



The Effects of Network Outages on User Experience in Augmented Reality Based Remote Collaboration - An Empirical Study

TOOBA AHSEN, Tufts University, USA

ZI YI LIM, Tufts University, USA

AARON L. GARDONY, US Army Combat Capabilities Development Command Soldier Center & Center for Applied Brain and Cognitive Sciences, USA

HOLLY A. TAYLOR, Tufts University & Center for Applied Brain and Cognitive Sciences, USA

JAN P. DE RUITER, Tufts University, USA

FAHAD DOGAR, Tufts University, USA

Augmented Reality (AR) applications can enable geographically distant users to collaborate using shared video feeds or interactive 3D holograms, and may be particularly useful in the socially distant context of the Covid-19 pandemic. However, a good user experience is key for their success and could be negatively impacted by network impairments, which are an inevitable occurrence in today's best-effort Internet. In this paper, we present the findings of an empirical user study, aimed at understanding the effects of network outages, on user experience and behavior, in a collaborative AR task. We highlight how network outages affected users in different ways depending on their role in the collaborative task, and how giving users explicit information about poor network conditions helped them deal with some of these negative effects. Furthermore, we report the strategies that users themselves adopted, to deal with outages, such as batching instructions, or shifting to a different spatial referencing style when communicating with their partners. Lastly, based on our findings, we present some design implications for future remote-collaborative AR applications.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**; **Mixed / augmented reality**; • **Networks** → **Network experimentation**.

Additional Key Words and Phrases: Augmented Reality; Mixed Reality; Remote Collaboration; Distributed Collaboration; Shared Holograms; Spatial Referencing; Common Ground; Network Impairments; Network Outages; User Experience

ACM Reference Format:

Tooba Ahsen, Zi Yi Lim, Aaron L. Gardony, Holly A. Taylor, Jan P. de Ruiter, and Fahad Dogar. 2021. The Effects of Network Outages on User Experience in Augmented Reality Based Remote Collaboration - An Empirical Study. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 313 (October 2021), 27 pages. <https://doi.org/10.1145/3476054>

Authors' addresses: Tooba Ahsen, tooba.ahsen@tufts.edu, Tufts University, Medford, Massachusetts, 02155, USA; Zi Yi Lim, zi_yi.lim@tufts.edu, Tufts University, Medford, Massachusetts, 02155, USA; Aaron L. Gardony, agardony@centerforabcs.org, US Army Combat Capabilities Development Command Soldier Center & Center for Applied Brain and Cognitive Sciences, USA; Holly A. Taylor, holly.taylor@tufts.edu, Tufts University & Center for Applied Brain and Cognitive Sciences, Medford, Massachusetts, 02155, USA; Jan P. de Ruiter, jp@eecs.tufts.edu, Tufts University, Medford, Massachusetts, 02155, USA; Fahad Dogar, fahad@cs.tufts.edu, Tufts University, Medford, Massachusetts, 02155, USA.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

© 2021 Copyright held by the owner/author(s).

2573-0142/2021/10-ART313. <https://doi.org/10.1145/3476054>

1 INTRODUCTION

Supporting remote collaboration has always been important, but the Covid-19 pandemic, and the subsequent shift to work-from-home and remote learning, has made it critical for our society. With many of these trends predicted to outlast the pandemic [42], it is important to consider (and rethink) how remote collaboration can be best supported in future.

Compared to traditional means of collaboration (e.g., video conferencing), augmented reality (AR) and mixed reality (MR) technology can allow for more immersive and content-rich collaborative experiences. For example, remote users have used augmented reality headsets to visualize and analyze complex geo-spatial data, using virtual holograms, and participate in collaborative sense-making [52]. Instructors have also used shared, life-like virtual models of the human body to teach remote students how to perform medical procedures [36], and experts have used AR/MR technology to provide remote assistance to novice users [25, 28, 64]. With the increasing availability of various AR/MR devices [3], as well as services and toolkits to create shared holographic experiences [1, 2, 70], AR-based remote collaboration has the potential to become mainstream in future.

While previous studies have explored various aspects of AR-based remote collaboration [7, 34, 43], the role of the *network* (i.e., the Internet), which allows remote users to communicate, is typically ignored. Given the best-effort nature of today's Internet (i.e., there are no guarantees on data delivery or quality of service), network impairments are inevitable - most of us have experienced this during the pandemic, with the increasing reliance on audio/video calls. Therefore, we posit that network impairments can play a crucial role in determining the effectiveness of remote AR collaboration. Prior research work has explored network impairments in the context of interactive applications, such as audio or video conferencing [37, 40]. Although some of these findings are applicable to certain aspects of an AR collaboration (users may be communicating over voice, or sharing video feeds), AR has some unique characteristics. For example, in AR collaborations involving untethered devices (e.g., HMDs) and interactive, immersive visualizations, users may choose to walk around the room and observe virtual holograms from various perspectives. This could lead to perspective shifts that are unlikely to occur when one user simply observes an egocentric video-feed from another user's device. This may impact tasks that are *spatial* in nature. In light of these characteristics, it is important to conduct first hand studies that *empirically investigate* how *network impairments* affect *user experience and behavior* for *remote AR collaboration*.

Although various kinds of network impairments exist (e.g., packet loss, jitter, etc), our study focused on network outages, which are arguably the most detrimental to user experience [4]. Outages can last from several seconds to minutes, with no data being transferred over the network for the duration of the outage [41, 65]. While the networking research community is actively looking at mitigating the effects of outages - using both proactive and retroactive approaches - outages are still concerning [26, 37, 44]. Anecdotal evidence and network measurement studies show that outages occur all over the world, and could be more severe in the context of developing regions where networking infrastructure may be poor [19, 23, 50, 53].

To summarize, the goals of this research were to understand how outages affected *user experience and user behavior* in AR-based collaboration through shared virtual objects / holograms, and analyze potential strategies for improving user experience (e.g., giving users feedback about outages). To achieve these goals, we implemented an AR-based collaborative application, deployed it on Microsoft HoloLens headsets [54], and used it to conduct a small-scale user study.

Our user study emulated a remote collaborative setting, where pairs of users solved a joint object placement task using 3D virtual objects. Our users formed *worker-helper* dyads; one user acted as the *helper* and provided instructions, while the other acted as the *worker* and performed actions on the virtual objects. We emulated network outages within the task. During the outage

no information (verbal or visual) could be exchanged between the users. Our users completed the task under different conditions; In one condition they were given feedback about the state of the network. The main findings of the study were as follows:

- Our study highlighted the various ways in which network outages affected user experience, based on the role that users play within the interaction. For example, *helpers* expressed frustration at being unable to assist their partners during outages, and *workers* felt uncertain about what actions to take to proceed with the task.
- We found that the majority of users preferred being given feedback about network state as they had more clarity on when communication would be difficult, and it acted as an indicator that work could eventually resume. In contrast, some users preferred it when no feedback was provided because of their perception of outages and personal preference for interfaces. The latter provides avenues for future research and will be discussed more in section 8.
- Based on the analysis of video recordings of the interactions, we found that giving users explicit feedback about the presence of outages could potentially reduce their collaborative effort *during* an outage, by reducing the need to ground task and communication state.
- Finally, our study also revealed that the *helpers* employed common strategies in order to cope with network outages: this included *batching* instructions in order to convey more information during periods of connectivity, and changing the content of instructions by shifting to another *spatial referencing style*. This allowed *workers* to autonomously work even during outages.

While our study focused on remote AR collaboration through shared virtual objects / holograms, some of our findings may be applicable to other forms of collaboration, e.g. video-mediated AR collaboration, or tabletop collaboration. We discuss these possibilities in section 7.

In the subsequent sections of this paper, we discuss related work, our collaborative context, and the design of our AR application and user study. We then discuss our findings regarding the effects of outages, the implications of giving users feedback about network state, and users' coping mechanisms. In the last leg of the paper, we discuss the implications of our findings for the design of future remote-collaborative AR applications, highlight the limitations of our study, and provide future directions for this research.

2 RELATED WORK

Our work consolidates concepts from research conducted in the domains of AR collaboration, networks, spatial referencing and conversational grounding.

2.1 AR-based Collaboration

Prior research in AR-based collaboration can be broadly categorized based on whether the collaboration occurred through shared virtual objects/holograms or through a shared video feed.

When collaborating through shared virtual objects/holograms, users use AR devices/headsets to visualize shared virtual content, and any changes to or manipulations on the shared objects can be seen by all the collaborators. Prior studies have looked at this type of collaboration in the context of education [27, 45, 59], medical training and practice [13, 36, 38] and for visualization of complex data [14, 52]. For example, in the study by Brun et. al, two surgeons simultaneously viewed a 3D hologram of the human heart, using HoloLens headsets, with the project leader manipulating the 3D heart as their co-located collaborator observed [13]. Hamza et. al developed a training tool that uses AR and shared virtual holograms to teach *remote* students how to perform endotracheal intubation [36]. In the context of immersive visualizations and data analysis, Mahmood et al, carried out a study where two users, located in separate rooms, used the Microsoft HoloLens headsets to collaboratively visualize and analyze complex geo-spatial data [52]. Recent work has also looked

into the privacy concerns surrounding the sharing of holographic/virtual AR content between multiple co-located and remote users [67].

In contrast, video-mediated AR collaboration involves remote users viewing the live video feed from a local user's device and providing guidance when necessary (remote assistance). Prior research has explored techniques to make this type of collaboration more efficient; for example, studies have explored the use of shared gaze [8, 9, 51], and different pointing techniques [7, 10, 71] to improve the remote user's situational awareness. Other studies have investigated how audio and tactile cues can be used to guide a local user's gaze or actions [35] and how a remote user can augment the local user's live camera feed with 3D annotations [28, 29]. While our study does not focus on video-mediated AR collaboration, some of our findings may still be applicable here, and will be briefly discussed in section 7.

Although one research study did consider network latency whilst exploring the various factors that affect shared gaze in an AR-based identification task [58], network impairments have largely been overlooked when it comes to remote collaboration and problem solving through AR.

2.2 Network Impairments & User Experience

Traditional audio and video conferencing applications have been the focus of various network QoE (Quality of Experience) studies in the past. For example, prior studies have explored the network-level metrics that could be used to calculate user satisfaction for VoIP services, like Skype [18]. Moreover, prior work has explored back-end solutions to improve the quality of voice calls over the internet [40], and presented insights from measurement studies that could be useful when designing future video conferencing applications [75]. While findings from some of these previous studies regarding audio/video conferencing *may* be applicable to video-mediated AR collaboration, it is difficult to say whether they could also be applicable to scenarios where collaboration occurs through shared virtual objects/holograms. Our research study, therefore, focused on the latter.

To the best of our knowledge, no study has thoroughly focused on investigating the effects of network impairments (e.g., network outages, jitter, etc) on remote collaborative AR applications. A poster paper provided some preliminary findings regarding collaborative AR and network impairments [6]. However our work looks at this phenomenon in greater detail by exploring user experience, coping strategies and the benefits of giving users network state information, through extensive qualitative analysis.

2.3 Virtual Objects and Spatial Referencing

We leverage insights from prior studies that have explored how participants use shared virtual objects to make more effective spatial references, during collaborative AR tasks.

In the field of spatial referencing, research indicates that neutral descriptions (references that can be understood in the same way by all users, no matter their perspective) are more effective when talking about distance/location, as compared to deictic/egocentric references (references that are made from a speaker's own perspective) [69]. In the context of remote AR collaboration, a prior study revealed that virtual objects/landmarks, act as 'universal referencing options' or as 'absolute references' [57], thus facilitating the use of neutral descriptions. A study also revealed that users reduced their use of deictic/egocentric expressions (e.g., "here", "over there"), in co-located AR collaboration, when virtual objects were present in the environment [56].

In our study, we observed that participants also used virtual objects as spatial reference points. This allowed their partners to work autonomously when communication was disrupted, thereby mitigating the negative effects of outages (see section 6.3).

2.4 Grounding

Findings from our study also relate to the concept of grounding. In communication and collaboration, common ground refers to the body of shared information or mutual knowledge, about the task, that is accessible to both parties [20, 21]. This could be about the current state of the task or what steps to perform next in order to achieve the end goal. Grounding, or establishing common ground, is the process of creating or updating this body of shared information and has to be performed continuously throughout the task [21].

Prior studies have explored the challenges involved in dispersed and technology-mediated collaboration, such as users being uncertain about the meaning of silences, or information being unevenly distributed [22]. Studies have also explored mechanisms that support grounding in dispersed collaboration. For example studies have looked at how shared visual information (e.g., a shared view of the task space [33, 66], and a partner's actions [31]) can provide evidence that instructions have been understood and followed, thus providing evidence of grounding.

In our study, we observed users' grounding behavior, during outages, and analyzed how this behavior changed when users were given feedback about network state (see section 6.2.2).

3 OUR COLLABORATIVE CONTEXT

Our collaborative context resembled that of prior studies described in section 2.1. We considered AR-based collaborative scenarios, where remote users interacted with shared virtual objects/3D holograms, and verbally communicated through their AR headsets. We focused on types of collaborations with *asymmetric* roles between users: we had users who performed a largely observatory role, and provided help and guidance, while others manipulated virtual objects. Therefore, our user formed *worker-helper* dyads. We chose this set-up as it may be analogous to common collaborative relationships, such as teacher-student, or expert-novice.

4 RESEARCH QUESTIONS

The overall goal of the study was to understand the effects of network outages on user experience and behavior, within the collaborative context we described above. Specifically, we sought answers to the following high level research questions:

- **RQ 1: What were the effects of network outages on user experience, in light of user roles?** We sought to understand how various aspects of the collaborative experience (e.g., manipulation vs. observation) may be affected differently due to network outages. Did the role that a user played within the collaboration, influence how they were impacted by network outages?
- **RQ 2: Did users prefer being given feedback about network state?** What were the reasons behind a user's preferences for feedback, or a lack thereof?
- **RQ 3: How did user behavior change, if at all, when users were given feedback about network state?** In particular, did users change the number of attempts they made at re-establishing common ground?
- **RQ 4: How did users adapt to network outages?** What coping strategies (if any) did users adopt in order to deal with the negative effects of network outages?

5 EXPERIMENTAL METHODOLOGY

To help answer our research questions, we designed and implemented an AR application that focused on a collaborative problem solving task. We then proceeded to conduct a user study, using the application, over a period of one month. Our study was approved by the Institutional Review Board (IRB). The details can be found below.

5.1 Participants

We recruited 16 users (8 pairs), aged between 18 to 65, from within and around our university community. All users were fluent English speakers, had normal or correct-to-normal vision, and had no history of migraines, epilepsy or seizures.

5.2 Apparatus and Implementation Details

We implemented the application in Unity [74] and deployed it on Microsoft HoloLens headsets [54]. The headsets allowed users to see virtual objects in their real world environment. Their gesture recognition capabilities allowed users to grab these virtual objects using the ‘pinch’ hand gesture (Fig 1), and move them by moving one’s hand forwards/backwards, up/down, or right/left (3DOF).

Network Sharing - The network sharing service, in the Mixed Reality Toolkit, allowed the headsets to communicate over the network [55]. When the application was first launched on a headset, it connected to a broadcast server running on a laptop computer. Each instance of the application, running on the headsets, had its own copy of a few virtual objects. Any time the user moved a virtual object, their headset sent data packets containing the object’s updated cartesian coordinates (positions) to the broadcast server. The server then forwarded the packets to the second connected headset. The instance of the application, running on the second headset, used these data packets to update the position of its own copy of the virtual object that was moved. We also used the “VoiceChat” feature in the Mixed Reality Toolkit so that users could talk to each other over the network, through the microphones in their headsets.

Outages - Outages resulted in a complete cut off of communication between users. The mechanism was as follows: when the *worker* moved the current virtual object beyond a certain distance (mentioned below), the *worker*’s headset would send a message to the *helper*’s headset (via the server), telling it to trigger an outage. No data packets were forwarded to the server, by either headset, for the duration of the outage. At the end of the outage, the *worker*’s headset would send a data packet to the *helper*’s headset (via the server), indicating that the outage has ended and both headsets resumed communication as usual. Each individual outage lasted for 10 seconds. The outages were triggered by distance. As the distance between each virtual object and its target location was approximately 3.5 meters (m), an outage would occur when a virtual object covered a distance of 0.2m, 2.0m and 3.3m. This ensured that users experienced an outage at key points in the task - when the *worker* first started moving the object, when the object was approximately halfway to its target location, and when the object was very near to its target location.

Interpolation - To ensure eventual worldview consistency, and to prevent objects suddenly disappearing out of a *helper*’s line of sight due to position updates after an outage, our application also supported *interpolation*. The interpolation mechanism, which has its roots in online gaming, has previously been used to deal with the after effects of an outage [12]. In our application, we used the object’s current position (before the outage) and updated position (after the outage), to quickly move the object to intermediate positions in the path, until the updated position was reached. This enabled the *helpers* to follow the object’s movements. Given that previous studies have also highlighted the importance of having a consistent view of the task space [12, 31, 47], we believe that interpolation must be a standard feature of any collaborative AR application.

5.3 Study Task: Object Placement Task

We chose an object placement task for two reasons; Firstly, we wanted to keep the task simple, to offset the difficulty some users might have faced if they were new to AR collaboration and working





Fig. 1. The 'pinch' hand gesture allowed users to grab and move 3D virtual objects



(a) *Helper's View* (The *helper* can see both the object to be moved, the red apple, and its target location, the white apple)

(b) *Worker's View* (The target locations for the virtual objects (the white versions of the objects) cannot be seen)

(c) Canvas showing the order of the sub-tasks

Fig. 2. View of the Application

with HMDs. Secondly, re-positioning and placing objects is a basic sub-task used in many types of collaborative and problem solving activities (e.g., urban planning, assembly tasks etc.).

The task involved a *worker-helper* dyad working together to move three virtual objects (a laptop, an apple and a pair of books) from their initial locations to target locations. The key point was that the target locations could only be seen by the *helper*. The *helper* could not move the objects by themselves. They therefore needed to provide suitable instructions to the *worker*, who would then have to place the objects at their correct locations. Any action the *worker* performed on their virtual objects (such as moving the apple to the right), was mimicked in the objects seen by the *helper* (the *helper's* apple would move to the right as well). Figure 2a and 2b show the different views of the *helper* and *worker*, respectively. When a virtual object aligned with its target location, the target location would disappear for the *helper*, indicating that the users could move on to the next object. To ensure consistency across trials, the objects *had to* be moved in a specific order. For example, if the *worker* tried to move the apple, before correctly positioning the laptop, the apple would not move. The order was displayed on a canvas that could be seen by both *helpers* and *workers* at all times (Fig. 2c).

5.4 Study Design

We chose a within-subject design for the study. Each pair of users had to undergo three trials. The first trial acted as a baseline, giving users some experience with the task. As data from trial 2 and trial 3 had to be compared directly, these trials were counterbalanced, to account for learning effect.




Fig. 3. Feedback Mechanisms

Trial 1 - No Outages: In the first trial, the users performed the object placement task, in *perfect* network conditions. The purpose of the trial was to familiarize the users with their partners, and with collaboration in an AR setting. It also gave them an idea of how the collaboration would proceed, if there were no network disruptions.

Trial 2 - Outages: In the second trial, network outages were introduced within the task. A total of 3 outages were introduced for each of the 3 virtual objects that had to be moved. Each outage lasted 10 seconds, during which the audio on both devices was cut off and no data packets containing updated object positions were exchanged. Users were therefore unable to hear each other over their headsets. Moreover, if the *worker* moved a virtual object in any direction, those movements could not be seen by the *helper*. The *helper*'s virtual object would remain in its current position, until the outage ended, and their application received information about the updated position of the object. The interpolation mechanism was active in this condition.

Trial 3 - Outages With Feedback: The third trial was similar to the second trial, but users were also given feedback about the outage while it was underway. Three different feedback mechanisms were used. For the sub-task with the laptop, a static sound would start playing in the users' headsets when the outage started, and would continue to play until the outage ended. For the sub-tasks involving the apple and the books, a banner, showing a message, would appear in the task space when an outage started (Fig: 3a). It would persist until the outage ended, and would then disappear. For the sub-task with the books, the banner also contained a counter that counted down the time until the end of the outage (Fig: 3b). The interpolation mechanism was active in this condition as well.

Pattern of Outages - One consequence of triggering the outages by distance was that if an object was moved too far too quickly, outages would occur back to back. For example, if an object covered a distance of 2.5m very quickly (crossing two of the previously mentioned distance thresholds one by one), the second outage would immediately start after the first outage ended, resulting in the outage lasting for 20 seconds instead of the usual 10 seconds. This meant that for each object they moved, users either experienced three outages of 10 seconds, or one outage of 20 seconds and one of 10 seconds, or just one outage of 30 seconds. The pattern of outages experienced by all the user pairs, based on how they manipulated each virtual object, is shown in Figure 4. For example, the figure shows that when user pair 1 was manipulating the laptop in trial 2, they first experienced a 10 second outage, and then a 20 second outage. On the other hand, user pair 2 experienced three 10 second outages. Having said that, while the pattern of outages may have been slightly different for some users pairs, each pair experienced outages for a cumulative time of 90 seconds, per trial.




Fig. 4. The figures above show the pattern of outages experienced by each user pair, for each of the 3 virtual objects, in each trial. Each cell represents a single outage, and the color depicts the length of the outage.

5.5 Procedure

Users were recruited individually and randomly paired to form *worker-helper* dyads. The roles of *helper* and *worker* were assigned randomly to users within a pair. Informed consent was obtained before the study began.

Users were first shown a video that explained how to grasp and move 3D virtual objects when wearing the HoloLens headset, using the ‘pinch’ hand gesture (Fig 1). They were then asked to practice moving virtual objects using a practice game that we developed specifically for this purpose. Within the game, users had to grab virtual objects (differently colored 3D cubes), and move them to set locations. The game was finished individually. To ensure that users had acquired a standard skill level in moving objects, they were only allowed to proceed with the study once they were able to finish the practice game in under 5 minutes. After completing the practice game, users were moved to separate rooms (to emulate remote settings) so that the trials could begin.

During each trial, videos were recorded through each user’s headset, thus providing the research team with a look into both sides of the interaction (the *helper’s* and the *worker’s*). At the end of the study, users had to fill out an end-of-study questionnaire (discussed further in section 5.6). They were then debriefed about the in-depth goals of the study and any questions were answered.

5.6 Measures

The end-of-study questionnaire was used to gauge users’ opinions about trial 2 and 3. Users were asked to respond to certain questions while keeping in mind their role within the task (*helper* or *worker*). Their responses were used to gain qualitative insight into their experience with outages (to answer RQ 1), their preferences for either trial, and the reasons for their preferences (to answer RQ 2). We used the videos recorded through the headsets to analyze and code for users’ grounding behavior (to answer RQ 3), and to collect anecdotal evidence of coping strategies (to answer RQ 4). Trint, an online video transcription tool, was used to generate transcripts of the video

recordings [73]. To ensure accuracy, the transcripts were then manually reviewed and corrected by the research team.

6 RESULTS

6.1 Effects of Outages on User Experience

The first aim for the study was to understand how users were affected by outages and whether these effects varied depending on the role users played within the interaction. Therefore, we asked users questions about their experiences with the outages that occurred in trial 2, in the end-of-study questionnaire. Since the word ‘outage’ may be confusing for users, we used the phrase ‘breaks in communication’, to refer to the periods of time when an outage was occurring. The relevant questions and responses are discussed below.

Question: Were there moments in the task when you felt that there was a break in the communication (audio) with your partner? If so, please elaborate.

The goal of the first question was to gauge the users’ general impressions of outages. When asked if they noticed any breaks in communication in the task, all users, irrespective of their roles, replied in the affirmative. Outages were described as ‘moments of complete silence’, ‘moments with no communication’, or periods of time when ‘the sound dropped out’ or when the users ‘could not hear (their) partner at all’. One *helper* stated that the ‘breaks in the communication were similar to being disconnected on a cell phone call’.

Besides referring to them as periods of silence, some users described outages in the context of the negative effects they had on task progress and collaborative effort. One *helper* talked about how there were ‘multiple times when the audio cut out’ while they were in the ‘midst’ of ‘giving directions’ to their partner, thereby elaborating on how outages interrupted task progress. Similarly, one *worker* stated that, ‘Yes, there were times when she (the *helper*) went silent and had to restate her instructions when we were connected again after the break’. These responses indicated that *helpers*’ instructions were interrupted due to outages, and were either only partially heard by their partners, or not heard at all. As these instructions had to be repeated when the outage ended, we posit that users had to perform extra work due to the presence of outages. We discuss this further later in this section and in section 6.2.2.

We then asked our users how they might have been affected by outages based on their role in the collaborative task (RQ 1). We divided the responses into two groups according to user roles (*helper* and *worker*), and looked for common themes in each group. Relevant quotes are mentioned in the discussion below as supporting evidence.

Question: In your capacity as a (*helper*, *worker*), how did these breaks affect you and your communication?

Helpers - Based on the responses, outages affected *helpers* in the following ways:

- They could not carry out their core responsibility (giving instructions)
- They had to perform extra work (repeating instructions)
- They could not accurately interpret situational cues or task state

Helpers could not carry out their core responsibility (giving instructions) - *Helpers*’ responses indicated that outages were affecting their ability to provide complete instructions to their partners, which was their core responsibility within the task. One *helper* stated that ‘it was annoying’ when

they could not convey the 'last part of their sentence or direction' before the outage started. They described feeling as if they were 'leaving their partner hanging'. Another user commented on the unpredictability of outages; they stated that as 'the breaks in communication were random', they were unable to 'anticipate' when they might happen, thus leaving them 'unable to complete' their instructions. The same *helper* also stated that, as they were 'unable to instruct' their partner, their partner would sometimes 'keep moving the object', thus sometimes taking actions that would be 'counter-productive' to the overall progress of the task.

Helpers had to perform extra work (repeating instructions) - *Helpers* alluded to the extra work they had to perform, when repeating instructions, as they did not always notice that an outage had started. They would therefore be unaware that their partner could not hear them and had to repeat instructions once the outage ended. One *helper* stated, 'I didn't know when my partner couldn't hear me and so I would continue giving instructions without her responding'. Another *helper* expressed the same sentiment. They stated that, 'Each break meant I had to keep the instruction I was about to say in mind, and often had to say it several times, first time not realizing the other party could not hear, and then again after the break was over.'

Helpers could not accurately interpret situational cues or task state - *Helpers* also alluded to their inability to accurately interpret situational cues during outages. For example, a *helper* stated that they could not determine whether the lack of response from the *worker* was because they were 'struggling to pick up the object', or because they had 'just stopped with the object', or 'if it was a technical issue'. Therefore, guiding their partner would sometimes be 'a little difficult'. Another *helper* stated that they felt frustrated at having to constantly confirm the position of the object and state of the task, after each outage ended. Their exact words were as follows; 'The breaks became annoying as the study continued. I felt I needed to backtrack my last train of thought with my partner to ensure we were on the same page as far as where the object was last placed.'

Workers - While the *workers* retained their ability to move objects even when the network was unavailable, they reportedly experienced the following effects of outages:

- They experienced increased uncertainty in their actions
- They were unable to ask for help
- They were forced to slow down or stop their actions

Workers experienced increased uncertainty in their actions - Several *workers* mentioned feeling unsure about how to proceed with the task during outages as instructions were cut off. One *worker* stated that, 'I felt at a loss for what to do and unsure what to do to communicate (to their partner) that I could no longer respond to them.' Another *worker* expressed the same sentiment of uncertainty by stating that 'there were times when I was in the middle of moving an object in one direction and I wouldn't be sure if I should keep moving it or not'. They went on to say that the randomness of the breaks 'made it harder to know if (they) were about to lose communication'. One *worker* expressed uncertainty, but more so for their partner and less so in their own actions. They stated that they were 'not sure' if their partner could see what they were seeing, which was their 'only concern'.

Workers were unable to ask for help - *Workers* highlighted their inability to ask for help or acquire more information about the task, which hampered task progress. One *worker* stated that outages affected their 'ability in completing the tasks' since they were 'missing information'. Another *worker* stated, 'I kept asking for some instructions and it kept cutting out when he was in the

middle of explaining something or telling me if I was heading in the right direction’.

Workers were forced to slow down or stop their actions - As mentioned previously, *workers* were unable to seek help from their partners during outages. As a result, they were forced to change the way they manipulated objects; either by slowing down object movements, or stopping them altogether. One *worker* stated that, ‘I couldn’t hear the instructions of my partner and therefore couldn’t know where to move the objects. I had to simply wait for her to be reconnected so that I could continue with the task.’ Similarly, a *worker* stated that they had to ‘pause during the breaks’ and that they ‘went much slower’ with their movements to ensure that they were not ‘messing up’.

Out of all the responses, only one *helper* and one *worker*, from different user pairs, expressed minimal concern about outages. The *helper* stated that they did not feel that the breaks affected them in a ‘significant way’ and the *worker* described the outages as a simple ‘time obstacle’.

To summarize, it seems that outages affected users differently based on their roles. The *helpers’ instructions were being cut off*, they experienced *frustration or difficulty* with keeping updated on the state of the task, and had to *perform extra work* to repeat instructions when they were not fully received by their partner. The *workers*, while still being able to move objects, became *uncertain* with how to proceed with the task, *could not ask for help* to reduce this uncertainty, and as a result, *changed their behavior* (stopping or slowing down object movements).

In the next section, we discuss how providing users with feedback about network state helped mitigate some of these negative effects, and caused a change in users’ collaborative behavior.

6.2 Providing Users with Feedback About Network State

We adopted a twofold approach when evaluating the effectiveness of providing users with feedback about network state. Firstly, we analyzed the users’ responses to the end-of-study questionnaire, to broadly highlight the benefits of feedback. We then used the video recordings of the interactions and behavioral coding strategies, to highlight the change in the users’ collaborative behavior (specifically, their attempts at re-establishing common ground), on receiving feedback about network state.

6.2.1 User Preferences: Feedback Versus No Feedback.

To gauge whether users themselves saw any benefit in the given feedback about network state (RQ 2), we asked our users, in the end-of study questionnaire, which trial they preferred, the one with feedback or the one without any feedback.

A majority of our users (11) preferred the condition with feedback, whereas 4 users preferred the condition without any feedback, and one user critiqued both conditions but did not state a clear preference. We divided the responses based on user preferences (feedback or no feedback), and then looked for reasons for these preferences. Responses that alluded to the same reasons are grouped together in the discussions below.

Reasons for a Preference for Feedback - Users preferred the trial with feedback because of the following reasons:

- Clarity about the state of communication
- Users were warned about times when communication would be difficult
- Users could seamlessly resume communication

Clarity about the state of communication - Many of the users who preferred the trial with feedback, did so because they had a better idea about the state of communication. For example, one *worker* stated that they preferred the trial with feedback because ‘there were messages saying what was

going on'. Another *worker* talked about how the system seemed 'more aware of when things went wrong', no doubt referring to the static sound and message banners that appeared when outages occurred. Another *worker* stated that 'even though the (static) sound was startling', they preferred it to the silence, as they knew when communication was cut off and did not have to 'continue asking (their) partner whether she could hear (their) voice'. Some users also appreciated knowing the reason behind the communication breakdown. For example, one *worker* stated that 'it was more clear to me that the disconnection was due to the network issue', so they would 'stop and wait' until 'it re-connected'. Moreover, one user stated that the term 'network error', that was displayed on the message banners, was 'familiar' to them. The familiarity of the term might have helped them understand why they were suddenly unable to communicate with their partner.

Users were warned when communication would be difficult - One *worker* preferred the trial with feedback because they were given more warning when something went wrong. They stated that, 'Even though the (static) sounds and lost internet tags would pop up, whenever I noticed them it acted as a warning that communication would be hard for the moment, and I knew to stop moving the object very much.' They stated that in the other trial, they 'would lose connection but without warning so it was harder to know whether or not to keep going.' Therefore, the feedback acted as a warning sign and helped them make better judgements about whether they should continue with the task or not.

Users could seamlessly resume communication - Various users alluded to the ease with which communication could resume when feedback was provided. One *helper* mentioned that the banner with the 'countdown' gave them an idea of 'when the connection would be restored'. Similarly, the disappearance of the static sound and message banners, served as an indication that the 'communication was back on'. In the words of two of the *helpers*, this allowed them to 'continue guiding (their partner) immediately', or to 'continue where (they had) left off', thus allowing them to seamlessly proceed with the task. Another *helper* also talked about how the feedback served as a 'positive indicator', projecting the idea that work could eventually resume.

Reasons for a Preference for No Feedback - Users who chose the trial without any feedback, gave the following reasons for their preference:

- Annoyance at having to wait for the countdown to end
- Personal preference for feedback mechanisms
- Perception of the severity of outages

Annoyance at having to wait for the countdown to end - Out of the four users who preferred the trial without any feedback, one *helper* stated that 'it was a bit annoying having to wait for the 10 seconds to go by in those internet issues', alluding to the counter that counted down from 10 seconds until the outage ended. The *helper* stated that in the previous condition, 'the audio didn't really have an effect on the object's movement. Once the audio came back, my partner was still moving and I could help her more easily'. The presence of the counter might have made this user pair believe that they had to wait for the outage to end, before continuing to move the object (even though the *worker* could still move the objects). This comment seemed interesting as, in contrast, another *helper* mentioned that they preferred the feedback precisely *because* it prevented their partner from making unnecessary movements or taking 'counterproductive' actions. They stated that, 'we both lost ability at the same time and could continue where we left off instead of having to add more instructions because of extra movement by the *worker* during loss of communication'. This begs the question of whether *workers* being able to autonomously work, while communication

is disrupted, is always good for the overall progress of the task. Our discussion in section 6.3 investigates this further by considering how giving instructions using virtual object-based references may give users a better idea of distance, and help them work autonomously. Alternatively, since AR applications can be content-rich, this may also serve as motivation to look into feedback that not only makes users aware of network issues, but diverts a user's attention to local sub-tasks that do not require much assistance from the *helper*. This could allow the task to progress, all the while preventing actions that might be counterproductive. We discuss this further in section 7, in the context of design implications.

Personal preference for feedback mechanism - One *helper* stated that they preferred the trial with no feedback as 'nothing appeared on the screen'. Similarly, the user who remained neutral stated that the feedback 'felt so forced, and the timers were quite long', and that 'just losing connection', in the other condition, 'at least felt natural'. As these comments had to do with the interface design, it might be beneficial to explore the best interface to use when displaying information in AR, considering that multiple sensory modalities are at play (auditory, and visual). Headsets that make use of handheld controllers might have tactile feedback available to them as well. As we cannot answer these questions at present, we leave this for future work.

Perception of the severity of outages - Only two users alluded to the severity of the network disruptions; the *helper* from pair 5 stated that the trial without feedback had 'fewer network disruptions', and the *worker* from pair 6 stated that the 'audio issues weren't as disturbing'. As both conditions had outages for a total of 90 seconds, the overall time for which the users were disconnected from each other remained the same. Having said that, as outages for each object were triggered by the distance the object had moved, some outages occurred back to back, thus appearing to be longer than 10 seconds (as mentioned in section 5.4). This might have resulted in these comments about the severity of the outages. Looking at Figure 4, it seems that users from pair 5 experienced *fewer, but longer*, outages in the trial without feedback. Similarly user pair 6 experienced *more* outages in the trial without feedback, but the outages were *shorter in length*. This suggests that some users may tolerate outages that are longer if they happen less frequently, while other might have a higher tolerance for shorter, more frequent, outages. Therefore, the effects of the duration or frequency of outages on user experience may not be as simple as a one size fits all policy. Future studies should explore this in more detail.

In summary, a majority of our users preferred the trial with feedback, because the feedback made the state of communication clear, and acted as a warning sign. Moreover, the feedback was transitory, and provided them with an indication of when the task could resume. Recall that one of the negative effects of outages on *helpers* was that they were unable to accurately interpret situational cues or task state. Providing users with feedback seemed to have helped in this regard. Users who preferred the trial without feedback, did so because they either felt that they had to stop with the task and wait for the outage to end or did not like the feedback mechanism itself or they perceived the outages to be longer. These reasons provide interesting directions for future work, and are discussed in more detail in section 8.

6.2.2 *Change in Users' Collaborative Behavior on Receiving Feedback About Outages.*

Our next goal was to understand how our users' collaborative behavior changed when they were provided feedback about network state (RQ 3). To that end, we first looked through the video recordings of the trial with outages but no feedback, and observed user behavior *during* the outages. We found that if the *helpers* had asked a question or given an instruction, and did not receive

either an explicit or an implicit response, our *helpers* would often repeat the question or instruction. They would continue to do so at small time intervals until their partner responded, thus signalling the end of the outage. The example below shows a *helper* asking questions repeatedly, during an outage, to determine if their partner could hear them. Outages in the transcript are indicated by the keywords **Outage starts** and **Outage ends**. Note that the statements that occur in the middle of the outage go unheard by the opposite party.

Outage starts

Helper: How's it going?

Helper: Can you hear?

Helper: Hello?

Helper: Hello?

Helper: Hello?

Outage ends

Similarly, in the example below, the *helper* kept repeating the same instruction, until either the *worker* responded, or the *helper* saw the object moving. They then proceeded to give more instructions.

Worker: Yeah yeah. It's to the right.

Outage starts

Helper: All right More right.

Helper: More right.

Helper: More right.

Outage ends

Helper: [Takes partner's name in a questioning tone]

Worker: Yeah yeah yeah.

[*Helper* saw the object moving]

Helper: Yeah hold it hold it hold it. Bring it inwards towards you.

Each of these questions or statements that were made during the outage, acted as an attempt by the user to re-establish common ground with their partner. A lack of response indicated that communication was still not possible. The user then opted to repeat instructions or questions (constantly probing the network), until a response was received, and grounding occurred. This added to the work that they had to perform during the collaborative task, which may have added to their collaborative effort.

This begged the question, of whether exposing network state to the users, through some feedback mechanism, helped reduce this extra work. The idea being that the feedback would add to a user's knowledge about the state of communication, so they might not expend effort in trying to re-establish common ground at a time when it would be difficult to do so. Therefore, we hypothesized that we would see a reduction in the number of attempts at re-establishing common ground, when users were given feedback about network state (trial 3), as compared to the trial when no such feedback was provided (trial 2). Prior studies have also looked at ways of reducing collaborative effort by providing users with visual information and allowing them to use language more efficiently [32], although they did not do so in the context of network impairments.

Behavior Coding Process - To test our hypothesis, we went over the video recordings for the trials with and without feedback, and counted the number of attempts made by each user to re-establish common ground *during outages*.

We first demarcated the users' verbal speech into separate turns, using the definitions of turns and turn construction units that were presented in the work by Sacks et al., and are widely used in the field of conversational analysis [68]. Turn-taking refers to the process of handing over the floor to other conversational partners so that they may speak. The piece of speech that a turn comprises of, is known as a turn construction unit (TCU). At the end of a TCU, the turn may be handed to other conversational partners or may be taken up again by the same person (self-selection), should they not receive a reply [68]. For example, if we look back to the first transcript mentioned in this section, the user asked 5 questions in the middle of the outage with pauses in between each question (where they waited for a reply). These count as 5 individual turns. Two coders (where one coder was blind to the study conditions), independently demarcated each user's speech during outages into individual turns. This was done for both trials. Any disputes were resolved by discussion until an agreement was reached.

Once the individual turns were demarcated, we analyzed their content (what the user was actually saying). We found that in some turns, users were actively trying to re-establish common ground with their partners by asking them a question ('Can you hear me?', 'Hello?'), or giving them an instruction ('Move it to the right'). However, in other turns, users were simply verbally expressing their frustration at being unable to communicate. For example, two users said, 'Oh Oh' or 'Ugh', when they encountered yet another outage. Similarly, one user, during an outage, realized that their partner could not hear them and mumbled to themselves, 'You can't hear me. Okay, this is brilliant'. These statements were not intended for a user's partner, and therefore, were not counted as valid attempts at re-establishing common ground. The two coders categorized the turns into valid and invalid attempts at re-establishing common ground. When calculating inter-rater reliability, the value of Cohen's Kappa was 0.81, thus indicating a good amount of agreement between the two coders [49]. We used the turns that counted as valid attempts at re-establishing common ground, in the analysis described below.

Coding Analysis - We analysed the data to see whether there was a change in the number of users' attempts at re-establishing common ground, during outages, when feedback was either present or absent. Since this was a two trials (no feedback and feedback) into two roles (*helpers* and *workers*) design, we performed a Two Way Mixed ANOVA test. The significance level was set to 0.05.

Meeting Assumptions For a Mixed ANOVA - The Levene's test of homogeneity indicates that variances were homogeneous for both levels of the within subjects factor - no feedback ($F(1,14) = 3.028$, $p = 0.104$) and with feedback ($F(1,14) = 1.014$, $p = 0.331$) - as both p values are above the 0.05 threshold (thereby supporting the null hypothesis). Similarly, the Shapiro Wilke's test confirmed that the data was normally distributed for both levels, with p values for both trials - no feedback ($F(16) = 0.939$, $p = 0.339$) and feedback ($F(16) = 0.906$, $p = 0.099$) - being above the 0.05 threshold. Since we only had two levels of the within subjects factor, the test for sphericity (which requires at least 3 levels or more of a within subjects factor) is not applicable.

Results - The results of the mixed ANOVA found no significant interaction effect between role and trial type ($F(1, 14) = 0.206$, $p = .657$, $\eta_p^2 = 0.014$) and no significant main effect of role ($F(1, 14) = 0.977$, $p = 0.340$, $\eta_p^2 = 0.065$), the between subjects factor. However, there was a significant main effect of trial type ($F(1, 14) = 13.18$, $p = 0.003$, $\eta_p^2 = 0.49$), the within subjects factor, on the




Fig. 5. Mean (And Standard Error) For Number of Grounding Attempts Made By Users, For Each Trial

users' mean number of attempts at grounding. The mean values and associated standard errors are shown in Figure 5. The results show that our users (both *helpers* and *workers* collectively), made significantly less attempts at re-establishing common ground, in the trial with feedback (mean = 5.56), as compared to the trial without feedback (mean = 9.56).

Therefore, it seems that giving users information about the state of the network, during outages, was useful as they did not have to expend as much effort in probing the network and trying to re-establish common ground, as they otherwise had to expend when no feedback was provided. This may have led to a reduction in their collaborative effort. This is interesting, as the principle of least collaborative effort, found in grounding and communication theory, indicates that within a contribution, users try to minimize the total effort spent on that contribution [20]. Providing users with information about the state of the network, in our case, may have aided in doing just that.

Thus far, we have discussed how providing users information about outages, through a feedback mechanism, can help mitigate some of the negative effects of outages. In the next section, we discuss some strategies that users themselves adopted in order to deal with these negative effects.

6.3 Adapting to Network Outages: Coping Strategies

Recall that in section 6.1, *workers* mentioned their uncertainty about how to proceed with the task during outages. To alleviate some of this uncertainty, and their own inability to provide assistance during outages, *helpers* adopted coping strategies that provided *workers* with more information during periods of connectivity. This allowed *workers* to work more autonomously while outages were underway. In this section, we attempt to answer RQ 4 by discussing these two coping strategies: i) batching instructions and ii) using alternate spatial references. We provide anecdotal evidence for each strategy, using transcripts from our video recordings.

6.3.1 Batching Instructions. - Batching instructions means giving multiple instructions together instead of giving instructions one by one. Users began to batch instructions in order to get more information across to their partner when the network was available. There is various anecdotal evidence of this phenomenon scattered throughout the videos of the interactions¹. For example, the *helper*, from user pair 5, batched instructions when prompted by the *worker*. The excerpt from

¹Video Example: <https://toobaahsen.github.io/Collaborative-AR-Under-Internet-Issues/>

the transcript for this interaction is given below.

Outage ends

Helper: OK now you move it to the right.

Worker: Give me maybe more directions cause you're gonna cut out.

Helper: And you're gonna have to. Just a bit more to the right and then it's a bit behind.

In this instance, the *worker* prompted the *helper* to give more instructions because they may 'cut out' soon, thus highlighting how users had begun expecting network outages and were thus proactively trying to reduce their negative effects. Another example, from the same user pair, is given below.

Outage starts.

Helper: And now bring it forward. Like closer to your body.

Outage ends.

*Worker: I can't hear you so you better give your instructions **real quick**.*

Helper: Okay. Move it closer to your body. Bring it forward.

In this case, the *helper* was giving instructions when an outage occurred, so the *worker* was unable to hear the first instruction. Once the outage was over, the *worker* prompted the *helper* to quickly give more instructions. It is also interesting to note the state of urgency the *worker's* words induced in the *helper*. While the first time around, the two instructions the *helper* gave took 4.8 seconds to relay, the *helper* took less time (2.7 seconds) when repeating the same instructions after the outage.

In other instances, *helpers* batched instructions not only to relay more information, but also as a corrective measure in case *workers* moved too far away from the target while the outage occurred, or if they still thought that the previous instruction was valid and continued to follow it. An example is given below.

Helper: Next is Apple.

Worker: Mm hmm.

Helper: So you want to go basically just to your left for now.

Outage starts

[*Helper* remains quiet]

Outage ends

Worker: To the left?

Helper: Hold on. You want to go down and to your left just a little bit. Like maybe a foot.

In this case, after the outage, the *worker* asked 'To the left?', being under the impression that the object still had a ways to go. The *helper* replied with a batch of corrective instructions to make sure the *worker* did not move the object in the wrong direction.

Therefore, *helpers*, whether on their own or after prompts from the *workers*, seemed to be batching their instructions as a way to proactively reduce the effects of outages.

6.3.2 Adapting Spatial Referencing Styles. - *Helpers* shifted their spatial referencing style when giving instructions, to allow the *workers* to work more autonomously during outages.

While reviewing the video recordings, we observed that *helpers* typically made use of both *directional references* and *virtual object-based references* [77]. For example, in some instances, a

helper asked the *worker* to move the object ‘to the right’, or to ‘move it up’. These references were largely egocentric, with the *helper* giving instructions relative to their own self. In contrast, virtual object-based referencing involved the use of 3D virtual objects², such as the ‘laptop’ or the ‘apple’, as spatial reference points. For example, some *helpers* gave the following instructions to their partners: ‘(the books go) behind the laptop’, or ‘the apple needs to go to the left. So near where you picked the laptop up from’ or ‘You want to bring the laptop - you’re gonna go past the apple. Just to give you like a reference point.’

As directional references were more general, *helpers* often augmented them with instructions such as ‘keep going’ or ‘more right’, until the object was moved to the correct location. This was so because directional references did not help one gauge distance. The *worker* may have moved the object to the right, but needed the *helper*’s assistance in determining how far to the right. In contrast, if users used the virtual objects as spatial markers, or checkpoints, then *workers* had a better idea of the distance that needed to be covered, or the approximate spatial position that the object had to be moved to. Therefore, they could work autonomously even if outages occurred and communication broke down³.

We went over the video recordings and collected anecdotal evidence of when *helpers* might have switched from directional to object-based references when encountering an outage, and how they were being used by the *helpers* to give their partners more autonomy. Examples are given below. Virtual object-based references have been underlined.

Helper: You’ve got the computer?

Worker: Yes.

Helper: Cool. We’re bringing it down and to the right again.

Outage starts

Helper: Cool. How’s that coming?

Outage ends

Helper Hold on one sec.

Worker: Okay.

Helper: Yeah it’s cutting. It’s cutting us out.

Worker: Okay.

Helper: Okay. So is your computer currently near the apple?

Worker: Yes.

In this example, the users experienced a network outage, and the *helper* switched from directional to object based referencing, to better gauge the object’s spatial position.

Moreover, spatial references could also be used to break down the task into smaller components, such as moving the object to some halfway point first. This allowed *workers* to autonomously continue with the task even if an outage occurred. Once this was done, the *helper* could give more instructions. An example, from the perspective of the *worker*, is given below.

Helper: Ok do you remember where we dropped off the laptop or do you see the laptop?

Worker: Yeah.

Helper: Take the books -

²Besides using virtual objects as referents, one pair also tried to use the actual physical objects in their environment, such as doors and windows, in order to orient themselves and give instructions. However as the environments the users were in differed, this user pair had some difficulty with perspective shifts, and they eventually shifted to object-based referencing using virtual objects.

³Video Example: <https://toobaahsen.github.io/Collaborative-AR-Under-Internet-Issues/>

Worker: Do you want me to put them over there?

Helper: Yeah put them like there.

Outage starts

Worker: [Moves the books up to the laptop]

Worker: And now books are by the laptop, but you are gone.

Outage ends

Helper: OK. Hello?

Worker: Hi. I can hear you.

Helper: OK, I can hear you OK. Yes. So more. OK. OK, cool. Go up I think.

In this example, the *worker* kept moving the books for the majority of the outage, until they reached the laptop ('And now books are by the laptop...'), and then waited for the connection to be restored. When the outage ended, the *helper* saw that the books were near the laptop ('OK. OK cool.'), and continued with the next instruction ('Go up I think.'). Recall that *workers* mentioned feeling uncertain with their actions when outages occurred (section 6.1). Having a spatial reference point seemed to have alleviated the uncertainty in this case.

In one interaction, besides using other virtual objects as reference points, the *helper* also used the dimensions of the virtual object that was being moving (the books) as a measure of distance.

Helper: Yeah keep moving it. More more more. About 5 book lengths towards your right.

To summarize, it seemed that *helpers* consciously changed their behavior (batching instructions, or switching to virtual object-based references) to deal with the negative effects of outages. They initiated this change in behavior without any prompts or intervention from the application.

7 DISCUSSION

The goals of our study were to understand how users' collaborative behavior was affected by network outages, and to analyze strategies to mitigate some of these adverse effects (e.g, providing users with feedback about network state). In this section, we use insights from our findings to discuss *design implications* for future remote collaborative AR applications.

How to Best Utilize Free Time During Outages - Recall that one benefit of providing users with feedback during outages was that it reduced their attempts at re-establishing common ground during outages (section 6.2.2). However, while this does reduce collaborative effort, one can posit that having to simply stay quiet and wait for outages to end (especially for longer outages), may cause users some frustration, uncertainty, or a feeling of making no progress on the shared task. This begs the following question: How do we best utilize this time, during outages, when users may be quiet and waiting for the connection to be restored? We discuss two potential strategies in this regard:

- Shifting a user's focus to 'private' virtual spaces
- Distracting the user through gamification

Private and shared virtual spaces - One strategy is to shift a user's focus away from 'shared' virtual objects to those in a 'private' virtual space. For example, users could have replicas of the virtual objects being used for collaboration, in their 'private' space (or they could create these replicas at will). When the application detects an outage, it can prompt users to shift focus from the shared space to the private space and perform local sub-tasks (e.g, reviewing the steps to a surgical

procedure that a professor just demonstrated using shared 3D holograms). This concept of ‘private’ and ‘shared’ spaces appeared in the MR application developed by Mahmood et. al [52], where each user could create 3D data visualizations locally in their private space and manipulate them at will. Users could then move private visualizations into the shared space to allow their partner to interact with them.

Distracting the user through gamification - Application designers could install a handful of short AR-based games into their applications. When an outage begins, users could be given the option to select and play any of these games. Considering the immersive nature of AR technology, the games might be engaging enough that users do not feel overly frustrated when waiting for the network to be restored. Moreover, adaptive avatars have been used in MR collaboration to facilitate non-verbal cues between a local AR and a remote VR user [61]. Future work could explore how these avatars could be used to distract the user, or mimic the behaviour of a remote user, while an outage is underway.

How to Deal with Other Types of Network Impairments - While our study only considered full outages, users may encounter other types of network impairments:

- Partial outages - Only audio or visual communication available
- Packet loss/short frequent outages - Distorted audio/virtual objects showing jumpy movements

Partial outages (no visual communication) - In the case where users can hear each other but cannot see how virtual objects are being moved by their partners, users will have an inconsistent view of the task space. Therefore, *helpers* may require *workers* to verbally acknowledge their understanding of the instructions and describe the current state of task space [33]. Using virtual objects as landmarks and spatial references may be an effective strategy here as well (as mentioned in section 6.3.2). *Workers* could use these virtual objects to describe the current state of the task space (‘The apple is by the books’), and the *helper* could use them to give instructions (‘Bring the apple to around where the books are’).

Partial outages (no audio communication) - In cases where users can no longer hear each other, but can still see model manipulations, non-verbal cues could be used as an alternate method of communication. Gaze and gesture information could be displayed to collaborators through virtual avatars [46, 63]. A *helper* could then gaze at, or gesture to the object they want the *worker* to move. Application designers could also allow users to create virtual landmarks (3D arrow, checkpoints, etc) in the task space, to indicate an object’s movement trajectory or end position / orientation. That said, although prior studies have used gaze or gestural cues *alongside* audio [11, 60], future studies may need to investigate their effectiveness when audio is completely cut off or distorted. In fact, in their study on sharing awareness cues in remote MR collaboration, Piumsomboon et al. noticed that participants were verbally describing their thoughts or actions to their partner and constantly verbally communicating with each other [60]. It would be interesting to explore how the suggested awareness/non-verbal cues could help users progress in the task when this constant verbal communication is interrupted due to network impairments.

Packet losses, short frequent outages (distorted audio and jumpy virtual objects) - Although strategies like batching instructions (section 6.3.1) may be ineffective in the face of distorted audio, non-verbal cues, as mentioned above, could still aid in communication. Similarly, the ‘interpolation’ method (discussed briefly in section 5.2), could smooth out an object’s jumpy transformations (suddenly changing position, orientation or size) when packet losses are high [12]. Another strategy

is to utilize network prioritization, a mechanism that is implemented in various network transport protocols [17, 48]. The application would prioritize sending either visual or audio data so that at least some information makes it through to one's collaborator. One challenge here is to determine which data gets prioritized over the other, if at all. This decision could be made by the application designers, or left up to the users themselves (through a suitable interface).

Which Feedback Mechanisms Are More Effective? - While we cannot definitively say whether one type of feedback (audio vs visual) was more effective than the other - a shortcoming of our study design - we discuss some design implications based on user responses and our observations.

Audio and Visual Feedback - When giving users information about network state, familiar terms might make the meaning of the message more clear (refer to section 6.2.1). However, visual feedback may not be preferred by everyone. Recall that one user did not like the visual banners that appeared in the task. As our application displayed the visual banners right next to the virtual objects themselves, users might have been thrown off when something new appeared in their task space. An alternative may be to display a network connectivity icon on the screen of the user's device, or place the banner in a fixed location in the virtual task space. While users may run into a situation where they are too busy with the task to notice any changes to the icon or fixed banner, one way to solve this is to use temporary indicators (e.g. an arrow that points towards the icon/banner), to draw a user's attention towards the icon/banner depicting the network state.

Significance of Countdowns - The 'countdown' based feedback mechanism was perhaps more forward-looking in nature. We posited that if future applications could predict the duration of outages, and convey them to users, it may help improve collaborative behavior. During our study, we did observe that a few *helpers* would focus on the countdown and immediately start giving instructions when the counter hit 0. However, one user did express annoyance at having to wait for the countdown to finish, and another thought that the mechanism felt 'too forced' (refer to section 6.2.1). There is also a question of how prediction accuracy may impact user experience. For example, users may get frustrated if the outage ends later than predicted. Future research could explore the effects of prediction accuracy (for outage duration) on user experience, in greater detail.

When Should Network State Be Exposed To Users - It is possible that by exposing the entirety of outages through our feedback mechanism, users were noticing outages that they would have otherwise missed. For example, if a *worker* was focused on moving a specific object or the users were currently not talking, they might have only noticed a part of the outage or not all. This begs the question; can we hide some outages from users? One approach is described below.

Predicting An Object's Movements - One way to hide the effects of outages, on object movements, is to predict where it will end up based on its current position and velocity. This technique, called dead reckoning, is used in multiplayer games to hide the effects of latency from users [5]. For example, when an outage starts and the application detects that users are not currently talking, dead reckoning could be used to show the *helper* (the observer of the object manipulations), a predicted path that the virtual object is being moved in. If the users start talking, the application could inform them of the network impairment. If the outage ends without the users attempting to talk, then the model's position will simply be updated in real time, and the users will have come away from this experience without noticing that an outage occurred. This may give users a better

overall impression of the collaboration.

This strategy may also be useful in reducing the network overhead incurred from sending frequent update information for virtual objects, especially when network resources are constrained. Prior work has investigated this in the context of human avatars and virtual environments [16].

Applicability of Our Findings on Other Types of Remote Collaboration - While our user study focused on collaborative AR applications using virtual objects, some of our findings may be applicable in other collaborative scenarios.

Video-Mediated AR Collaboration - In video-mediated AR collaboration, outages could disrupt verbal and visual communication; the remote user's video feed might freeze and their audio might cut out. Short but frequent outages or packet losses, could lead to distorted audio and video as well. Similar to our findings, users may face feelings of uncertainty and might exert effort in trying to constantly re-establish common ground. Some of the techniques we mentioned previously might be useful here; for example interpolation could be used to 'fill the gaps' between the latest and last received video frame [39], or one type of data (audio or video) could be prioritized over the other. Having said that, one way video-mediated collaboration is different from AR collaboration through virtual objects, is that when an outage occurs and the video freezes, remote users lose access to the shared task space (the local user's environment). One way to deal with this is to allow the remote user to view a 3D reconstructed model of the local user's environment, independently from the live camera feed, as done in prior work on remote AR and MR collaboration [28, 29, 72]. While the outage is underway and the remote user is not getting any updates to the live video feed, they could use this 3D model to analyze the local user's environment and plan next steps. Moreover, the emergence of 360 videos in mixed reality [62, 72], means that remote users have access to a larger view of the local user's environment. While outages will still prevent any updates to the live video feed, remote users could use the latest version of the 360 video they obtained to explore the local user's environment until communication between the users is restored.

Tabletop Collaboration - Tabletop collaboration can take various forms; users may interact with virtual or physical objects (within the confines of the tabletop), with the remote user's hand gestures being projected on top of the table through video [30, 76]. As in any type of video-mediated collaboration, network impairments could cause audio and/or video to freeze or become distorted. Techniques like interpolation or data prioritization may be useful here as well. For example, users may prioritize seeing their partner's hand gestures over hearing their voice, in poor network conditions. Recent studies have also looked at hybrid systems that combine tabletop interaction techniques familiar to users, with complex 3D visualizations that can be viewed above the tabletop using AR headsets [15, 24]. While these systems are for co-located users, it would be interesting to see how many of our insights (related to remote AR interaction through virtual objects) apply to future work that explores these hybrid systems in a remote setting.

8 LIMITATIONS & FUTURE WORK

Our study had some limitations, and raised some interesting research questions for the future. Firstly, our study only looked at *worker-helper* dyads. Now that some insights common to each role have been gained, it may be useful to conduct studies where users have more variety in their interactions. Moreover, our study only focused on one type of network impairment (outages), and one type of action that a user could perform on the virtual objects (moving the object). It may be beneficial to look at other types of network impairments in the context of remote collaboration.

Lastly, our study did not look at the effects of the duration of outages on user experience. Do users perceive outages as being more or less severe, based on their lengths? Some of our users talked about how they were unable to notice when some outages started, in the trial without feedback. This begs the question of whether we should always tell the users when an outage is underway. Could we hide shorter outages, that would otherwise go unnoticed by the user, so as to give them a more positive impression of the overall interaction? While we briefly discussed some of these questions in section 7, future work could look at them in more detail. The answers to these questions could provide further design implications for collaborative AR applications.

9 CONCLUSION

This paper provided empirical evidence of the effects of outages on user experience and behavior, in a collaborative AR task. Our user study highlighted how users were affected differently based on their role in the interaction, and how providing users information about outages, using some feedback mechanism, may have alleviated some of these negative effects. We also provided anecdotal evidence of strategies that users adopted, in order to cope with outages. Using these insights from our study, we proposed some design implications for future remote collaborative AR applications.

10 ACKNOWLEDGEMENTS

This research was partially sponsored by NSF (Award: 1618321) and the Combat Capabilities Development Command Soldier Center (under Cooperative Agreement Number W911QY-15-2-0001). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of NSF, the Combat Capabilities Development Command Soldier Center, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] [n. d.]. <http://arsharing.cs.washington.edu/index.html>
- [2] [n. d.]. <https://objecttheory.com/platform/>
- [3] 2020. 10 BEST Augmented Reality Glasses (Smart Glasses) In 2021. <https://www.softwaretestinghelp.com/best-augmented-reality-glasses/>
- [4] Giuseppe Aceto, Alessio Botta, Pietro Marchetta, Valerio Persico, and Antonio Pescapé. 2018. A comprehensive survey on internet outages. *Journal of Network and Computer Applications* 113 (2018), 36–63.
- [5] Sudhir Aggarwal, Hemant Banavar, Amit Khandelwal, Sarit Mukherjee, and Sampath Rangarajan. 2004. Accuracy in dead-reckoning based distributed multi-player games. In *Proceedings of 3rd ACM SIGCOMM workshop on Network and system support for games*. 161–165.
- [6] Tooba Ahsen, Fahad R Dogar, and Aaron L Gardony. 2019. Exploring the Impact of Network Impairments on Remote Collaborative Augmented Reality Applications. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, LBW2619.
- [7] Deepak Akkil and Poika Isokoski. 2016. Accuracy of interpreting pointing gestures in egocentric view. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 262–273.
- [8] Deepak Akkil and Poika Isokoski. 2016. Gaze Augmentation in Egocentric Video Improves Awareness of Intention. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1573–1584.
- [9] Deepak Akkil, Jobin Mathew James, Poika Isokoski, and Jari Kangas. 2016. GazeTorch: Enabling Gaze Awareness in Collaborative Physical Tasks. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1151–1158.
- [10] Judith Amores, Xavier Benavides, and Pattie Maes. 2015. Showme: A remote collaboration system that supports immersive gestural communication. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. 1343–1348.
- [11] Huidong Bai, Prasanth Sasikumar, Jing Yang, and Mark Billinghurst. 2020. A user study on mixed reality remote collaboration with eye gaze and hand gesture sharing. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.

- [12] Yahn W Bernier. 2001. Latency compensating methods in client/server in-game protocol design and optimization. In *Game Developers Conference*, Vol. 98033.
- [13] Henrik Brun, Robin Anton Birkeland Bugge, LKR Suther, Sigurd Birkeland, Rahul Kumar, Egidijus Pelanis, and Ole Jacob Elle. 2019. Mixed reality holograms for heart surgery planning: first user experience in congenital heart disease. *European Heart Journal-Cardiovascular Imaging* 20, 8 (2019), 883–888.
- [14] Wolfgang Büschel. [n. d.]. Challenges in Collaborative Immersive Visualization. ([n. d.]).
- [15] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, trends, and outliers: How immersive technologies can facilitate the collaborative analysis of multidimensional data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [16] Tolga K Capin and Igor S Pandzic. 1997. A dead-reckoning algorithm for virtual human figures. In *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*. IEEE, 161–169.
- [17] Armando L Caro, Janardhan R Iyengar, Paul D Amer, Sourabh Ladha, Gerard J Heinz, and Keyur C Shah. 2003. SCTP: a proposed standard for robust internet data transport. *Computer* 36, 11 (2003), 56–63.
- [18] Kuan-Ta Chen, Chun-Ying Huang, Polly Huang, and Chin-Laung Lei. 2006. Quantifying Skype user satisfaction. *ACM SIGCOMM Computer Communication Review* 36, 4 (2006), 399–410.
- [19] Marshini Chetty, Srikanth Sundaresan, Sachit Muckaden, Nick Feamster, and Enrico Calandro. 2013. Measuring broadband performance in South Africa. In *Proceedings of the 4th Annual Symposium on Computing for Development*. ACM, 1.
- [20] Herbert H. Clark. 1996. *Common ground*. Cambridge University Press, 92–122. <https://doi.org/10.1017/CBO9780511620539.005>
- [21] Herbert H Clark and Susan E Brennan. 1991. Grounding in communication. (1991).
- [22] Catherine Durnell Cramton. 2001. The mutual knowledge problem and its consequences for dispersed collaboration. *Organization science* 12, 3 (2001), 346–371.
- [23] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. 2011. Analysis of country-wide internet outages caused by censorship. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 1–18.
- [24] Barrett Ens, Sarah Goodwin, Arnaud Prouzeau, Fraser Anderson, Florence Y Wang, Samuel Gratzl, Zac Lucarelli, Brendan Moyle, Jim Smiley, and Tim Dwyer. 2020. Uplift: A Tangible and Immersive Tabletop System for Casual Collaborative Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* (2020).
- [25] Barrett Ens, Joel Lanir, Anthony Tang, Scott Bateman, Gun Lee, Thammathip Piumsomboon, and Mark Billinghurst. 2019. Revisiting collaboration through mixed reality: The evolution of groupware. *International Journal of Human-Computer Studies* 131 (2019), 81–98.
- [26] Brian Eriksson, Ramakrishnan Durairajan, and Paul Barford. 2013. Riskroute: a framework for mitigating network outage threats. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*. ACM, 405–416.
- [27] P Esteban, JE Jaramillo, N Álvarez, J Restrepo, and H Treftz. 2004. Augmented Reality: a space for the understanding of multi-variate Calculus. In *Proceedings of the IADATE conference. Spain*.
- [28] Steffen Gauglitz, Benjamin Nuernberger, Matthew Turk, and Tobias Höllerer. 2014. In touch with the remote world: Remote collaboration with augmented reality drawings and virtual navigation. In *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology*. 197–205.
- [29] Steffen Gauglitz, Benjamin Nuernberger, Matthew Turk, and Tobias Höllerer. 2014. World-stabilized annotations and virtual scene navigation for remote collaboration. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 449–459.
- [30] Aaron M Genest, Carl Gutwin, Anthony Tang, Michael Kalyn, and Zenja Ivkovic. 2013. KinectArms: a toolkit for capturing and displaying arm embodiments in distributed tabletop groupware. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 157–166.
- [31] Darren Gergle, Robert E Kraut, and Susan R Fussell. 2004. Action as language in a shared visual space. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. ACM, 487–496.
- [32] Darren Gergle, Robert E Kraut, and Susan R Fussell. 2004. Language efficiency and visual technology: Minimizing collaborative effort with visual information. *Journal of language and social psychology* 23, 4 (2004), 491–517.
- [33] Darren Gergle, Robert E Kraut, and Susan R Fussell. 2013. Using visual information for grounding and awareness in collaborative tasks. *Human-Computer Interaction* 28, 1 (2013), 1–39.
- [34] Raphaël Grasset, Philip Lamb, and Mark Billinghurst. 2005. Evaluation of mixed-space collaboration. In *Proceedings of the 4th IEEE/ACM International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, 90–99.
- [35] Sebastian Günther, Sven Kratz, Daniel Avrahami, and Max Mühlhäuser. 2018. Exploring Audio, Visual, and Tactile Cues for Synchronous Remote Assistance. In *Proceedings of the 11th Pervasive Technologies Related to Assistive Environments Conference*. 339–344.

- [36] Felix G Hamza-Lup, Jannick P Rolland, and Charles Hughes. 2018. A distributed augmented reality system for medical training and simulation. *arXiv preprint arXiv:1811.12815* (2018).
- [37] Osama Haq and Fahad R Dogar. 2015. Leveraging the power of cloud for reliable wide area communication. In *Proc. ACM HotNeis*.
- [38] Tomoki Itamiya, Toshinori Iwai, and Tsuyoshi Kaneko. 2018. The Holographic Human for surgical navigation using Microsoft HoloLens. *EPiC Series in Engineering* 1 (2018), 26–30.
- [39] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. 2018. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9000–9008.
- [40] Junchen Jiang, Rajdeep Das, Ganesh Ananthanarayanan, Philip A Chou, Venkata Padmanabhan, Vyas Sekar, Esbjorn Dominique, Marcin Goliszewski, Dalibor Kukoleca, Renat Vafin, et al. 2016. Via: Improving internet telephony call quality using predictive relay selection. In *Proc. SIGCOMM'16*. ACM.
- [41] Wenyu Jiang and Henning Schulzrinne. 2003. Assessment of voip service availability in the current internet. In *Proceedings of the 4th International Workshop on Passive and Active Network Measurement (PAM 2003)*.
- [42] Laura Casado Joey Hadden. 2020. 21 major companies that have announced employees can work remotely long-term. <https://www.businessinsider.com/companies-asking-employees-to-work-from-home-due-to-coronavirus-2020>
- [43] Steven Johnson, Madeleine Gibson, and Bilge Mutlu. 2015. Handheld or handsfree?: Remote collaboration via lightweight head-mounted displays and handheld devices. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 1825–1836.
- [44] Ethan Katz-Bassett, Colin Scott, David R Choffnes, Ítalo Cunha, Vytas Valancius, Nick Feamster, Harsha V Madhyastha, Thomas Anderson, and Arvind Krishnamurthy. 2012. LIFEguard: Practical repair of persistent route failures. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. ACM, 395–406.
- [45] Hannes Kaufmann. 2002. Construct3D: an augmented reality application for mathematics and geometry education. In *Proceedings of the tenth ACM international conference on Multimedia*. 656–657.
- [46] Seungwon Kim, Allison Jing, Hanhoon Park, Gun A Lee, Weidong Huang, and Mark Billinghurst. 2020. Hand-in-Air (HiA) and Hand-on-Target (HoT) Style Gesture Cues for Mixed Reality Collaboration. *IEEE Access* 8 (2020), 224145–224161.
- [47] Kiyoshi Kiyokawa, Mark Billinghurst, Sohan E Hayes, Anoop Gupta, Yuki Sannohe, and Hirokazu Kato. 2002. Communication behaviors of co-located users in collaborative AR interfaces. In *Proceedings of the 1st International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, 139.
- [48] Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasie, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan Iyengar, et al. 2017. The quic transport protocol: Design and internet-scale deployment. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. 183–196.
- [49] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. Chapter 11 - Analyzing qualitative data. In *Research Methods in Human Computer Interaction (Second Edition)* (second edition ed.), Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser (Eds.). Morgan Kaufmann, Boston, 299 – 327. <https://doi.org/10.1016/B978-0-12-805390-4.00011-X>
- [50] Jun Li and Scott Brooks. 2011. I-seismograph: Observing and measuring Internet earthquakes. In *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2624–2632.
- [51] Stephan Lukosch, Heide Lukosch, Dragos Datcu, and Marina Cidota. 2015. On the spot information in augmented reality for teams in the security domain. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 983–988.
- [52] Tahir Mahmood, Willis Fulmer, Neelesh Mungoli, Jian Huang, and Aidong Lu. 2019. Improving information sharing and collaborative analysis for remote geospatial visualization using mixed reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 236–247.
- [53] Shawn McKee, Mike O'Connor, Soichi Hayashi, Marian Babik, Brian Tierney, Ilija Vukotic, Les Cottrell, and Henryk Gienza. 2016. *ICFA SCIC Network Monitoring Report*. Technical Report.
- [54] Microsoft. [n. d.]. Microsoft HoloLens. <https://www.microsoft.com/en-us/hololens> [Online; accessed 21-September-2018].
- [55] Microsoft. [n. d.]. Mixed Reality Toolkit. <https://github.com/Microsoft/MixedRealityToolkit-Unity> [Online; accessed 21-September-2018].
- [56] Jens Müller, Roman Rädle, and Harald Reiterer. 2016. Virtual Objects as Spatial Cues in Collaborative Mixed Reality Environments: How They Shape Communication Behavior and User Task Load. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1245–1249.
- [57] Jens Müller, Roman Rädle, and Harald Reiterer. 2017. Remote collaboration with mixed reality displays: how shared virtual landmarks facilitate spatial referencing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 6481–6486.

- [58] Nahal Norouzi*, Austin Erickson*, Kangsoo Kim, Ryan Schubert, Joseph LaViola, Gerd Bruder, and Greg Welch. 2019. Effects of shared gaze parameters on visual target identification task performance in augmented reality. In *Symposium on Spatial User Interaction*. 1–11.
- [59] Lyn Pemberton and Marcus Winter. 2009. Collaborative augmented reality in schools. (2009).
- [60] Thammathip Piumsomboon, Arindam Dey, Barrett Ens, Gun Lee, and Mark Billinghurst. 2019. The effects of sharing awareness cues in collaborative mixed reality. *Frontiers in Robotics and AI* 6 (2019), 5.
- [61] Thammathip Piumsomboon, Gun A Lee, Jonathon D Hart, Barrett Ens, Robert W Lindeman, Bruce H Thomas, and Mark Billinghurst. 2018. Mini-me: An adaptive avatar for mixed reality remote collaboration. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.
- [62] Thammathip Piumsomboon, Gun A Lee, Andrew Irlitti, Barrett Ens, Bruce H Thomas, and Mark Billinghurst. 2019. On the shoulder of the giant: A multi-scale mixed reality collaboration with 360 video sharing and tangible interaction. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–17.
- [63] Thammathip Piumsomboon, Youngho Lee, Gun Lee, and Mark Billinghurst. 2017. CoVAR: a collaborative virtual and augmented reality system for remote collaboration. In *SIGGRAPH Asia 2017 Emerging Technologies*. 1–2.
- [64] Ronald Poelman, Oytun Akman, Stephan Lukosch, and Pieter Jonker. 2012. As if being there: mediated reality for crime scene investigation. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*. 1267–1276.
- [65] Lin Quan, John Heidemann, and Yuri Pradkin. 2013. Trinocular: Understanding internet reliability through adaptive probing. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 255–266.
- [66] Fraser JM Reid and Susan E Reed. 2007. Conversational grounding and visual access in collaborative design. *CoDesign* 3, 2 (2007), 111–122.
- [67] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2019. Secure multi-user content sharing for augmented reality applications. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 141–158.
- [68] Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson. 1978. A simplest systematics for the organization of turn taking for conversation. In *Studies in the organization of conversational interaction*. Elsevier, 7–55.
- [69] M.F. Schober. 1995. Speakers, addressees, and frames of reference: Whose effort is minimized in conversations about locations? *Discourse Processes* 20, 2 (1995), 219–247.
- [70] Ronell Sicat, Jiabao Li, JunYoung Choi, Maxime Cordeil, Won-Ki Jeong, Benjamin Bach, and Hanspeter Pfister. 2018. Dxr: A toolkit for building immersive data visualizations. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 715–725.
- [71] Rajinder S Sodhi, Brett R Jones, David Forsyth, Brian P Bailey, and Giuliano Maciocci. 2013. BeThere: 3D mobile collaboration with spatial input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 179–188.
- [72] Theophilus Teo, Louise Lawrence, Gun A Lee, Mark Billinghurst, and Matt Adcock. 2019. Mixed reality remote collaboration combining 360 video and 3d reconstruction. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–14.
- [73] Trint. [n. d.]. Trint. <https://trint.com/>
- [74] Unity. [n. d.]. Unity. <https://unity3d.com> [Online; accessed 21-September-2018].
- [75] Yang Xu, Chenguang Yu, Jingjiang Li, and Yong Liu. 2012. Video telephony for end-consumers: measurement study of Google+, iChat, and Skype. In *Proceedings of the 2012 Internet Measurement Conference*. 371–384.
- [76] Naomi Yamashita, Katsuhiko Kaji, Hideaki Kuzuoka, and Keiji Hirata. 2011. Improving visibility of remote gestures in distributed tabletop collaboration. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. 95–104.
- [77] Jeffrey M Zacks and Pascale Michelon. 2005. Transformations of visuospatial images. *Behavioral and Cognitive Neuroscience Reviews* 4, 2 (2005), 96–118.

Received October 2020 ; revised April 2021 ; accepted May 2021