The Future VR/AR Network – Towards Virtual Human/Object Teleportation

NSF Workshop on Networked Virtual and Augmented Reality Communications†

Washington, DC, April 23-24, 2018

http://www.jakov.org/NSF_NetworkedVRAR_Workshop/

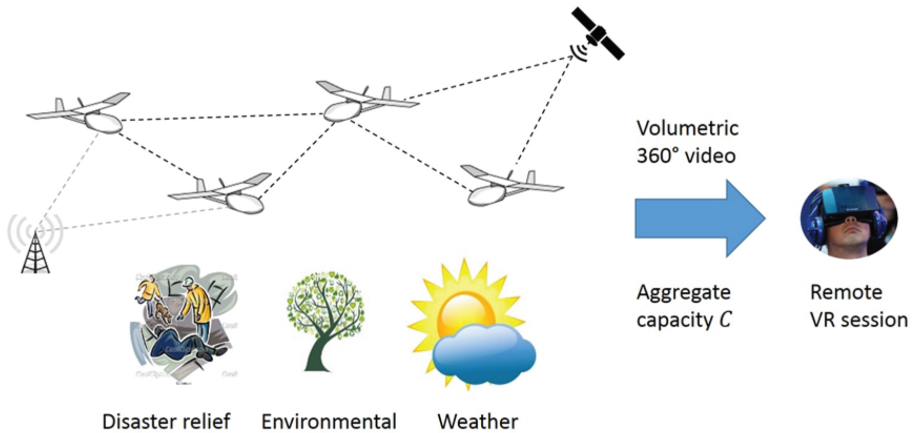

Principal Investigator:
Jacob Chakareski, University of Alabama
www.jakov.org, jacob@ua.edu

Workshop Report

This material is based on work supported by the National Science Foundation (NSF) under grant number 1821875. Any opinions, findings, and conclusions or recommendations expressed here are those of the principal investigator and workshop participants and do not necessarily reflect the views of the NSF.

†**Gold sponsors:** Adobe Research and the University of Alabama College of Engineering.

The Future VR/AR Network – Towards Virtual Human/Object Teleportation
Visioning Workshop on Networked Virtual and Augmented Reality Communications

Jacob Chakareski
University of Alabama

Executive Summary

This is a report on **The Future VR/AR Network – Towards Virtual Human/Object Teleportation**, the first Visioning Workshop on Networked Virtual and Augmented Reality Communications, held in Washington, DC on April 23-24. The event was sponsored by the National Science Foundation (NSF). Gold sponsors for the event were Adobe Research and the College of Engineering at the University of Alabama.

Virtual and augmented reality (VR/AR) hold tremendous potential to advance our society. Together with another pair of emerging technologies, 360° video and holographic video, they can suspend our disbelief of being at a remote location or having remote objects/people present in our immediate surrounding, akin to **virtual human/object teleportation**. Presently limited to offline operation and synthetic content and targeting gaming and entertainment, VR/AR are expected to reach their potential when deployed online and with real remote scene content, enabling novel applications in disaster relief, the environmental sciences, transportation, and quality of life. Networked VR/AR applications will play a major role in the envisioned/emerging global Internet of Things framework and are expected to represent the foundation of the anticipated 5G tactile Internet ecosystem. However, there are considerable challenges ahead, in the form of technology limitations and infrastructure costs.

The objective of the workshop was to gather a forum of experts and interdisciplinary practitioners to help identify the most promising horizons to explore to address the above challenges and contemplate a community agenda that will integrate emerging threads of related work into collaboration. This report describes the topics discussed at the workshop and summarizes the key points of the discussions. Material related to the workshop is available online at www.jakov.org/NSF_NetworkedVRAR_Workshop/.

There was strong agreement among the workshop participants that overcoming the present challenges and technology limitations is essential for enabling the next generation societal VR/AR applications. Similarly, there was uniform agreement that this will necessitate departing from traditional networking approaches that solely aim to increase data rates or lower transmission delays, as the performance gap between networked VR/AR requirements and present/upcoming networking technologies is only expected to increase. Instead, holistic solutions should be investigated that go beyond the traditional networking domain to **closely integrate capture, coding, networking, and user navigation** of VR/AR data.

There was strong consensus at the workshop that it is essential to invest now in research that aims to meet this objective. Moreover, steady investment should be made in developing publicly available benchmark datasets, source code, evaluation settings, and testbeds, to help accelerate such research and drive reproducibility and standardization, at the same time.

Background

The principal investigator (PI) made his first observations about the expected societal importance of VR/AR early on in 2017. Since then, he shared his vision broadly with colleagues in academia and industry, aiming to understand whether there was a community agreement about it. In particular, he traveled to three representative conferences in the field, IEEE INFOCOM in May 2017, ACM MobiSys in June 2017, and ACM SIGCOMM in August 2017, where he presented early work in this area [1, 2, 3, 4] and discussed his VR/AR perspective extensively. The feedback he received was uniformly very positive and conveyed a sensation of excitement about the promise of these technologies. There was also a sense of urging to bring their present challenges to the forefront of attention and start exploring networked systems aspects of VR/AR[1], to help deploy such applications online and enable their societal benefits.

These experiences and outcomes helped the PI solidify his belief of the forward-looking nature of networked VR/AR as a research agenda by the early fall of the same year. They also echoed the observations the PI made regarding a few preliminary international technical events on networked VR/AR that he organized in 2016/2017 [6, 7], which all turned out to be highly successful. Motivated by all these positive developments, the PI then approached the NSF with a white paper proposing the organization of a related visioning workshop and justifying the need for it. The white paper was followed with a longer research proposal, which received equally strong support from three NSF programs, NeTS, CSR, and CIF.

Given the broad NSF support, workshop participants were solicited from the three supporting communities, broadly categorized as Networks, Systems, and Theory. The PI introduced two additional categories into the workshop program under which participants were solicited, Industry and Interdisciplinary, to match well the nature and scope of the themes the workshop aimed to explore. The workshop gathered strong interest in response to its public call for participation that was periodically disseminated online over the four months leading to the workshop. A number of participants from each of the five participating categories above were previously invited by the PI directly. Overall, more than 140 researchers from diverse disciplines expressed interest in participating, out of which 40 were selected, given the projected size of the event. Seven NSF program directors from several programs also participated in the workshop program.

Strongly positive feedback has been received uniformly from all participants about their experience at the workshop and the opportunities for exciting research and collaboration it created and stimulated. The workshop also motivated follow-up technical events, e.g., special sessions and invited talks, on the topics it explored and its key findings, at related international conferences and technology forums.

Figure 1: A group photo of the event participants (courtesy of Glenn Ricart).

The PI acknowledges the support in preparing this report of many workshop participants who kindly shared their discussion notes and provided input as the report was being developed. The PI also acknowledges the support of the IEEE in publicizing the event online through the IEEE Future Directions program[2] and its IEEE Digital Reality initiative/portal [9] and social media channels.

[1]For instance, the INFOCOM keynote of Dr. Edward Knightly exclusively focused in its first half on VR/AR as next generation applications and the challenges they raise for our present and upcoming networks [5].

[2]Introduced to promote initiatives in emerging technologies identified as primary focus areas within the IEEE [8].

Summary of Key Findings

At the high level, the workshop established that overcoming the present networked VR/AR challenges and technology limitations that will help lead society into the envisioned VR/AR future necessitates departing from the traditional (**business as usual**) networking approach of addressing emerging applications by releasing a new generation of networking technology that offers higher data rates or lower transmission delays. That is because the performance gap between networked VR/AR requirements and present/upcoming networking technologies is only expected to increase [10, 5]. Instead, holistic solutions should be investigated that go beyond the traditional networking domain to **closely integrate capture, coding, networking, and user navigation** of VR/AR data. There was strong consensus at the workshop that it is essential to invest now in research that aims to meet this objective. Moreover, investment should be made in developing publicly available benchmark datasets, source code, evaluation settings, and testbeds, to help accelerate such research and drive reproducibility and standardization, at the same time.

Networked VR/AR applications will involve data capture across space and time. Integrating the dynamics of user navigation will be critical in deciding how the network should support this process effectively. Moreover, data representation methods that scale across space and time in sync with the dynamic network allocation will be necessary for efficient operation to enhance the end-to-end quality of experience. Simultaneously, they can help integrate system/network scalability and end-to-end reliability for heterogeneous VR/AR clients/applications.

Conventional capture involves a small number of collocated static sensors. This precludes **exploring fundamental limits that can be pursued by dynamic sensors** that adapt their locations and network rates over time, as the user navigates the remote scene. This will also have cost repercussions, expanding the scope of networked VR/AR applications considerably, since the alternative of traditional dense/static capture of remote scenes will be generally impractical or unfeasible. Integrating the user navigation and understanding how the network will need to dynamically adapt to support dynamic sensor location will be critical here. Are networked drones the needed cost-effective medium to support dynamic sensor capture and help usher in the next generation of societal VR/AR applications?

Understanding and characterizing the networked immersion quality of experience and its dependence on various system/network/client aspects is essential. This will require accounting for the user navigation patterns, in contrast to conventional quality measures that consider the fidelity of the reconstructed data only. Will interactivity (latency) matter more in this context than the sheer volume of data delivered to the user? It was established that edge computing will feature a very prominent part here to help address prospectively long end-to-end delays, and end-to-edge designs should emerge. Due to the stringent demand on real-time communication and low tolerance to delay jitter, providing deterministic, low-latency communication service will be critical, especially at the wireless edge. Moreover, to assure quality of experience (QoE) delivered to the user, more measurement and analysis are needed to characterize the QoE associated with different VR/AR applications. Similarly, it is essential to investigate effective and accurate mapping of QoE application requirements to respective networking/communication system quality of service (QoS) requirements, to ensure a holistic end-to-end operation. Publicly available benchmark datasets[3], source code, evaluation settings, and testbeds are sorely needed and will need to be developed to accelerate the research described herein, while helping drive reproducibility and standardization, at the same time.

Legacy wireless technologies are insufficient to meet the Gbps data rates needed for VR/AR applications. **Investigating spectrum bands with large unlicensed bandwidth** may offer a prospective solution. In particular, techniques utilizing high frequency radio (such as millimeter-wave) and/or free-space-optical bands should be explored. Generally, VR/AR-aware utilization of various wireless bands could be researched for better exploitation of their PHY features. Moreover, novel protocols and architectures are needed to ensure predictable communication quality in dynamic, uncertain wireless networks (e.g., due to node mobility and wireless channel fading). Similarly, it is important to understand what contextual information needs to be shared during a mobility event, and what context needs to be transferred from previous access point (AP) to the next one, to enable effective operation of future VR/AR applications.

[3]Comprising VR/AR raw/compressed data, QoE responses, and user navigation patterns.

Caching will have significant effect on VR/AR performance and novel multi-level caching architectures may be required, operating synchronously at the client device (personalized view cache), edge, and cloud. In particular, when caching is pushed to the edge, the traditional understanding of law-of-large-numbers caching may not apply any longer. Moreover, instead of traditional caching methods, personalized and viewport-driven strategies should be investigated, to capture the spatial and temporal locality induced by user navigation upon the VR/AR data. Similarly, it is important to understand the impact of the interaction of virtual and physical features in such applications on caching, as another new source of prospective data locality to be leveraged. Finally, it appears that prefetching may be the prevalent general approach to caching in this context, due to the extremely stringent latency constraints of VR/AR.

Traditional video is delivered via TCP/IP. Will the same approach be efficient for VR/AR? To explore the need for **protocol designs customized to VR/AR**, the traffic patterns that such applications will generate should be studied. It is expected that these patterns will change relative to traditional video, at least in the uplink, since considerable additional information will be sent from the client to the source, amplified by the fact that devices not only consume but also produce significant amounts of data in VR/AR. Moreover, novel protocols might be needed to better support "horizontal communication", since for some AR/VR applications spatial adjacency is important and end systems might want to exchange data in a peer-to-peer (P2P) fashion. This will likely require near-field and device-to-device (D2D) communications. In addition, TCP is designed for one-to-one communication, but AR/VR interactions will be many-to-many, as it expected that VR/AR environments will integrate multiple interacting end devices like cameras, sensors, and displays. This raises further challenges, as TCP is not only restrictive for end user/systems communication, but also for group (multi-device) D2D communication. Finally, with many devices involved in the creation of content/scenes, the challenges of synchronization between streams, rapid network orchestration, and dynamic resource provisioning become very important and protocols have to support them all well.

It was established that **edge computing** will feature a very prominent role to help address prospectively long end-to-end delays, and end-to-edge designs should emerge. Due to the stringent demand on real-time communication and low tolerance to delay jitter, providing deterministic, low-latency communication service will be critical, especially at the wireless edge. Understanding the trade-offs in placing the computation and resources/data at end-user, edge, and/or cloud sites will be essential. Preliminary progress in this direction has been established in [4]. It is generally anticipated that placement of computing tasks and data/resources will be driven by latency and related computing/storage constraints, as also verified by [4]. Limited energy on mobile devices must be integrated into the application design and edge computing strategies. One related challenge that arises is whether trade-off configurations of computing task and data/resource placement states should be explored and how should transition of such states be carried out during mobility? To enable collaborative mobile edge computing in this context, novel edge infrastructures and diverse network topologies should be explored. For software defined systems, exploring a control plane across the cloud, edges, and end-users for coordination and monitoring of tasks, data, bandwidth, and latency will be essential. Similarly, it is anticipated that designing edge infrastructure that will simultaneously support multiple applications, besides VR/AR, will be important for network providers.

Regarding novel services and applications, integration of additional modalities, e.g., haptic data, will lead to next generation VR/AR experiences. Major VR/AR societal applications that are envisioned include training/education, medical rehabilitation, first responders, remote monitoring for the weather and environmental sciences and large-scale infrastructures, and transportation. Extending the scope and deployment of such applications can benefit from integration with other telemetry sensors and emerging technologies such as unmanned aerial systems and machine learning.

1 Introduction and motivation

Virtual and augmented reality (VR/AR) hold tremendous potential to advance our society. Together with another pair of emerging technologies, 360° video and holographic video, they can suspend our disbelief of being at a remote location or having remote objects/people present in our immediate surrounding, akin to **virtual human/object teleportation** [11]. Presently limited to offline operation and synthetic content and targeting gaming and entertainment, VR/AR are expected to reach their potential when deployed online and with real remote scene content, enabling novel applications in disaster relief, the environmental sciences, transportation, and quality of life, as illustrated in Figure 2. Networked VR/AR applications will play a major role in the envisioned/emerging global Internet of Things framework and are expected to represent the foundation of the anticipated 5G tactile Internet ecosystem [12, 13, 14, 15, 16, 17, 18].

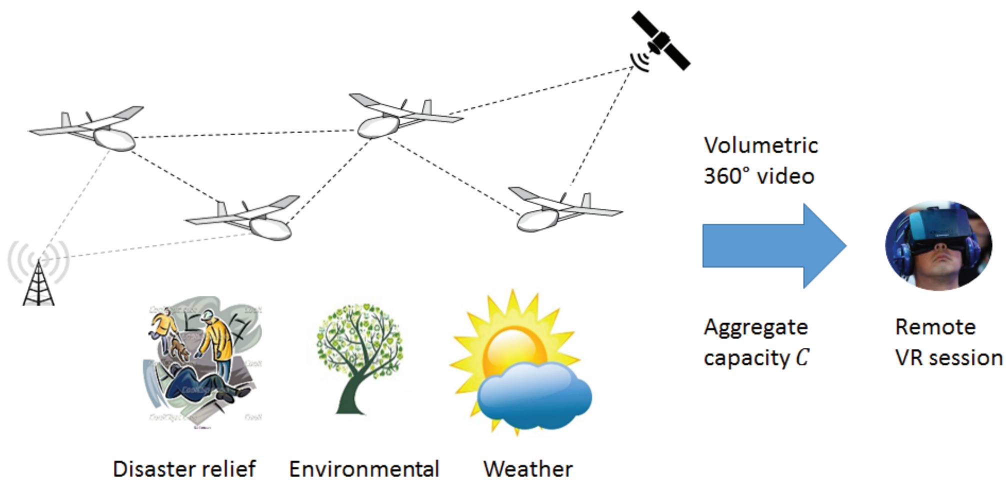

Figure 2: Remote 360° VR/AR immersion for diverse societal applications.

There are considerable challenges ahead in the form of technology limitations and infrastructure costs, however. The raw data rates (5-60 Gbps) required to enable an online immersion experience indistinguishable from real life dramatically exceed the FCC requirements for future broadband networks [10, 5]. Thus, simply introducing more bandwidth (**business as usual**) will not bridge this gap as the scales demand vs. supply are very different. This necessitates exploring holistic solutions that go beyond the traditional networking domain and **integrate capture, coding, networking, and user navigation** of VR/AR data. Moreover, emerging services, e.g., YouTube/Facebook 360° [19, 20] are extremely inefficient in bandwidth utilization and data management, thereby considerably degrading the user experience, due to their heuristic design choices. Finally, further critical aspects such as wireless operation, ultra-low latency, client/network heterogeneity, system scalability, edge computing/caching, and end-to-end reliability are yet to be considered.

2 Scope and objectives

The objective of the workshop was to gather a forum of experts and interdisciplinary practitioners to help identify the most promising horizons to explore to address the above challenges and build the foundations for a community agenda that will integrate emerging threads of related work into collaboration. By **setting a rigorous long-term vision platform of science and technology**, the workshop will help chart the field technically. In turn, this will help accelerate the introduction and broad adoption of related societal applications. The workshop advances will also facilitate **fundamental research in the general application area of high-volume high-speed/low-latency data transfer in emerging settings**, where for the first time the temporal and spatial dimensions of the data capture need to be closely explored and tightly integrated with the user navigation actions, to maintain the desired quality of experience for the end user, given the limited available system resources.

The workshop comprised a broad set of participants, as several disciplines and closer collaboration of researchers of diverse backgrounds needed to be involved to realize the above vision. The workshop program comprised a mix of breakout sessions, panel discussions, and short talks, scheduled over two days. The program included on the first day a panel discussion featuring industry leaders in the field and interdisciplinary practitioners, and an expert panel on the second day to identify the grand challenges in networked VR/AR over the next five years. The workshop participants were invited or selected from five different categories: **Networks, Theory, Systems, Industry, and Applications**, where the first three categories correspond to the respective NSF programs that supported the workshop, NeTS, CIF, and CSR.

The participants were integrated into the workshop program according to their category and the themes and topics of discussion it comprised. Each theme was associated with a respective breakout session, focusing on the individual topics identified to comprise that theme, and prospectively multiple presentations given during the rest of the program. The detailed program and participant list of the workshop are included in the appendices of this report.

3 Themes and topics of discussion

The following themes and research questions were explored at the workshop:

- What are the fundamental performance limits/bottlenecks of networked VR/AR? Will innovations be required at all layers of the protocol stack to help overcome them and what will be the required novel network architectures? Are network bandwidth and latency the only major performance factors and which one is more relevant? How can system and network scalability and end-to-end reliability be introduced for such demanding applications?

- Networked VR/AR applications will involve data capture across space and time. Integrating the dynamics of user navigation will be critical in deciding how the network should support this process effectively. Moreover, data representation methods that scale across space and time in sync with the dynamic network allocation will be necessary for efficient operation to enhance the end-to-end quality of experience. Simultaneously, they can help integrate system/network scalability and end-to-end reliability for heterogeneous VR/AR clients/applications.

- Conventional capture involves a small number of collocated static sensors. This precludes exploring fundamental limits that can be pursued by dynamic sensors that adapt their locations and network rates over time, as the user navigates the remote scene. This will also have cost repercussions, expanding the scope of networked VR/AR applications considerably, since the alternative of traditional dense/static capture of remote scenes will be generally impractical or unfeasible. Integrating user navigation and understanding how the network will need to dynamically adapt to support dynamic sensor location will be critical here. Are networked drones the needed cost-effective medium to support dynamic sensor capture and help usher in the next generation of societal VR/AR applications?

- Understanding and characterizing the networked immersion quality of experience and its dependence on various system/network/client aspects will be essential. This will require accounting for the user navigation patterns, in contrast to conventional quality measures that consider the fidelity of the reconstructed data only. Will interactivity (latency) matter more in this context than the sheer volume of data delivered to the user?

- Untethered, in the wild, wireless operation. Can edge computing help us overcome some of the fundamental challenges in terms of latency or bandwidth in this context? How can the available resources (rates/computing/storage/power) be traded for one another here and assigned across space, time, and the network hierarchy to maximize the system efficiency?

- What novel/creative service and business models will be required to accelerate adoption?
- What will be the requirements of diverse interdisciplinary networked VR/AR applications?
- Should user-centric design be integrated into the network hierarchy and where? Furthermore, early advances motivate holistic approaches to networked VR/AR systems that promote novel synergies between the network and the space-time nature of the data and the user navigation patterns [2, 3, 4]. What are the pros. vs. cons. of both of these strategies? Are they necessary to help introduce networked VR/AR applications that meet the present expectations?

4 Summaries and conclusions of breakout sessions

The discussions of these themes and questions were carried out first within smaller breakout sessions that focused in parallel on 1-2 individual themes and their related research questions and challenges. These activities took place during the first day of the workshop. Each breakout session comprised a subset of workshop participants that were interested in the themes and topics explored by that session. A session chair selected from the participants compiled the most important aspects and conclusions of the discussions carried out during that session. In a second round of discussions and interaction, on the second day, the entire set of workshop participants met together and the summaries and conclusions from each session were presented and read out individually, one by one. Time was allocated to follow the presentations of the individual breakout session summaries with further interactive discussion among the workshop participants to supplement the summaries and conclusions of each session. These follow-up discussions also integrated further input and knowledge gathered from the industry panel and brief talks that took place in the interim. Below, the final breakout session summaries and conclusions are included individually, for every session.

4.1 Session 1: Networked VR/AR: Fundamental performance limits and bottlenecks. Novel network architectures and protocols.

Session discussion team: Cedric Westphal, Murat Yuksel, Michael Zink, Hongwei Zhang, K.K. Ramakrishnan, Glenn Ricart, Ravi Prakash, Yong Liu, Hulya Seferoglu, and Lixia Zhang.

Session chair: Murat Yuksel.

The session team discussed a number of aspects and challenges, as summarized below:

Quality of Experience (QoE) and Quality of Service (QoS) Requirements of VR/AR

- VR/AR applications are delay and throughput sensitive, and different types of VR/AR applications (e.g., video on demand, live streaming, interactive VR/AR) have different requirements (with interactive AR/VR applications being most demanding);
- To help address potentially long end-to-end delay, edge computing support is important, and end-to-edge designs should emerge;
- Due to the stringent demand on real-time communication and low tolerance to delay jitter, providing deterministic, low-latency communication service will be critical, especially at the wireless edge;
- To assure user experience, we need more measurement and analysis to quantify QoE for different VR/AR applications, and we need to map QoE requirements to communication QoS requirements.

Mobility and Last Hop Wireless

- How to handle node mobility and associated issues such as network functions adaptation and virtualization is a key design dimension.

- Legacy wireless technologies are far from satisfying the Gbps data rates needed for VR/AR applications. Most VR setups require wired viewport devices for satisfying the data rate requirements.
- Spectrum bands with large unlicensed bandwidth may offer a reasonable solution. Techniques utilizing high frequency radio (such as millimeter-wave) and/or free-space-optical bands may be explored. In general, VR/AR-aware use of various wireless bands could be researched for better exploitation of the PHY features of these wireless bands.
- Novel protocols and architectures are needed for ensuring predictable communication quality in dynamic, uncertain wireless networks (e.g., node mobility and wireless channel fading).
- What contextual information needs to be shared during a mobility event? What context needs to be transferred from the previous access point (AP) to the next AP?

Caching and Edge Computing Effects

- Caching will likely have significant effect on VR/AR performance.
- Caching architecture should be multi-level at client device (personalized view cache), edge, and cloud.
- When the caching is pushed to the edge, the traditional understanding of law-of-large-numbers caching may not apply.
- Due to the extremely stringent latency requirements of VR/AR, prefetching seems to be a more prevalent approach to caching.
- Instead of traditional caching algorithms, more personalized and viewport-based caching designs will likely prevail.
- Caching designs will heavily depend on user navigation/head movement within the context of the VR/AR application.
  - *Roller Coaster Scenario*: Head looks in one direction – very pronounced **spatial locality** among tiles (spatial regions) of 360° videos. Caching spatially collocated tiles of accessed tiles may be beneficial. **Temporal locality** (tiles accessed in the past may likely be accessed again) will exist here as well and should be leveraged by caching popular VR/AR tile content.
  - *Wedding Scenario*: More random head movements – temporal locality may be the only major source of locality.
- What are the key new sources of locality for VR/AR?
  - Due to heavy dependence on human movement, spatial locality is likely to be pronounced
  - Field-of-view (FOV) locality
  - Horizon level, equatorial plane locality
  - How do the virtual and physical features interact with each other for the purpose of caching?

Security and Privacy

As a new way of human-cyberspace interaction, AR is finding its way into a variety of applications, raising from commercial advertisements to telemedicine. Strong security protection becomes as critical as ever, yet a number of unique features in perceive AR/VR systems also make security solution development more challenging.

First, the performance requirements of AR mandates edge computing. Today, end users trust cloud servers by verifying their certificates issued by a small number of Certificate Authorities. Edge computing resources are likely to be provisioned by multiple stakeholders, raising questions of whether it is feasible to follow todays practice or new solutions will be needed.

Second, security and privacy requirements may be application specific. This can be difficult to foresee; we will learn more as AR/VR applications get more widely deployed. At this time, we believe that all AR/VR systems should at least support the three basic security measures:

(1) Data authentication: being able to verify the data source

(2) Data confidentiality through encryption

(3) Data availability

Third, there are at least two ways that security solutions may impact AR/VR applications. AR/VR applications have stringent real time performance requirements. We must be able to quantify the performance impact of data authentication and encryption/decryption through developing pilot AR/VR applications. In addition, usability of security solutions has been a well-known challenge. AR/VR applications might offer both new urgency, and new potentials, to the solution development.

Last but not the least, AR/VR production raises new issues regarding the Digital Rights Management (DRM) consideration. For example, in an AR system, everyone can become a content producer. Different from individual user inputs into YouTube, AR apps may employ edge computing to process and merge content from multiple users (e.g. videos from different angles into a wide-angle view), making it non-trivial to pin down the producers of specific content.

Economics and Business Models

- Legacy cloud computing approaches will likely not scale to VR/AR applications. We will need to decentralize computing resources and put them close to the edge and end-user.
- A key difference from legacy network applications is the extensive need for computation. How to place these computing resources along with network resources is a critical design issue. New concepts like Edge Stations where heavy computing resources could be co-located with base stations need to be explored.
- A better understanding of economic and regulation approaches appropriate for VR/AR is needed.
  - Should it be mostly SILOs[4] or horizontally integrated business boundaries?
  - What are the right places to establish economic and business boundaries for emerging VR/AR hardware and software technologies?
  - If VR/AR requires a significant infrastructure investment to reduce network latency and bring computation closer to the users, how is this funded? Can VR/AR be deployed as an OTT play?
  - Can there be edge operators as a business model? In a way similar to Internet eXchange Points are operated, VR/AR edges may be operated in an exchange model.
- What are potential killer VR/AR applications beyond entertainment? Examples include educational, observational (looking what is behind), and safety/first responders applications.

Protocol Design Issues

During the breakout session the topic of protocol design for VR/AR came up. The basic question that was discussed was the need for a new custom protocol for AR/VR. Traditional, single-view, 2D video is streamed via TCP/IP. Will this approach also be efficient for VR/AR? To answer this general question, additional questions have to be answered first. First of all, it is not clear to what extent AR/VR apps generate

[4] Businesses that are vertically integrated over a variety of technologies that may be necessary to deliver an end product. In the VR/AR case, for instance, this would mean that a single company is handling all the hardware and software technologies needed to deliver a VR/AR application. The VR helmet, the edge computing software, the positioning systems would all be handled by one company. A non-SILO approach would mean that these (sub)technologies are developed and innovated separately (or in parallel) by different companies. One of the motivations for the layered architecture of the network protocol stack was to prevent SILOs.

traffic patterns that are vastly different from the ones we see in traditional streaming? One conjecture we made is that the pattern will definitely change in the uplink since additional information is transmitted from the client to the source. This is amplified by the fact that devices will not only consume but also produce significant amounts of data. In addition, protocols might be needed that better support "horizontal communication", since for some AR/VR applications spatial adjacency is important and end systems might want to exchange data in a P2P fashion. This will likely require near-field and device-to-device communications. In addition, TCP is designed for one-to-one communication but AR/VR interactions will be many-to-many. This presents further challenges, as it not only restricts communication between end user devices but also for D2D communication, since AR/VR environments might be composed of many end devices like cameras, sensors, and displays. Finally, with many devices involved in the creation of content/scenes, the challenge of synchronization between streams becomes very important and protocols have to support this synchronization well.

Other related issues that were discussed here include the need for and the challenge of very rapid network orchestration and dynamic resource provisioning that were found to be necessary for emerging and anticipated VR/AR settings. Similarly, it is expected that such applications will induce a mix of broadcast, multicast, and unicast network sessions, and addressing their needs effectively will be equally important.

4.2 Session 2: Dynamic spatiotemporal capture and data representation

Session discussion team: Gonzalo Arce, Zhu Li, and Yao Wang.

Session chair: Zhu Li.

The challenges of achieving 3 Degrees-of-Freedom (DoF) and 6 DoF immersive media experiences are multi-faceted. The core issues are **high fidelity** VR/AR content capture, **high efficiency** immersive content compression, and *fine granular scalable* representations that can offer rich set of adaptation tools to deal with network latency, loss and bottlenecks. The session team identified a number of specific challenges that are summarized below:

- 360° video provides 3 DoF immersion experiences: Can we have optimal local projections for efficient coding? Existing projection geometry options are arbitrary and non-adaptive.
- Foveated vision 360° region-of-interest (ROI) coding, (Deep) learning based solution for optimal pre-filter design.
- Spatio-temporal scalability, QoE metrics within a FOV, FOV stream switching support, and related caching and indexing challenges.
- New active sensors like 77GhHz radar support for outdoor point cloud/volumetric video capture. Deep learning algorithms for post processing.
- Point cloud/volumetric video is a graph signal, optimal sampling and compression shall be conducted under a graph signal processing framework, graph spectral band limitation would be the key constraint.
- Point cloud/volumetric video geometry coding is currently tree partition based, would there be a more flexible, efficient solution?
- Light field data is highly redundant, can there be new sparse signal based representation and compression?

The team further identified challenges specifically related to point clouds used in AR applications and 360° video used in VR applications. In particular, 3D sensors (Structured light, Time-of-Flight, multi-cameras) provide rich point cloud information of scenes. Simultaneously, RGB, hyper-spectral, and temperature or motion sensors provide content signals that can be thought of as living on such point clouds. The research challenges that arise here are related to 3D acquisition, representation, sampling, and compression of multi-modal point clouds, as summarized below.

Multi-modal point cloud capture

- How can multi-modal information from a multitude of sensors/cameras be captured on point clouds efficiently in real time?
- How can multi-modal point clouds be efficiently displayed on emerging display technologies?
- How can multi-modal point clouds be efficiently represented in a multi-resolution framework?

Point cloud representation and compression

- Graph representations of raw dense point clouds
- Optimal graph sampling for accurate approximation of such signals on the graph
- Graph Laplacian formulation that maximizes the sparsity of the induced graph transform coefficients
- Multi-resolution representation and bitstream scalability
- Space-time sampling of point cloud data
  - Motion prediction in graphs in 4D volumes
  - Sparse signal processing on graphs

360° video representation and compression

- How to optimally represent a 360° sphere's surface
- Optimal 360° video signal sub-sampling based on the graph spectrum of band-limited common signals or sparsity constraints

4.3 Session 3: Edge computing, user-centric design, and related trade-offs

Session discussion team: Torsten Braun, Sebastian Friston, Klara Nahrstedt, Zhisheng Yan, Tao Han, Bo Han, Maxine Brown, Sujit Dey, Philip Chou, and Jacob Chakareski.
Session chair: Klara Nahrstedt.

The session team discussed a number of challenges. They are summarized in the following for every individual topic comprising this theme that was considered in the discussion.

Placement and tradeoffs

- Tradeoffs of computation at end-user, edge and/or cloud sites
  - Consider **prioritization of application requirements** and QoS, and then partition the tasks to resources across user-edge-cloud
  - Agreement where **tasks and data are placed on user/edge/cloud**
  - Challenges: Heterogeneity, application requirements; classification of edges
- (a) Classification of edge: Local (cell phone, glasses), network, or vehicular edge.
- (b) Classification of edge with regard to ownership: Cloud-edge or user-edge

Task and data placement

- Network edge vs. cloud

  - Depending on latency, place some tasks close to users (on users)
  - Consider not only task placement but also data placement
  - Consider task workflow partitioning as well as data partitioning
- Challenges: **In-cast problem**, **high fidelity of data**; **synchronization** if one partitions tasks and data

Energy and tradeoffs

- Energy for edge computing
  - If smart phones are edges, energy is an important resource
  - **Application design should consider energy**
  - Energy on VR /AR devices must be considered or on wearable cameras and for energy reasons offload tasks/data
- Challenge: should we have **tradeoff configurations**?
- **Transition of states for tasks/data during mobility**
  - Transition on edges: Consider handoff and tradeoffs

    If one does, e.g., object detection, should one do soft or hard handoff?
  - To support mobile edge computing: **Need edge infrastructure and diverse topologies** to enable collaborative edge computing
    * Consider cost of **handoff and state placement**
    * Consider which tasks/data to use to place on edge devices: Compression? Point cloud processing? Rendering on glasses or edges? What enhancement functions to run on edge?

Latency

- Placement of tasks/data will be driven by latency
- Use hardware for local processing (GPU, acceleration hardware)
- Consider **access to lower level APIs** (not just socket interface)
- Consider **adaptation for throughput and latency**

  The MPEG DASH streaming standard [] has extensive bit-rate adaptation capabilities
- How far can network providers provide real-time/low latency?
  - For entertainment, it is OK to rely on network providers
  - For tele-surgery training one cannot rely on network providers

Computing paradigms

- Docker / Micro-service Technology
  - Utilize Docker, docker swarm, micro-services, Kubernetes – lightweight cloudlet/edge platforms for real-time resource/service processing
  - Develop application services in evolutionary way to run in dockers
- Software-defined systems
  - Consider control plane across cloud/edges/end-users to coordinate, monitor tasks/data/bandwidth/latency,
  - Decide on effective placement of tasks/data

- Orchestrate end-to-end VR/AR: How would an orchestrator for AR/VR look like?

- Server-less computing paradigm
- CDN/P2P models

Networked edge infrastructure

- Group of edge devices

  Consider **handover: Avoid handover between edge devices**

- Consideration of **storage**
  - Cache vs static storage
  - Cache placement in multiple places

    Tradeoffs between cache hit and bandwidth usage to enable cache consistency vs latency

- Design of edge infrastructure for multiple applications (important for network providers, e.g., AT&T)
  - Edges for TV, Phone, Data Services and VR/AR services
  - Edges utilized across multiple applications
    * Vehicles, 360° video, health-care,
    * Multi-services

4.4 Session 4: Novel services, business models, and interdisciplinary application requirements

Session discussion team: Lei Liu, Juan Wachs, Philip Chou, Kimberly Genareau, and Meghan Houghton.
Session chair: Kimberly Genareau.

The session team discussed the following challenges and topics:

- Need for haptic data in addition to audio and video to provide accurate VR experience.
- VR can provide live broadcasts of entertainment events, sports events, interviews for a mass audience.
- VR can be useful in several types of training: medical procedures, sports, rehabilitation
- For rehabilitation, VR can provide safer alternatives for learning skills following disability or change in ability. Specialists can also be provided in VR.
- Real-time data is necessary for informing patients about layout of city and timing of traffic lights, perhaps through a wearable device connected to the IoT.
- Teleportation can provide human contact to people who cannot travel or reach rehabilitation services.
- Hardware for enabling teleportation must be easy to access and at a reasonable cost for people who can benefit from using it the most. Package or kit sold in stores?
- Drones for analysis of volcanic eruptions: AR control of drone suite equipped with multiple analytical instruments for examining grain size distribution, charge structure, hydrometeor concentrations as they change in time and space.

4.5 Session 5: Networked Immersion quality of experience: Trade-offs between latency and data volume

Session discussion team: Magda El Zarki, Vishy Swaminathan, Nikhil Balram, Feng Qian, Gwendal Simon, Tamay Aykut, Balakrishnan Prabhakaran, and Brendan John.

Session chair: Brendan John.

The session team had extensive discussions across multiple topics. The following were identified as **biggest present challenges:**

- There is currently no direct mapping between quality of streaming to quality of experience (QoE) for immersive VR content as in traditional 2D video.
- Results from existing 2D metrics applied to 360 video are inconsistent since subjective experiences vary greatly within VR HMDs, and new metrics are needed for 360 video that quantify how we perceive content. These metrics depend heavily on the use case, and different applications will have different responses for good enough quality.
- Openly available datasets of QoE responses and head movements are lacking, and future researchers need to make data and evaluation code open source to drive reproducibility and standardization.
- Since this is an unexplored field we should tackle the problem of 3DOF (rotation only) 360 video first in specialized contexts to establish QoE evaluation methods that work for particular use cases, before generalizing methods and expansion to 6DOF, volumetric, and interaction metrics.

The following are **detailed notes** of the discussions that took place:

- At the heart of the problem is turning subjective measurements into quantitative ones.
- We have concluded that selecting an appropriate metric depends on the audience and content.
- Field by field this varies, for example in networking the 1st priority is preventing stalls/freezes, followed by resolution quality, followed by the rate at which quality degrades.
- For visual content small changes can be barely noticeable, and a lower quality video that doesn't vary much can be perceived as higher quality than a high resolution video that suddenly drops to low for a small number of frames and returns.
- In streaming bandwidth can go up and down based on many factors, creating difficulty in identifying what role latency plays into the degradation of experience.
- In many scenarios we have to treat the network as a black box and only consider latency of the end to end, or end to edge system.
- To wrap this up we see the need to identify specialized metrics for specific use cases, and establish these first instead of a general standard.
- We also note that individual behavior differs with response to content and task which must be controlled for comparisons.
- Large datasets are needed with consistent and uniform content that can reproduce consistent patterns for testing systems utilizing head prediction methods.
- A benchmark based on previously collected data and QoE metrics that don't required subject responses can then be used to compare different methods on equal footing.
- We found many analogs to gaming and network experiences where 3D based metrics like level of detail (LOD) can be applied to individual meshes, but not the overall 3D experience.

- We don't currently have a mapping or understanding of how quality of service translates to QoE.
- Our discussion revolved around 360 video, and mostly monoscopic. Raw Stereo 360 certainly increases bandwidth demands, but at improvement for QoE. Omnistereoscopic vision allows the perception of depth and significantly improves the level of immersion.
- We didn't deep dive into the volume of data, including 6DOF (rotation + translation) content, acknowledging a sound understanding of 360 videos would serve as a starting point.
- Interaction within 3D space remains a harder problem as well. There are few local multiplayer VR games, and even less networked experiences, unlike gaming where many systems for first person shooters or massive online games like World of Warcraft are available.
- At the end of the day we are establishing a need for consistent metrics that can reliably measure specific flavors of QoE for networked experiences that test how we perceive latency, and in video streaming how much is enough quality.
- We also have to encourage open source evaluation code and datasets for reproducibility and consistency across the field.

5 Summary of panel discussions

Two panels took place during the workshop. On the first day, there was an industry panel that discussed VR/AR applications and systems from an industry perspective. On the second day, a grand challenges panel was assembled to identify and discuss what the major research challenges and opportunities may be over the next 5 years. These are the summaries of the related discussions that took place during each panel.

(1) Industry panel: Applications and Systems

  - Cost is a major concern for companies.
  - Industry players want to thread carefully and slowly.
  - Companies prefer to wait for the killer app to hit the consumer market.
  - The form factor of the VR/AR device is important and can constrain/challenge the rest (adoption of the technology).
  - Standardization of volumetric video is undergoing and is important for broader adoption.

(1) Grand Challenges panel

  - Raise the engagement in VR, to keep the user active and involved, to broaden the application scope beyond entertainment.
  - Give the user a chance to dynamically engage with the content.
  - Immersion quality is important to understand and characterize. Very application specific. How different immersion modalities (beyond VR/AR) contribute to the entire experience? Multi-sense modalities (touch, pressure, even smell?).
  - How to relate network metrics to quality of immersion?
  - Native 6DOF should be explored. Certain applications (virtual traveling museum visits, collaborative work) may deliver higher quality of experience in this context. How to overcome the challenges of enabling stereo 360 video?
  - Understand different levels of semantics: Detection beyond objects. Synergies with machine learning in knowledge discovery in this context.
  - Security and privacy important for adoption, in particular, for AR.

6 Concluding remarks

The following aspects were identified as important in a concluding discussion at the end of the workshop.

- Shall we explore novel methods for dynamic data representation that dispense with prior advances in video and image compression? How to integrate them natively with the user navigation actions?
- Dynamic capture that captures the minimum amount of data required to reconstruct the current viewpoint of interest? How to address the uncertainty that arises here, when the user takes an action that is not anticipated?
- How to optimally integrate emerging technologies (mmWave, FSO) with existing state-of-the-art for synergistic performance? How could this be carried out together with the data capture and coding, for holistic analysis?
- The need to allow the user to select its viewpoint or interact with the VR environment is important in multiple context, save for maybe passive content consumption live movies.
- Synchronization of other sensing modalities with VR is important to maintain quality of immersion experience. Trade-offs between inter-domain and intra-domain delays.
- 360 video is harder to deal with than VR due to the bigger uncertainty induced by the former relative to VR (you know the entire data/environment ahead of time). Live 360 video is even more challenging due to the lack of future knowledge of the data and no prior history of user navigation actions.
- It may benefit us to address 360 video first and then move on to 6DOF video. However, for 6DOF we also need to predict the user spatial position in the virtual environment.
- Another challenge that needs to be explored and understood is how to keep the user engaged in VR environments. Maybe there are analogies with gaming that can be leveraged here.
- Certain settings can benefit from 360/VR relative to conventional video (the need for seeing everywhere can bring benefits): collaborative work (teams of people working together), scientific studies, first responders, firefighters, transportation, remote sensing, forensics, assisted living, patient rehabilitation and healthcare, real estate, large infrastructure monitoring, training and education, stereoscopic and asymmetric rendering, for enhanced depth perception
- User specific vs. application specific view prediction. In AR applications, the context information is also important and can be targeted to a specific user dynamically.
- AR applications may require varying latency, depending on the application and context. This can be leveraged to do better utilization of resources.
- The Internet is best effort, however, some of these applications have strict performance criteria. Though the former may not be critical for mass consumer applications, it has been argued that a broad adoption of one may lead to developments that may enhance more real work related applications that have tight constraints.

References

[1] J. Chakareski, "Aerial UAV-IoT sensing for ubiquitous immersive communication and virtual human teleportation," in *Proc. INFOCOM Workshop on Communication and Networking Techniques for Contemporary Video*. Atlanta, GA, USA: IEEE, May 2017.

[2] X. Corbillon, A. Devlic, G. Simon, and J. Chakareski, "Viewport-adaptive navigable 360-degree video delivery," in *Proc. Int'l Conf. Communications*. Paris, France: IEEE, May 2017, **(best paper award)**.

[3] J. Chakareski, "Drone networks for virtual human teleportation," in *Proc. MobiSys Workshop on Micro Aerial Vehicle Networks, Systems, and Applications (DroNet)*. Niaga Falls, NY, USA: ACM, Jun. 2017.

[4] ——, "VR/AR immersive communication: Caching, edge computing, and transmission trade-offs," in *Proc. SIGCOMM Workshop on Virtual Reality and Augmented Reality Network (VR/AR Network)*. Los Angeles, CA, USA: ACM, Sep. 2017.

[5] E. Knightly, "Scaling Wi-Fi for next generation transformative applications," Keynote Presentation, IEEE INFOCOM, Atlanta, GA, USA, May 2017.

[6] J. Chakareski (TP Chair), "IEEE Int'l Packet Video AR Workshop @ IEEE Int'l Conf. Multimedia & Expo," Seattle, WA, USA, Jul. 2016. [Online]. Available: http://www.icme2016.org/

[7] J. Chakareski (Organizer/Chair), "Special Session on Immersive Experiences in Virtual and Augmented Reality @ ACM Int'l Conf. Multimedia Systems," Taipei, Taiwan, Jun. 2017. [Online]. Available: http://mmsys17.iis.sinica.edu.tw/

[8] "New Technology Connections: IEEE Future Directions." [Online]. Available: https://www.ieee.org/about/technologies.html

[9] "The IEEE Digital Reality Initiative: Go-to platform at the forefront blend of virtual reality (VR), augmented reality (AR), mixed reality (MR), 360-degree media, immersive, and other future potential reality technology spaces." [Online]. Available: https://digitalreality.ieee.org/

[10] B. Begole, "Why the Internet pipes will burst when virtual reality takes off," Forbes Magazine, Feb. 2016.

[11] J. G. Apostolopoulos, P. A. Chou, B. Culbertson, T. Kalker, M. D. Trott, and S. Wee, "The road to immersive communication," *Proceedings of the IEEE*, vol. 100, no. 4, pp. 974–990, Apr. 2012.

[12] P. F. Drucker, "Internet of Things position paper on standardization for IoT technologies," Jan. 2015. [Online]. Available: http://www.internet-of-things-research.eu/pdf/IERC_Position_Paper_IoT_Standardization_Final.pdf

[13] D. Evans, "The Internet of Things: How the next evolution of the Internet is changing everything," Apr. 2011. [Online]. Available: https://www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf

[14] M. Kranz, P. Holleis, and A. Schmidt, "Embedded interaction: Interacting with the Internet of things," *IEEE Internet Computing*, vol. 14, no. 2, pp. 46–53, 2010.

[15] P. Vlacheas, R. Giaffreda, V. Stavroulaki, D. Kelaidonis, V. Foteinos, G. Poulios, P. Demestichas, A. Somov, A. R. Biswas, and K. Moessner, "Enabling smart cities through a cognitive management framework for the Internet of things," *IEEE Communications Magazine*, vol. 51, no. 6, pp. 102–111, 2013.

[16] W. Barfield, *Fundamentals of wearable computers and augmented reality*. CRC Press, 2015.

[17] G. P. Fettweis, "The tactile Internet: Applications and challenges," *IEEE Vehicular Technology Magazine*, vol. 9, no. 1, pp. 64–70, Mar. 2014.

[18] Nokia Siemens Networks, "2020: Beyond 4G Radio Evolution for the Gigabit Experience," White Paper, Aug. 2011.

[19] "Facebook 360: A stunning and captivating way to share immersive stories, places and experiences." [Online]. Available: http://facebook360.fb.com

[20] "YouTube: 360° Videos." [Online]. Available: https://www.youtube.com/

Appendix A: Participants and Presentations

Jacob Chakareski, The University of Alabama
Virtual and Augmented Reality: The New Frontier (oral)

Tao Han, University of North Carolina Charlotte
Orchestrating Edge Network Resources for Mobile Augmented Reality (poster)

Murat Yuksel, University of Central Florida
Multi-Element High-Frequency RF and FSO Modules for Mobile VR/AR Networking (oral)

Klara Nahrstedt, University of Illinois Urbana
Teleimmersion and Video 360 Challenges (oral)

Hulya Seferoglu, University of Illinois Chicago
Coded Cooperative Computation for Augmented and Virtual Reality (oral)

Vishy Swaminathan, Adobe Research
Making Immersive Video Experiences Viable (oral)

Nikhil Balram, Google Research
A Perspective on Virtual and Augmented Reality (oral)

Cedric Westphal, Huawei Research and University of California Santa Cruz
Joint Field of View and Rate Adaptation for 360° VR Streaming (oral)

Amitabh Varshney, University of Maryland
Immersive Telepresence (oral)

Lixia Zhang, University of California Los Angeles
Supporting Augmented Reality: Looking Beyond Performance (oral)

Gonzalo Arce, University of Delaware
Coded Aperture Spectral Imaging for Augmented Reality Projection Systems (oral)

Balakrishnan Prabhakaran, University of Texas Dallas
Estimating Quality of Multi-modal Teleportation Experiences (oral)

Bo Han, AT&T Labs – Research
Immersive Computing for Mobile: Challenges and Research Opportunities (oral)

Sujit Dey, University California San Diego
Enabling Mobile Immersive Experiences Using Edge Computing: Challenges and Research Directions (oral)

Glenn Ricart, US Ignite
US Ignite Smart and Connected Community Experiences with Streaming VR/AR and Community Edges (oral, informal)

Mohamed Hefeeda, Simon Fraser University
Adaptive Streaming of Immersive Multimedia Content (oral)

Yong Liu, New York University
Multi-path Multi-tier 360-degree Video Streaming in 5G Networks (oral)

Hongwei Zhang, Iowa State University
Safety-Critical AR/VR Systems, Wireless Networks and Cross-Community Collaboration (oral)
Predictable Wireless Networking for AV/VR Systems: Perspectives from Wireless-Networked Augmented 3D Vision (poster)

Maxine Brown, University of Illinois Chicago

Magda El Zarky, University of California Irvine

Brendan John, University of Florida
Streaming, Gaze Guidance, and Perceptual methods for 360° Content (poster)

Ravi Prakash, University of Texas Dallas
Network Support for Adaptive 360-degree Video Streaming (poster)

Kimberly Genareau, University of Alabama
Volcanic Lightning: Current Knowledge and Emerging Questions (oral)
Volcanic Lightning vs. Thunderstorm Lightning (poster)

Liu Lei, University of Alabama Birmingham
Improving Accessibility and Affordability of Low Vision Rehabilitation through Networked VR/AR (oral)

Philip A. Chou, 8i Labs, Inc.
AR/VR just had its JPEG Moment (oral)

Jack Brassil, National Science Foundation
Emerging NSF Perspective on VR/AR and Opportunities (oral)

Darleen Fisher, National Science Foundation

Mimi McClure, National Science Foundation

Monisha Ghosh, National Science Foundation

Samee Khan, National Science Foundation

Ann Von Lehmen, National Science Foundation

Meghan Houghton, National Science Foundation

Zhisheng Yan, Georgia State University

Yao Wang, New York University

Zhu Li, University of Missouri Kansas City

Michael Zink, University of Massachusetts Amherst

Craig Gotsman, New Jersey Institute of Technology

K.K. Ramakrishnan, University of California Irvine
Characterizing VR/360° Videos for Network Delivery (oral)

Tamay Aykut, Technical University Munich
Real-Time Stereoscopic 360° Remote Reality (oral)

Sebastian Friston, University College London
Clock Synchronisation Issues for Future Digital Virtual Environments (oral)

Juan Wachs, Purdue University
Telementoring Using the Star System – Teleporting Surgical Expertise (oral and poster)

Gwendal Simon, Institute Mines Telecom Atlantique
Streaming for Multi-Viewpoint 360° Videos (oral)

Torsten Braun, University of Bern
Mobile Edge Computing and Service Function Chaining for Augmented and Virtual Reality (oral)
Mobile Edge Computing for Augmented and Virtual Reality (poster)

Feng Qian, Indiana University
Streaming 360° VR Videos: The Challenges (oral)

Appendix B: Program Schedule

Day 1: April 23

- 8:00am Breakfast and registration
- 8:30am Opening remarks: Jacob Chakareski and Jack Brassil
- 8:45am Short presentations (15 minutes each): *Towards networked VR/AR: Challenges and research opportunities*: Amitabh Varshney, Jacob Chakareski, Klara Nahrstedt, Nikhil Balram
- 10:00am Coffee Break and poster session
- 10:30am Short presentations (15 minutes each): *Towards networked VR/AR: Challenges and research opportunities*: Phil Chou, Sujit Dey, Lixia Zhang, Murat Yuksel, Balakrishnan Prabhakaran
- 12:00pm Lunch and short talks (15 minutes each): Gonzalo Arce, Juan Wachs
- 1:00pm Breakout sessions:
  - Networked VR/AR: Fundamental performance limits and bottlenecks
  - Novel network architectures and protocols
  - Dynamic spatiotemporal capture and data representation
  - Networked immersion quality of experience: Trade-offs between latency and data volume
  - Edge computing, user-centric design, and related trade-offs: Rate/computing/storage/power
  - Novel services, business models, and interdisciplinary application requirements
- 3:00pm Coffee Break and poster session
- 3:30pm Report back from breakout sessions and discussion
- 4:30pm Industry panel: *Applications and systems*: Nikhil Balram, Vishy Swaminathan, Cedric Westphal, Phil Chou, Bo Han
- 7:00pm Dinner and short talks (10 minutes each): Lei Liu, Kim Genareau, Sebastian Friston

Day 2: April 24

- 8:00am Breakfast
- 8:30am Short presentations (15 minutes each): K.K. Ramakrishnan, T. Braun, H. Seferoglu, T. Aykut
- 10:30am Coffee Break and poster session
- 11:00am Report back from breakout sessions, feedback from industry panel, and discussion
- 11:30am Short presentations (15 minutes each): Mohamed Hefeeda, Yong Liu, Gwendal Simon
- 12:30pm Lunch and short talks (15 minutes each): Feng Qian, Cedric Westphal
- 1:30pm Expert panel: *Grand challenges*: Amitabh Varshney, Jacob Chakareski, Sujit Dey, Lixia Zhang
- 3:00pm Report preparation plan and writing assignment
- 3:30pm Small group to assemble initial draft report document led by Jacob Chakareski