity. If

fossil-fuel-based power plants generate electricity, the carbon intensity would be high and show less variability. When there is no variability in the carbon intensity, carbon efficiency, and energy efficiency are the same. However, if carbon intensity shows spatiotemporal variability, the carbon and energy efficiencies diverge [36]; it not only matters how much energy is consumed but also when and where it is consumed. Most tasks, except collaborative ones, have almost no flexibility and cannot explicitly optimize carbon-efficiency.

**CARBON, ENERGY, AND PERFORMANCE TRADEOFFS** must be navigated to improve MR sustainability if energy- and carbon-efficiency improvements have been exploited.

MR’s performance in creating immersive experiences for users takes the front seat, but not all applications require the highest level of resources, and significant sustainability gains can be made with a favorable performance sacrifice. Modern MR applications create seamless interactions and realistic simulations, requiring detailed and interactive virtual environments based on complex rendering tasks. However, the bigger is not always better, and immersiveness gains from high-resolution designs can be marginal for many applications. The application can afford better energy and performance tradeoffs (EPT) by prioritizing function over form.

Electricity’s carbon efficiency exhibits spatiotemporal variations, and the energy efficiency of different components in the MR ecosystem varies. These variations can be leveraged to gracefully navigate the carbon and energy tradeoffs (CET). For example, many offline tasks, such as compression and caching, can be offloaded to low-carbon edge/cloud locations where they run on low-carbon electricity. While the headset and the network consume energy when transferring the data, the total energy consumed can still be smaller than the on-device processing’s overall carbon footprint. However, sustainability implications beyond energy and carbon must be considered, as water consumption and computing resource requirements can be significant. These tradeoffs are especially possible in collaborative MR scenarios, *where users may already be geographically distributed and have carbon intensity variations*.

Finally, the carbon-performance tradeoffs (CPT) can be exploited in multiple phases of the MR lifecycle. Prior work on design space optimization demonstrates that significant hardware-software co-design opportunities can help reduce the hardware requirement in MR headsets [6]. The reduced hardware specification significantly reduces the headsets’ embodied carbon footprint while improving operational efficiency. Similar opportunities exist for optimizations across the ecosystem, such as headset vs edge. It can also be exploited by offline and background tasks with loose latency requirements, e.g., data compression or trend analytics. These tasks can wait for low-carbon periods to run, which will reduce carbon footprint but increase the completion time for these tasks.

We acknowledge that our list of potential research directions is not exhaustive, and other opportunities to improve MR’s sustainability exist. However, our work takes an important step towards making MR sustainable and integrating sustainability as an optimization metric in MR research.

TABLE I  
MAJOR TASKS AND RELATED WORK IN MR WITH POTENTIAL FOR ENERGY EFFICIENCY IMPROVEMENT (EEI), CARBON EFFICIENCY IMPROVEMENT (CEI), ENERGY PERFORMANCE TRADEOFF (EPT), CARBON-ENERGY TRADEOFF (CET), AND CARBON-PERFORMANCE TRADEOFF (CPT).

| Tasks  | EEI | CEI | EPT | CET | CPT |
|--|-----|-----|-----|-----|-----|
| <b>1 – Data collection and processing</b>    |     |     |     |     |     |
| 1.1 Sensor calibration [11]                  | ✓   | –   | ✓   | –   | –   |
| 1.2 Sensor synchronization [12]              | ✓   | –   | ✓   | –   | –   |
| 1.3 Data filtering [13]                      | ✓   | –   | ✓   | –   | –   |
| <b>2 – System services</b>                   |     |     |     |     |     |
| 2.1 Power-aware system states [14]           | ✓   | ✓   | ✓   | –   | –   |
| 2.2 Idle state optimization [15]             | ✓   | ✓   | ✓   | –   | –   |
| <b>3 – User interaction</b>                  |     |     |     |     |     |
| <b>Gesture recognition</b>                   |     |     |     |     |     |
| 3.1 Eye-gaze tracking [10]                   | ✓   | –   | –   | –   | –   |
| 3.2 Hand gesture recognition [9]             | –   | –   | –   | –   | –   |
| 3.3 Voice recognition [17]                   | –   | –   | –   | –   | –   |
| <b>Object rendering and manipulation</b>     |     |     |     |     |     |
| 3.4 Rendering 2D/3D models [18]              | –   | –   | –   | –   | –   |
| 3.6 Direct object manipulation [19]          | –   | –   | –   | –   | –   |
| 3.5 Anchoring, aligning, & persistence [20]  | –   | –   | –   | –   | –   |
| <b>Display</b>                               |     |     |     |     |     |
| 3.7 Resolution management [21]               | ✓   | –   | ✓   | –   | –   |
| 3.8 Color calibration [22]                   | ✓   | –   | ✓   | –   | –   |
| 3.9 Adaptive brightness [23]                 | ✓   | –   | ✓   | –   | –   |
| 3.10 Foveated rendering [24]                 | ✓   | –   | ✓   | –   | –   |
| <b>4 – Understanding the world</b>           |     |     |     |     |     |
| <b>Scene understanding</b>                   |     |     |     |     |     |
| 4.4 Depth estimation [25]                    | ✓   | –   | –   | –   | –   |
| 4.6 Semantic segmentation [26]               | –   | –   | –   | –   | –   |
| 4.5 Object detection [27]                    | ✓   | ✓   | –   | –   | –   |
| <b>Spatial mapping and 3D reconstruction</b> |     |     |     |     |     |
| 4.1 Handle occlusion, avoid collision [20]   | –   | –   | –   | –   | –   |
| 4.2 Real & virtual world blending [28]       | –   | –   | –   | –   | –   |
| 4.3 Pose/head tracking [29]                  | –   | –   | –   | –   | –   |
| <b>5 – Collaboration</b>                     |     |     |     |     |     |
| <b>Network/edge offloading</b>               |     |     |     |     |     |
| 5.1 Content delivery [30]                    | ✓   | ✓   | ✓   | –   | –   |
| 5.2 Content caching on edge [31]             | ✓   | ✓   | ✓   | –   | ✓   |
| 5.3 Compression [32]                         | ✓   | ✓   | ✓   | –   | ✓   |
| 5.4 Cloud-based processing [33]              | ✓   | ✓   | ✓   | ✓   | ✓   |
| <b>Multi-user experience</b>                 |     |     |     |     |     |
| 5.5 Remote experiences [34]                  | ✓   | –   | –   | ✓   | ✓   |
| 5.6 Blending co-located & remote users [35]  | ✓   | –   | –   | –   | –   |

#### IV. CONCLUSION AND FUTURE WORK

MR’s environmental sustainability implications are growing as it is deployed for applications beyond leisure. Prior work has explored specific aspects of MR sustainability, but significant opportunities remain, especially in the collaborative MR ecosystem. We map the collaborative MR landscape and discuss potential opportunities and their implications. We posit that efficient MR needs to integrate lessons from broader computing with targeted innovations tailored to the unique demands of MR systems. While there have been improvements, future work should optimize hardware to support novel software algorithms, enhancing sustainability and user experience. Each step forward contributes to more energy-efficient MR technologies and aligns with broader goals to reduce the carbon footprint of digital systems globally.

#### ACKNOWLEDGEMENTS

This research is supported by NSF Grants 2105564, 2236987, 2346133, 2237485, 2230143, and VMware.

## REFERENCES

- [1] T. M. Gregory, J. Gregory, J. Sledge *et al.*, “Surgery Guided by Mixed Reality: Presentation of a Proof of Concept,” *Acta Orthopaedica*, 2018.
- [2] T. Hamada, A. Hautasaari, M. Kitazaki *et al.*, “Solitary Jogging with A Virtual Runner using Smartglasses,” in *IEEE VR*, 2022.
- [3] M. Finnegan, “AR/VR Headset Sales Slide,” <https://www.computerworld.com/article/3706918/idc-arvr-headset-sales-slide-could-rebound-with-apple-meta-device-launches.html>, 2023.
- [4] D. Jones, “Global Electricity Review,” <https://ember-climate.org/insights/research/global-electricity-review-2022/>, 2022.
- [5] S. Zhang, W. Y. B. Lim, W. C. Ng *et al.*, “Towards Green Metaverse Networking: Technologies, Advancements and Future Directions,” *IEEE Network*, 2023.
- [6] M. Elgamal, D. Carmean, E. Ansari *et al.*, “Design Space Exploration and Optimization for Carbon-Efficient Extended Reality Systems,” 2023.
- [7] B. Duinkharjav, K. Chen, A. Tyagi *et al.*, “Color-Perception-Guided Display Power Reduction for Virtual Reality,” *ACM TOG*, 2022.
- [8] E. Mills, N. Bourassa, L. Rainer *et al.*, “Toward Greener Gaming: Estimating National Energy Use and Energy Efficiency Potential,” *Computer Games Journal*, 2019.
- [9] T. Luong, Y. F. Cheng, M. Möbus *et al.*, “Controllers or Bare Hands? A Controlled Evaluation of Input Techniques on Interaction Performance and Exertion in Virtual Reality,” *TVCG*, 2023.
- [10] T. Li, Q. Liu, and X. Zhou, “Ultra-Low Power Gaze Tracking for Virtual Reality,” in *ACM SenSys*, 2017.
- [11] A. Jadid, L. Rudolph, F. Pankratz, and G. Klinker, “Utilizing Multiple Calibrated IMUs for Enhanced Mixed Reality Tracking,” in *ISMAR-Adjunct*, 2019.
- [12] S. Ganerwal, R. Kumar, and M. Srivastava, “Timing-sync Protocol for Sensor Networks,” in *SenSys*, 2003.
- [13] Q.-L. Han, D. Ding, and X. Ge, “Secure Control and Filtering for Industrial Metaverse,” *Frontiers of Information Technology and Electronic Engineering*, 2024.
- [14] F. Liu, Q. Pei, S. Chen *et al.*, “When the metaverse meets carbon neutrality: ongoing efforts and directions,” 2023.
- [15] M. Lyu, R. D. Tripathi, and V. Sivaraman, “MetaVRadar: Measuring Metaverse Virtual Reality Network Activity,” *ACM SIGMETRICS*, 2023.
- [16] L. E. Sibert and R. J. Jacob, “Evaluation of Eye Gaze Interaction,” in *ACM CHI*, 2000.
- [17] R. M. Hanifa, K. Isa, and S. Mohamad, “A Review on Speaker Recognition: Technology and Challenges,” *Computers & Electrical Engineering*, 2021.
- [18] F. Massa, B. C. Russell, and M. Aubry, “Deep Exemplar 2D-3D Detection by Adapting from Real to Rendered Views,” in *IEEE CVPR*, 2016.
- [19] V. Schwind, P. Knierim, C. Tasci *et al.*, ““These Are Not My Hands!”: Effect of Gender on the Perception of Avatar Hands in Virtual Reality,” in *ACM CHI*, 2017.
- [20] M. Alfakhori, H. Dastageeri, S. Schneider, and V. Coors, “Occlusion Screening Using 3d City Models as a Reference Database for Mobile AR-Applications,” *ISPRS Annals*, 2022.
- [21] K. Debattista, K. Bugeja, S. Spina *et al.*, “Frame Rate vs Resolution: A Subjective Evaluation of Spatiotemporal Perceived Quality under Varying Computational Budgets,” in *Computer Graphics Forum*, 2018.
- [22] P. Dash and Y. C. Hu, “How Much Battery does Dark Mode Save? An Accurate OLED Display Power Profiler for Modern Smartphones,” in *MobiSys*, 2021.
- [23] A. Shye, B. Scholbrock, and G. Memik, “Into the Wild: Studying Real User Activity Patterns to Guide Power Optimizations for Mobile Architectures,” in *Micro*, 2009.
- [24] R. Xiao and H. Benko, “Augmenting the Field-of-View of Head-Mounted Displays with Sparse Peripheral Displays,” in *ACM CHI*, 2016.
- [25] C. Diaz, M. Walker, D. A. Szafrir, and D. Szafrir, “Designing for Depth Perceptions in Augmented Reality,” in *ISMAR*, 2017.
- [26] S. Hao, Y. Zhou, and Y. Guo, “A Brief Survey on Semantic Segmentation with Deep Learning,” *Neurocomputing*, 2020.
- [27] Y. Chen, C. Armstrong, R. Childers *et al.*, “Effects of Object Size and Task Goals on Reaching Kinematics in a Non-Immersive Virtual Environment,” *Human Movement Science*, 2022.
- [28] J. B. Barhorst, G. McLean, E. Shah, and R. Mack, “Blending the Real World and the Virtual World: Exploring the Role of Flow in Augmented Reality Experiences,” *Journal of Business Research*, 2021.
- [29] W. Wang, Y. Hu, and S. Scherer, “Tartanvo: A Generalizable Learning-Based VO,” in *CoRL*, 2021.
- [30] J. Zheng, Q. Zhu, and A. Jamalipour, “Content Delivery Performance Analysis of a Cache-Enabled UAV Base Station Assisted Cellular Network for Metaverse Users,” *JSAC*, 2023.
- [31] K. Zhang, S. Leng, Y. He *et al.*, “Cooperative Content Caching in 5g Networks with Mobile Edge Computing,” *IEEE Wireless Communications*, 2018.
- [32] Y. Leng, C.-C. Chen, Q. Sun *et al.*, “Energy-efficient Video Processing for Virtual Reality,” in *ISCA*, 2019.
- [33] H. Kim, J. Park, S. Yang *et al.*, “Edge-Cloud Cooperative Image Processing by Partially Streaming ROI Data for Metaverse Applications,” in *ICCE-Asia*, 2022.
- [34] N. I. Gonzalez-Romo, G. Mignucci-Jiménez, S. Hanalioglu *et al.*, “Virtual Neurosurgery Anatomy Laboratory: A Collaborative and Remote Education Experience in the Metaverse,” *Surgical Neurology International*, 2023.
- [35] A. Schäfer, G. Reis, and D. Stricker, “A survey on synchronous Augmented, Virtual, and Mixed Reality Remote Collaboration Systems,” *ACM Computing Surveys*, 2022.
- [36] W. A. Hanafy, R. Bostandoost, N. Bashir *et al.*, “The War of the Efficiencies: Understanding the Tension between Carbon and Energy Optimization,” in *HotCarbon*, 2023.