

To Beam or Not to Beam: A Study of Remote Telepresence Attendance at an Academic Conference

Carman Neustaedter¹, Gina Venolia², Jason Procyk¹, and Daniel Hawkins¹

¹School of Interactive Arts & Technology, Simon Fraser University, Surrey, BC, Canada

²Microsoft Research, Redmond, WA, USA

carman@sfu.ca, gina.venolia@microsoft.com, jep4@sfu.ca, drhawkin@sfu.ca

ABSTRACT

Attending and presenting at conferences is a core part of academic and industrial research. Yet it can sometimes be challenging to attend conferences due to travel restrictions, time limitations, or accessibility challenges. To understand how people may be able to attend conferences remotely, we explored the use of commercial telepresence robots called Beams at Ubicomp/ISWC 2014. Seven people attended the conference remotely and used Beams to attend presentations, ask questions, and participate in break mingling activities. We collected data of their activities and conducted an online survey with the broader set of conference attendees. Results show that telepresence robots supported attendance at a basic level and even empowered those with accessibility challenges. However, issues related to identity, interaction, navigation, and privacy emerged. We present recommendations for future telepresence attendance at conferences focused on balancing these concerns with the benefits participants received.

Author Keywords

Telepresence; robots; academic conferences

ACM Classification Keywords

H.5.3 [Information interfaces and presentation]: Group and Organization Interfaces - *Computer Supported Cooperative Work*

INTRODUCTION

There now exists a whole host of ways to be present in a remote location through telepresence technologies. This includes a range of basic to advanced video conferencing setups, the use of virtual environments and avatars [2], and the more recent emergence of telepresence robots that can be remotely controlled [27,31,32,35]. Such technologies are now being used to support telecommuting, distributed teams [35], remote meeting attendance [23], and conference

attendance (e.g., [30]) where each activity has the goals of fostering interaction, awareness, and feelings of presence.

The 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (Ubicomp) and 2014 International Symposium on Wearable Computing (ISWC) were jointly held in Seattle, USA with approximately 800 attendees. As part of these conferences, we planned, organized, and studied telepresence attendance. Ubicomp is a multi-track conference focused on the study, design, and development of mobile, wearable, and sensor-rich technologies. ISWC is a multi-session symposium featuring cutting-edge research in wearable technologies. Together, Ubicomp and ISWC featured three parallel tracks of paper presentations, a poster and demo reception, and a suite of workshops and a doctoral consortium preceding the main conference program.




Figure 1: Beams at Ubicomp/ISWC 2014.

The goal of telepresence attendance at the conference was to allow remote attendees to see presentations and engage in networking and mingling during the conference's breaks. We used Suitable Technologies' BeamPro telepresence robots (Figure 1, hereafter called *Beams*) to support the remote attendance. Beams are telepresence robots that show a live video feed of a remote user and allow him or her to remotely drive the robot. Seven participants attended remotely using Beams. We studied the experiences of various stakeholders to inform future usage of telepresence at conferences, as well as future telepresence designs. We observed remote attendees' experiences throughout the conference and conducted semi-structured interviews with

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CSCW '16, February 27-March 02, 2016, San Francisco, CA, USA

© 2016 ACM. ISBN 978-1-4503-3592-8/16/02...\$15.00

DOI: <https://dx.doi.org/10.1145/2818048.2819922>

them. We surveyed the local attendees following the conference to understand their reaction to the telepresence attendance.

Our work has shown that both local and remote attendees saw value in offering remote attendance at Ubicomp/ISWC. Many reported on positives of the experience. Despite this praise, we found several challenges that should be mitigated in future telepresence attendance at conferences. These related to issues with presenting one's identity and identifying others in the conference venue, interacting with in-person attendees during breaks, navigating through crowded areas and maintaining spatial awareness, and privacy concerns resulting from being present in two different locations at once.

RELATED WORK

Telepresence Robots

Since Eric Paulos' work on "personal roving presence" in 1998 [25], a rich body of work has explored robotic telepresence devices. Such a device, sometimes called a *mobile robotic presence* or *mobile remote presence (MRP)* [31,32], combines aspects of video conferencing with remote-controlled mobility to give a single remote user the ability to interact autonomously with a group of collocated people [17]. Applications including office workplaces, various healthcare situations, aging in place, and school attendance have received formal study [17].

Most closely related to our work are studies of workplace telepresence robots. We caution, however, that this setting is not the same as the conference setting that we studied. Yet it provides a very useful and relevant comparison. In this research, Jouppi [15] found wide field of views, mobility, and audio 'whispers' during conversations to be important in workplace settings using MRPs. Follow-on work emphasized the importance of 360-degree views and the preservation of user head height and posture [16]. Kimura et al. [16] and Paepcke et al [23] illustrated the importance of visual feedback along with socially appropriate vocalization levels for the remote person. More recently, wide angle and panoramic views have been shown to improve peripheral awareness and task completion, though they require additional cognitive processing by users [13]. Studies have also shown the importance of robot mobility on task completion [26] as well as the height of the robot as it relates to persuasiveness [28].

Lee et al.'s [18] in-depth study showed how the mobility of telepresence robots afforded a sense of availability and presence awareness amongst co-workers. Telepresence robots helped create a sense of social connectedness between remote and local colleagues by enabling casual interactions, though new social norms needed to be created around negotiations of space and control of movement (e.g., is it okay to push a robot?). In addition, we see the importance of knowing how one looks in a remote robot [18]. Lewis et al. [19] provide a series of heuristics for evaluating MRP systems, including the minimization of

driving costs, allowing flexible use, and ensuring safety and courtesy.

Throughout our study findings, we point to past research in cases where our work validates previous telepresence robot studies or extends them.

Academic Conferences

Turning to academic conferences, we see a number of efforts to study the use of virtual environments for remote conference attendance. First, Jones [14] studied remote attendance at the Teaching in Community Colleges Online Conference (TCC 1999) and the V-Learn Track of the Avatars Conference (Avatars 1999), which both used the virtual environment Active Worlds. Results showed challenges with using stock avatars to represent people, time zone differences, as well as the use of a chat channel for all conference interactions [14].

More recently, Erickson et al.'s [5] 2011 paper describes the use of a version of Second Life for supporting avatar-only attendance at IBM's Academy of Technology conference by 502 employees. Employees used a self-created avatar and attended a series of single-track talks as well as plenary sessions. Social challenges included creating an avatar with a desired representation, recognizing the avatars of others, and moving into and out of conversations. Participants enjoyed the social mingling provided by the shared spaces and avatars. They felt that viewing talks via video streaming would have been more effective than watching with an avatar, but people especially enjoyed a poster session where they could walk around and view research posters.

Shirmohammadi et al. [29] explored mixed attendance at The Third International Workshop on Massively Multiuser Virtual Environments 2010. Here nine people attended via an avatar in a virtual environment while another 10-20 attended in person. Findings included positive experiences viewing presentations, yet there was a lack of social interaction in the virtual environment when compared to those attending in person [29].

McCarthy and boyd studied digital backchannel usage on Internet Relay Chat (IRC) at the ACM CSCW 2004 conference [20]. Findings showed most people used IRC to talk about the presenters and their presentation content. The largest amount of posts occurred during the opening plenary and panels, given that their content is often focused on generating discussions.

While telepresence robots have been used on a limited basis for remote participation in a number of conferences (e.g., Double telepresence robots were used at ACM CHI 2014 in Toronto, Canada), to our knowledge, there has not been a formal study of telepresence robots in academic conference settings. This paper presents a first step towards filling that gap by focusing on a major international conference with several hundred local attendees.

Social Theories

Goffman [10] defines identity from a sociological perspective as the mental model one has of him/herself. As our results will show, this mental model and self representation is important when one attends a conference using a telepresence robot. Goffman explains that one's identity is made up various characteristics including one's appearance, attitudes towards others, beliefs, emotions, etc. and this is presented to others through their interactions and actions (e.g., body language, speech) [10]. In this way, individuals are presenting themselves like an actor on a stage to those around them [10].

Our research is also in many ways a breaching experiment in-the-wild where we try to understand and articulate a set of social practices and norms that emerge through the introduction of a new technology that 'disrupts' the normal social order [3,6,21]. In our case, the Beams purposely disrupt the normal social order of the conference setting in an effort to increase remote attendance. Throughout our results, we showcase how this 'disruption' affects the behavior of both local and remote conference attendees.

Lastly, our results explore situations that are articulated by theories of privacy in video-mediated environments. Here we see that privacy is often a balancing act between providing awareness and opportunities for interaction with others [12], and having a sense of autonomy, maintaining desired levels of confidentiality, and having some solitude, if desired [1,4,12,23]. Our study explores how privacy is balanced in the mixed-context of remote conference attendance and what issues emerge.

STUDY METHODOLOGY

The goal of our study was to understand if and how telepresence robots would support remote conference attendance, what the benefits were, and what social and technical limitations existed. As the conference's telepresence chairs and support staff, we planned and organized the telepresence setup, and then subsequently planned and executed the study of it. The study was approved by Simon Fraser University's Office of Research Ethics.

Remote Attendees

We placed a general call for people interested in attending Ubicomp/ISWC remotely on the conference web site and also various SIGCHI mailing lists. We asked interested people to email us responses to a series of questions describing their job position (e.g., faculty, industrial researcher, graduate student), the number of times they had attended Ubicomp or ISWC, and their reasons for wanting to attend remotely. Based on these responses, we selected a diverse cross section of participants including four faculty, two graduate students, and one industry researcher with varying reasons for wanting to attend remotely.

Three participants faced accessibility challenges that either caused them to have very limited mobility (restricted to

being mostly at home) or reduced mobility that made conference attendance in person very tiring and/or difficult. One of the remaining participants was a paper presenter at Ubicomp/ISWC but unable to get a travel permit in time to present in person at the conference. Her presentation was done using a pre-recorded video at the conference and she used her Beam to address questions at the end of the talk. Two participants were partners and had recently had a baby; they shared a Beam throughout the conference.

Three participants were in the same time zone as the conference, three had a 9-hour time difference, and three had a 3-hour time difference. Four participants had attended Ubicomp three or more times and the remaining three were attending Ubicomp/ISWC for the first time.

All remote attendees paid to attend the conference at half the normal rate. Because the technology setup was still experimental, we did not want them to have to pay the full registration fee in case there were technical issues and the experience was less than satisfactory, yet we also wanted to ensure they were committed to attending. We felt a reduced registration fee achieved these goals.

We purposely selected only a small number of remote attendees. This was limited by the number of Beams at our disposal, plus it also reflected our desire to study the experiences of each remote attendee relatively closely, indicative of qualitative research more generally. In the remainder of the paper, participants are identified by P followed by their participant number. Some quotes are left non-attributed to preserve anonymity.

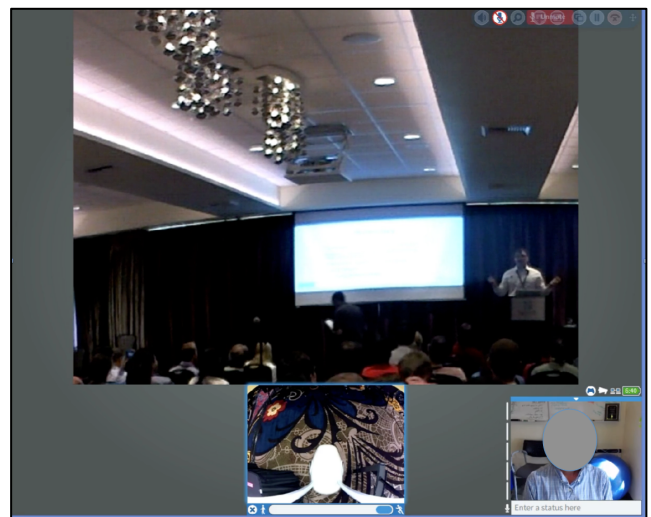


Figure 2: The remote attendee's view of the conference.

Telepresence Setup

Beams are telepresence robots, as shown in Figure 1. They are approximately five feet tall and contain a display at the top of the robot and two cameras. One faces outwards to provide the remote attendee with a first-person view and one faces down towards the robot's base to aid in navigation. The resolution of the webcam feed is 480p and field of view is 105°; the cameras were capable of 3x digital

zoom. Beams also contain an embedded microphone and speakers. The base of the robot contains wheels that allow it to move forward, backward, and turn.

Remote users can connect to an assigned Beam robot using a web interface, shown in Figure 2. The view from the first-person camera is shown at the top of the window. The top right of the interface contains controls for muting the microphone, adjusting volume, etc. The view from the floor-facing camera is shown at the bottom of the web interface along with a slider that adjusts the speed at which the Beam moves from slow to fast. Remote users can drive the Beam using a keyboard, mouse, or Xbox controller. The remote user can see how he or she appears on the robot's display with the mirror feature in the bottom right corner of the web interface (in Figure 2, the person's face is blocked to preserve anonymity).

Beams connect to a WiFi network to transmit and receive video. To try and alleviate network congestion issues during the conference, we had the conference hotel's tech support staff set up a dedicated wireless network for the Beams that was hidden from regular conference attendees. Occasionally Beams still lost connectivity or became unresponsive during the conference and in these cases we would manually reboot the robot.

Methods

Training and Orientation

Remote attendees first completed an informed consent form that told them about the study. During the week preceding the conference, they participated in a training session with Suitable Technologies staff to familiarize themselves with the operation of the Beams. We emailed them the hotel's map, which showed the conference room layout. The conference spanned the third and fourth floor of the hotel with stairs and elevators connecting the two floors. There was a large foyer space that was used for providing drinks and snacks during the breaks between sessions. There was no lunch provided, so in-person attendees went offsite.

Remote participants were instructed that they could attend all portions of the main conference (paper sessions) and would need to stay on the two floors of the meeting room area where the conference was being held. We did not allow participants to attend the conference reception because it was offsite and would be too difficult to transport the Beams; network connectivity was also poor at this location. Participants were also told they could not attend the Demo Reception on the Monday evening in the main meeting room area of the hotel because we felt it may be overly crowded for navigation. Despite this request, one person did attend the Demo Reception, albeit inadvertently (described later in our results).

Prior to the start of the conference, we used blue tape to section off areas of each presentation room floor where the Beams could park to watch the talk. This was always in the front left corner of the room (Figure 1, right). Our goal was

to have a space where the Beams could see the presenter/slides, while not blocking the local attendees' view of the talks. As discussed in our results, this layout changed over time due to poor viewing angles and the need for increased autonomy.

Personalization

Prior to the conference, we asked all remote attendees if and how they wanted to personalize their Beam. Our goal was to allow participants to choose their self-representation as best they could with the technology; in Goffman's terms, we were allowing them to create their 'actor' for 'performance' at the conference [10]. One person chose to have a scarf placed around the Beam's screen and another asked to have a printout of her school's logo affixed to her Beam. The rest opted to not decorate their Beams. In addition, we affixed each person's conference badge to his or her Beam under the main display, akin to how a regular attendee would wear the badge. During the conference we affixed stickers to the back of the Beam monitor with the remote participant's name and affiliation to facilitate identification from all sides.

Observations and Interviews

We observed remote and local attendees throughout the conference to see how they interacted with one another and to learn about people's behaviors. Observations involved us watching the remote attendees during all of the conference breaks and checking in on them periodically while they were in conference sessions. We did not stay with them throughout the sessions because they mostly sat parked watching the talks. We also observed them during several question periods at the end of talks, but not all.

We did not gain informed consent from the general conference attendees as part of our observations since the conference was in a public setting (this was approved by the ethics board). However, all attendees were told about the telepresence setup and the fact that we were studying it during the conference's opening plenary.

On the final day of the conference, we conducted one-on-one semi-structured interviews with each remote attendee using the Beams [11]. Questions focused on learning about what they did during the conference and understanding what they felt worked well and what did not work well about attending remotely. We also asked them for suggestions on future telepresence conference setups.

Throughout the conference, we had four people managing Ubicomp/ISWC's telepresence setup and this was easily doable with each person only spending a total of an hour or so per day troubleshooting problems. The remaining time was spent collecting study data. Fewer people may have been manageable, but it is difficult to know ahead of time if and how many technical problems will emerge and when they will occur.

Post-Conference Survey

After the conference had concluded, as part of the general post-conference survey, we asked all attendees questions about the telepresence setup. Most were open-ended. Questions focused on understanding how local attendees interacted with remote attendees, benefits to remote telepresence, and any challenges or limitations they faced as a result of having remote attendees at the conference. A link to the survey was sent to attendees via email. We received 210 responses (74% male). 50% of respondents were graduate students, 21% academic faculty, 13% industrial researchers, 4% post-doctoral researchers, 3% practitioners, and 2% each undergraduate students, programmers and designers. These numbers were representative of the general conference attendance with a large number of male attendees (~73%) and students (~40%). 17% of respondents had seen a Beam in usage before the conference and 2% had personal experience using a Beam. In the remainder of the paper, survey respondents are identified by S and a number.

Data Collection and Analysis

During our observations, we kept handwritten and typed notes and recorded pictures and video clips periodically to aid our analysis and memory of events. For example, we captured media about the locations where the Beams were parked during talks, how the remote users navigated through hallways, and how local and remote users interacted between talks and during breaks. Our observational data helped us understand our interview and survey data, and also acted as a lens throughout the analysis process. We did not explicitly apply any data coding procedures to our images and videos.

Interviews with each remote attendee were audio recorded and fully transcribed. We used open, axial, and selective coding [2] to draw out the main themes from the interview data. Here we iteratively went through each interview transcription and coded participants' quotes in an open fashion, generating codes as we went. We grouped codes/quotes through an axial coding process into themes related to navigation, identity, interaction, privacy, accessibility, and presence (along with various subthemes). Through selective coding, we selected the main results from our themes where we chose what appeared to be the most salient themes; these are presented in the paper. Coding of our interviews was performed by three researchers and each reviewed the codes and themes found by the others. Our analysis approach was based on analysis procedures found in Grounded Theory [2], but we purposely did not explicitly use Grounded Theory as a broader methodological approach given that we already had prior knowledge of how telepresence robots were likely to be used (based on past research in workplace settings).

Survey data was analyzed using thematic analysis (by a single researcher). Respondent answers were iteratively read and grouped according to common answers and

themes. We compared the findings from our survey with our interview results.

Without prompting, remote attendees created a backchannel chat using Google Hangouts on the first day of the conference and used it throughout the conference to share information among themselves and with our team. We as conference organizers also joined the backchannel and it became a valuable tool for providing technical support for the remote attendees. In comparison to McCarthy and boyd's study of an IRC backchannel at CSCW 2004 [20], our participants mostly talked about the Beams and technical or social tips rather than the presentations themselves. We recommend that conference organizers who are supporting remote attendance set up and participate in a backchannel for communication among the remote attendees. We analyzed the text within the backchannel's log using thematic analysis. Again, we compared the results with our interview and survey data.

In the following sections, we detail our study results. We begin by articulating the general reactions that we received from both remote and local attendees to the telepresence setup. Next, we step through each of the main social and technical themes that our data pointed to.

GENERAL REACTIONS

Overall, remote attendees reported enjoying the remote telepresence experience. Their time in attendance ranged from being at two to three talk sessions per day all the way to nearly the entire main conference (including plenary talks). Those with large time zone differences found it difficult to be at the conference into the evening and would sometimes finish days early as a result. The two participants who shared a Beam found it was easy to do so because they lived together. Here they could easily micro-coordinate which talks they wanted to attend.

Attendees valued being able to attend the conference without having to travel. Those with accessibility challenges felt especially empowered because they were able to attend the conference in a way that was not previously possible for them. It also alleviated mobility issues that they normally faced when attending in person.

The best part was that I was not able to come there so if I was not able and I missed the conference then it was really a pity. For me it was a great opportunity to at least be present partly in the conference and then don't miss it. - P5

I'm attending some conferences via Skype, but I never really felt as if I was there. I would just hear. Sometimes I couldn't see the presentations....So this is the first time I feel I'm a part of the conference....Attending through something like a web conferencing software, it isn't- you don't get the feeling you're there....Because you feel like you're just hearing maybe something on the radio, or something on TV. But using telepresence like this, I have the feeling I'm there. I'm able to meet new people. I'm able to see familiar faces. - P4

Actually it was a peculiar situation for us with a newborn baby and so for us was a bit of a way to participate as a family, so it was also that. I mean, we cannot be physically, but we can be there, all the three of us in a sense. It was nice, actually. It was fun. – Non-Attributed

Compared to using more standard video conferencing to remotely attend an event, participants generally felt that their experience with telepresence robots was more immersive. Some also felt they were advantaged because they could more easily fit the conference into their normal activities, such as multi-tasking work.

This is much better than video conferencing. I can tell when people are directing questions to me. There's certainly the ability to move around and direct my conversations toward people in a crowd, and the fact that I can go to different places. I do feel more immersed...I was able to guilt-free multitasking at a conference. If I go through the trouble of attending a conference I don't feel good skipping out on sessions doing other work, but I didn't feel bad about it.- P3

Local participants reported overall positive experiences with the telepresence experience. 89% of survey respondents recommended doing telepresence again at future Ubicomp/ISWC conferences. Many local attendees commented in the survey with similar reactions about how the robot telepresence was providing important ways for some people to attend remotely.

It was great. It was a key talking point of the conference, it allowed people with disabilities or family or visa issues to attend, and even for a remote presenter to take questions – S93

In addition to these generally positive experiences, we found a series of social and technical challenges emerge. Throughout the remaining results, we describe each of these in detail along with recommendations for conference organizers and telepresence designers. We have explicitly separated out the implications for each stakeholder as the implications are often different. It was also evident from our study that conference organizers can have an impact on the remote attendees' experience beyond the intent of the designers of the telepresence robots. Thus, it is critical to understand the implications for conference organizers. Moreover, we think that many of our results are valuable for those who might attend a conference remotely in the future for they shed light on best practices for attending remotely, presenting one's identity, interacting, and recognizing potential privacy issues.

IDENTITY AND THE PRESENTATION OF SELF

We found several challenges related to identity from the perspectives of the remote attendees and also those attending Ubicomp/ISWC in person. Here we explore aspects related to Goffman's presentation of self [10]. First, remote attendees faced challenges in being able to present their identity remotely in their Beam robot. Unlike real life and the way people can easily customize their appearance

[10], each robot was visually the same. Even though we had asked if people wanted to personalize their robot in some way, as mentioned, only two did so. Remote attendees also found that they had no way of seeing themselves to understand how they looked to people attending Ubicomp/ISWC in person [32]. Unlike situations where a person might look in a mirror to see herself, the remote attendees did not easily have such an option (and none chose to drive into a washroom to use a mirror). Knowing how one looked to local attendees was important to them as it could influence their conversations or the reactions from others.

I don't have a real good sense of myself in the scene so I can't see the scarf at all. I've had to ask people how long the base is, or how far it is from the base to their feet in order to get a better sense to not hit people's toes. I didn't know for a while where I had a badge, or if. I had to ask people a lot of questions about what I looked like to even get a sense of that so I don't really feel like I'm wearing a scarf. – P1

I wore power jewelry because I wanted to be certain that I didn't appear too dominant. The first day I wore red jewelry. There are ways that women can use jewelry and clothing to look feminine versus incompetent. I wanted to make sure that I looked assertive and powerful because I feel I always have to counteract the stigmas associated with being single and female. Now you add this tele-robotic dimension and look confident. - P2

Local attendees also liked to see the Beam robots personalized in some way (cf. [35]) so they could easily identify who was who, yet sometimes this was difficult to know because of the small size of the nametag or a lack of personalization of the Beams.

I loved that one person personalized their Beam with a scarf. It would have been nice if all of them had identifying traits like that - the name tags were hard to read from any distance. – S16

It was also difficult to identify the users of the non-augmented (i.e., bescarfed, name-tagged) Beams, especially from behind. – S138

Remote attendees talked about other ways that they thought they might have also liked to present their identity remotely. This moved beyond what one might commonly do when in person at a conference and included ideas such as attaching printouts describing one's research and RFID tags to identify passersby.

In addition to one's visual presentation of self, how one presented herself acoustically was also important; again, this is an important part of self-representation and identity [10]. Remote attendees found it was difficult to understand how loud they were when talking to others and if it was socially appropriate given the context of the situation they were in, as also found by Paepcke et al. [23] and Lee and

Takayama [18]. The remote attendees wanted to engage in a socially appropriate manner, but it was not always easy to achieve.

Like the visuals of their remote representation, they could not hear how they sounded in the remote environment. Local attendees sometimes complained that the remote attendees were talking too loudly or too softly. In one instance, a remote attendee was talking to an individual in the corner of a presentation room between sessions. The remote attendee thought the conversation was private with the person, but in actual fact it could be heard throughout the entire presentation room by others mingling about.

They don't adjust the volume of their voice according to surroundings. In large group conversations, the volume of the Beam was higher than surrounding people, so it was as if a person was standing in the room and shouting. – S7

In such circumstances, the remote attendees lost the ability to control aspects of their presentation of self [10].

Second, we learned that some people can develop a strong tie with their remote embodiment, almost as though it was their own 'body' that only they should be able to use.

You might hurt my bot, knock my bot over. It is physically disconcerting at this point. [It's like you] become the Mars rover, I really did feel like, immersed in this experience enough such that when somebody walks into the bot and it shakes, I feel my body tense. – P2

In somewhat of a contrast, as organizers of the conference, we felt we were providing a series of robots that were somewhat interchangeable between attendees. In this sense, we felt that if one stopped working for whatever reason, and another attendee was away or not planning to be present for a portion of the day, we could use that person's robot for someone else. Thus, as organizers, we had a relatively loose sense of tie between the robots and a particular person. Most remote attendees had the opposite reaction where they felt strong ownership over 'their' robot. Moving between robots, or realizing they had been connected into a different robot because of system configurations on our end, was disconcerting to them.

Similarly, it was very disconcerting when I woke up in the wrong body. I was really expecting to wake up in the middle of the room. At first I thought you'd taken me home. I'm like, "Why did <tech support> take me home? He knew I was talking to an NPR reporter. – P2

oh i didn't realize you guys name swapped too. oops. man it was weird being 'in the wrong body' and coming to life in the wrong place! is the town hall still happening? or did i miss it. – P2, Backchannel

Recommendations

We recommend that conference organizers and telepresence designers should explore ways to convey the identity of the remote attendees from all angles. Conference organizers

should encourage remote attendees to personalize their representation in the conference space and remote attendees should in turn think of ways to do so. Remote attendees should be given a clear idea of how they appear to the on-site attendees, whether through mirrors or infrastructure cameras (although Takayama and Harris [32] found that these interventions may have negative consequences). We found that remote participants bonded with the particular hardware assigned to them; conference organizers should endeavor to support that identification.

Beams require the remote person to set the speaker volume, yet do not provide adequate feedback for this task. Like Tsui et al. [34], we recommend that designers of telepresence systems either give the missing feedback [16], empower local attendees to change the volume (though this may be socially problematic [18]), automatically set the speaker volume based on the ambient noise level [7], or add sidetone to the audio channel [23]. It's unclear which solution (or combination) would address the problem.

These ideas align with social theories of the presentation of self to convey how people desire to present themselves in a telepresence robot (as an 'actor') and what is valuable for others with whom they interact (their 'audience') [10].

NAVIGATION AND A LACK OF SPATIAL AWARENESS

Navigation was sometimes challenging for remote attendees. This was especially the case in crowded rooms and hallways, such as when breaks were occurring. Unexpected obstacles such as power cords on the floors were difficult to avoid. Remote attendees talked about lacking spatial awareness of the larger conference setup. We had provided them with overview floor plan maps, but this did not necessarily help that much since it did not show what the space looked like in more of a first-person view. The cameras on the Beams tended to lack a peripheral view, which further made it difficult.

Movement between presentation rooms was generally left to the breaks because remote attendees found it hard to unobtrusively move between talks and 'session hop'. The Beams made a low noise as they moved and attracted visual attention. It was also even more difficult to avoid obstacles when trying to be discreet. Thus, going between rooms during a talk was easily seen as a distraction and highly noticeable. This was also partly because they had to park at the front of the room to get a good view of the presenter and slides.

In addition, some remote attendees felt that they lacked a larger sense of where the conference was being held. They were restricted to just two floors of the hotel (the meeting room space) and could not see or go out into the city more broadly. As a result, one remote attendee asked us for help to navigate to an open window where she would be able to see the city of Seattle to know what it looked like.

Sometimes navigation issues created unique opportunities for social engagement and interactions with others. When

Beams got stuck on power cords, for example, attendees may strike up a conversation together. Remote participants also needed to ask attendees for help with the elevators.

Recommendations

Conference organizers should be sure that their event will not be overcrowded and that there are clear navigation paths for the remote attendees. They should provide overview maps of the venue and give the remote participants adequate time to explore their environment in advance of the conference.

Designers of telepresence robots should explore practical solutions for indicating the robot's current position on the venue map (suggested by Tsui et al. [34]), and perhaps include the locations of other robots at the event on it. Other research has shown the value of wide angle and panoramic camera views [13].

ACCESSIBILITY AND AUTONOMY

As stated, several of our remote attendees faced accessibility challenges in their daily lives due to reduced mobility. The remote attendance allowed them to stay in a location such as their home or work office where mobility challenges were much less likely to surface. They could also easily disconnect from the conference if they needed additional rest. Those with accessibility challenges generally felt a strong sense of empowerment and additional feelings of autonomy that they often did not normally feel when attending conferences. Instead of feeling frustrated and challenged with moving throughout the conference, which they sometimes felt in-person, their accessibility issues were somewhat removed and much less 'present'. One remote attendee even described the feeling as one of having 'superpowers.'

It feels magical, because it's like your teleporting, right? I am in <city> but I'm still in Seattle, some part of me, right?
- P4

P4 was accustomed with moving around in a wheelchair and she felt this helped her use the Beam.

It's very wonderful, I see Beam as a great starting foundation to develop further robotic technologies in the future. Even now, I feel that it's still more like a wheelchair and it has Skype features, right? It has to have additional features, like maybe an arm that connects, to extend to shake somebody's hand, more interaction instead of just this visual presence... - P4

On the other hand, the use of the Beams created new accessibility challenges for those who were remotely attending. The most prominent example of this issue resulted from the elevators that were necessary for the Beams to move between the conference venue's two floors. The remote attendees could not activate the elevator call buttons nor select their desired floor. As a result, they had to ask people standing near the elevators for help. Sometimes the local attendees did not know when or how to

help the remote attendees because they were not familiar with the situation. Others would engage with the remote attendees in somewhat awkward ways by pushing or pulling them in and out of elevators, even though they could do this movement on their own.

Recommendations

Any measures that the conference organizers can take to give the remote participants agency and autonomy should be strongly considered. For example, they could communicate to on-site attendees about the challenges of robot mobility and instruct them that helping out by keeping pathways clear and helping at doors and elevators is appropriate providing that autonomy is maintain (e.g., avoid pushing a robot when helping). Alternately, designers of telepresence robots should consider solutions to allow remote users to control doors and elevators.

SOCIAL NORMS

Throughout the conference there was an evolution of social norms regarding how one should act as a remote attendee and how local attendees should interact with the remote attendees. This created several socially awkward situations, yet it also reflected the fact that our attempt to include Beams was a form of breaching experiment, with the likely effect of disturbing existing social norms [3,6,21].

First, remote and local attendees faced social challenges around the positioning of the Beams within the presentation rooms. As mentioned, we had originally created an area in the front corner of each presentation room where we wanted Beams to 'park' during a talk (Figure 1, right). Yet this created issues because the remote attendees had a poor view of the presenter and projected slides because of the angle. Over time, the remote attendees ended up moving around and trying different locations from which they could watch the presentations. This included the back of the room, which was too far from the screen to see them well. Eventually, most of the remote attendees found the best viewing location to be at the front of the aisle separating the rows of seats. This allowed them to be as close to the screen as was possible while still aligning one's Beam with the front row of the seats in the room. They also found they needed to use their camera's zooming features to help see the slides, even when close to the front of the room.

Well, at the beginning I thought that I had to park on the corner of the road on the place that they were, I think it was designated for robot. Then I had a very bad spot I couldn't see much of the slides. Then I moved to the middle and it was still difficult to read the slides, but I found that if I zoomed then I would see every word on the slide. I missed a lot of slides because I didn't know that I could zoom and see better. - P5

When parked in their location during a talk, the remote attendees were able to easily ask each other through the backchannel if they were blocking the view of another remote attendee in the Beam. However, it was more

difficult for the remote attendees to ask the local attendees if they were blocking their view. They lacked the ability to easily turn around and whisper or ask them in a low voice. As a result, many local attendees complained about Beams blocking their view of the presentations when the remote attendees decided to park in the front row of the room.

The Beams are unable to sit, so during talks they end up holding a privileged social position (effectively doing the equivalent of standing at the front of the room during a talk, since they weren't able to position themselves elsewhere and be able to see the screen). If a person did that, it would be considered very rude. – S7

Second, novelty was clearly a factor in terms of how local attendees reacted and interacted with the remote attendees during the initial day of the conference. Beams were a focal point during the first day of presentations. Local attendees could be seen talking amongst themselves and pointing to the Beams that were 'parked' in their session. During breaks, a common activity was local attendees taking photos ('selfies') with the Beams. Many local attendees would also surreptitiously capture pictures of the Beams for their own records or to post on social media.

The first day I couldn't get five feet without a selfie. The first day, the first morning ... - P2

These behaviors diminished considerably after the first day of the conference as people became used to having the Beams around. Until that point, they were, however, somewhat of a 'celebrity' but also a distraction for those who were trying to focus on the conference talks. As with many breaching experiments, such social 'disruptions' may go away with time as social norms adjust [3,6].

A lot of the audience focused on what they were doing just because it was novel and interesting ... might have taken attention away from some of the talks. – S21

I saw some people taking pictures and other such novelty interactions, but I also saw them working into regular group interactions during the breaks. – S187

Recommendations

Conference organizers should consider options for giving the remote attendees adequate audio and video of the talks. One option would be better positioning for the telepresence devices. Another would be to stream the presenters and slides directly through dedicated A/V equipment (or digitally share the slides in advance of the talk).

Designers of telepresence hardware should provide human-like aural and visual acuity. Based on our own informal measures, Beam acuity is approximately equivalent to 20/200 vision (i.e., the remote user can see something at 20 feet that an average person can see at 200 feet) when zoomed out, and 20/100 when zoomed in fully. Given this, it is not surprising that the remote attendees took up residence at the front of the room. Additionally, telepresence robots should be adjustable between a normal

human sitting height and a normal standing height, ideally under control of the remote user.

Those who attend a conference via a telepresence robot in the future should consider the social norms of how they move and park their robot while recognizing their likely impact on others.

SOCIAL ENGAGEMENT & MINGLING

The remote attendees generally felt that they were able to engage socially with other conference attendees through acts of 'mingling' during the breaks and at lunch (at least with those who did not go offsite for lunch). Interactions with others are integral to the creation of one's self-representation and identity [10].

It was a good interaction. It was someone who I had only previous been in email correspondence with, so it was nice to put a face to a name. The beams were very approachable (more so than people) which I really liked. – S75

Despite the positives, remote and local attendees also faced challenges when interacting with others.

First, it was difficult for both remote and local attendees to know if the other wanted to engage in conversation when they would see each other or it would appear like one was approaching the other. The normal body language that one might use to discern interest was difficult to ascertain from the remote and local attendees alike. For example, continuous eye contact over several seconds while approaching an individual might indicate the desire to converse. Yet such eye contact was hard to ascertain between remote and local attendees. Similarly, one's body posture and direction when approaching another person can help people determine if one is moving towards them to talk, or if they are simply navigating a crowded space. Such cues were hard to discern for local attendees because it was sometimes difficult to drive the Beam in a crowded area. Remote attendees could not easily see body language because the camera's quality was not high enough.

It was hard to use them in the coffee breaks. When a beam approached me I couldn't tell if they person was interested in talking to me or if they were trying to get past me. (no visible cues to make up for the lack of body language). – S204

Along a similar line, it was sometimes difficult for remote and local attendees to know when a conversation between the two should end. In a typical face-to-face conversation, people are often very good at determining this based on one's body language and gestures. However, these cues were hard to understand when the Beam was being used.

What I found that when I was facing someone and speaking it's difficult to end a conversation because you have to move really if you want to leave you have to move. If you don't say see you around, or if you are talking I felt that the other party was not, really wanted to continue with, it was awkward to end a conversation. - P5

Second, remote attendees faced a difficult time in identifying local attendees. Despite having a wide field-of-view camera on the Beams, they still felt like they lacked peripheral awareness to see who was around. As mentioned, it was also hard to recognize people at a distance, which included challenges in seeing and reading people's nametags.

And I have no peripheral vision so it's hard to spot people. It's also hard to recognize people, and I don't have the liberty of looking at their badge. ... I can't read badges even when I zoom in which most of time I don't have the opportunity to do before the conversations moved past that anyway. - P1

Even though we know the people and some people, if you know their voice very well then you could immediately pick up who was talking, but in many cases it really took a while. I sit also maybe turning left and right on this thing who is speaking. By the time you figured that out, maybe someone else was speaking, so that made it quite tiring. - P6

As a result of the challenges in identifying others, remote attendees who were more senior in the field generally felt it was easier to network and meet with people that they already knew, or that knew them and would approach them. One senior remote attendee had even developed a strategy where he would let people come to him, rather than approaching others, given the challenge in identifying people. Remote attendees who were job hunting found it difficult to do so because they would need to approach and talk with people that they may not have previously known. On the other hand, remote attendees who were relatively new to the field found it was beneficial to be using a Beam because of their initial 'celebrity status' at the conference. This made them more 'well known' at the conference than they would have been if they had attended in person.

Third, we learned just how important food and beverages are to a conference experience, especially when it comes to mingling and interacting during breaks and lunch. For obvious reasons, remote attendees are unable to eat and drink at the conference. Yet this means there is no reason to stand in a food or drink line where one may easily begin to chat with the person next to her in line. It also means that when 'standing around' and conversing, there is not the additional act of eating or drinking to augment the conversation with an activity. Instead, the conversation is focused on just the conversation. Food and drinks also represent a common and socially acceptable way to leave a conversation. For example, a person can tell another that they need to go get another drink or more food.

[I miss out during] the coffee breaks. (laughs)... The lunch and dinners... If there was some other activity planned for those who connect remotely, that would be nice. - P4

One might imagine that a remote attendee could attempt to resolve this situation by trying to eat or drink a beverage at

the same time as the local attendees. This might make them feel like they 'fit in' more with what was occurring at the conference site. Yet often this type of approach did not work well because there was a time-zone mismatch between the conference and attendee's location.

Remote attendees could not go to offsite restaurants for meals with attendees. They were also not able to come to the conference reception, which was held at a local offsite music museum. This meant their socializing time was limited and, perhaps, somewhat overly focused at the more 'academic' aspects of the conference: the presentations and the short breaks between them. By mistake, one remote attendee happened to attend the conference's Demo Reception (at the hotel) because she did not realize she was supposed to park her Beam for the night at the end of the day's talks. Instead, she drove around and mingled with other attendees. Overall, she described this as one of her favorite times at the conference because she was normally not able to attend large social functions like it due to her mobility issues. Upon reflection, the Demo Reception was just as crowded as the conference breaks and so it would likely have turned out to be fine having the Beams in attendance at it.

Recommendations

We recommend that conference organizers communicate with on-site attendees in advance of the conference to establish expectations and social norms. The remote experience could also be enriched by allowing the remotes to participate more fully in breaks and off-site activities.

As noted above, telepresence hardware developers should consider ways of improving visual acuity, to benefit mingling activities. At the same time, mingling requires even greater field of view. It may be possible to address this tension by combining peripheral and foveal views [35]. Mingling also benefits from the telepresence device's display and camera being near standing human eye height, further motivating adjustable height.

How to support beginning and ending conversations [33] in mingling situations that include telepresence devices is an open research question and is ripe for future work. Potential solutions may include utilizing the robot's mobility as a form of body language for starting/stopping conversations (e.g., slight movements may indicate the end of a conversation).

PRIVACY: THE MIXING OF TWO SOCIAL CONTEXTS

Remote attendees found themselves in two social contexts at one time during the conference and often needed to mitigate the challenges that this setup created. For example, P2 discussed the telepresence setup with her live-in partner prior to the conference starting. She explained how he could tell if her microphone was on and what areas of their apartment were captured by the camera. She also placed small pieces of paper over top of her camera when it was not in use. Thus, there was clear foresight and planning in

an effort to regulate how she and her partner participated in the space (or did not), thereby exercising autonomy [1].

[My partner] knows. I tried to tell him that if the headset is on the mic might be open. Don't talk unless you know the mic is closed ... We sort of discussed where the camera [was and what it captured]. - P2

Along a similar line, P1 occasionally connected into the conference from home so that she could also be with her children portions of the time or they could see her husband who was at the conference in person. At times she placed a tea bag label over her camera so she could see what was happening at the conference and watch a presentation, while not letting local attendees see into the more private confines of her home. This act allowed her to have a private view in solitude [1] and preserve some confidentiality by not revealing more than she wanted to [1].

I'm in my office. Most of my attendance have been from [my office], a little bit of it from home...Because I can close the door, and make sure that I'm not disturbed....I attended from home because the kids wanted to see my husband. - P1

P2 and P7 talked about how they purposely dressed in a way that they felt was acceptable in a conference setting, despite connecting in from home.

I actually put on a regular girl's makeup. It looked terrible. Then I actually sat down on the camera and redid my makeup. I did it like I did theater. I used to do a lot of theater, so I did my makeup like theater makeup. Then I actually looked like a normal woman. - P2

I was also kind of okay dressed. I don't think I would have done this in pjs or a t-shirt. I think I would have, I wouldn't say dress up, but I would have tried to be clothed appropriately. - P7

On the other hand, some participants were less concerned about being captured in their home. For example, one participant attended the entire conference from her home because she was confined to laying or sitting down as a result of mobility challenges. She was used to attending university classes like this using Skype and so it was not a problem for her if she was seen at home in this manner.

Despite attempts to mitigate privacy challenges, sometimes issues still arose. In one situation, for example, a remote attendee forgot her microphone was on while talking to her children. This was heard by local attendees and breeched the confidentiality [1] of the remote attendee by revealing more than she had intended to others at the conference.

Well I accidentally left my microphone on when I had the kids there, and the talk had already started. I was trying to tell them to leave quietly thinking my mic was off. So I had a parenting-in-public, and disturbing-a-speaker moment all at once...Somebody came up and waved in the camera, and I turned it off as soon as I could. - P1

Turning to survey responses from the local attendees, we did not see any concerns about seeing into the remote attendees' homes or otherwise private locations (e.g., an office). Similarly, nobody commented on concerns over having themselves visible by other people at the remote attendees' locations (e.g., by work colleagues if at work, by family members if at home). This is likely because the local attendees were already comporting themselves in a manner that they felt was acceptable for public viewing.

Recommendations

Conference organizers should work with remote attendees to instruct them of the potential for privacy implications, and suggest ways to mitigate them. Organizers should consider ways to allow remote attendees to: regulate their autonomy [1] by choosing when and how they participate in the space, easily move into and out of solitude [1] from local attendees, and sometimes mask camera views in order to reveal an appropriate amount of information about oneself (regulate confidentiality [1]).

Designers of telepresence devices should consider carefully whether an open-mic design is best, compared to push-to-talk or other solutions. If open-mic is provided, the current state of the mute toggle must be very apparent. Researchers should also explore options for mitigating privacy issues with what is shown in the background of the camera frame in such mixed contexts [1,20].

DISCUSSION AND CONCLUSIONS

We now summarize and reflect on our findings to suggest key directions for future conference telepresence explorations and research. While many of our findings may relate to office-based settings, our focus is strictly on conference settings. Thus, we are not attempting to generalize beyond the setting we explored even though some findings may be relevant to other situations. Future studies would be needed in office-settings to understand how our work generalizes to them (if at all).

Successes and Challenges

Overall, we found that the Beam telepresence robots were an effective solution for including a small number of remote attendees at an academic conference and would certainly advocate for similar setups at future conferences. Those with accessibility challenges found the experience especially valuable as it provided them with new opportunities for remote attendance to overcome mobility issues and fatigue. Remote attendees who were unable to travel due to family, visa, or time issues similarly thought the experience was better than other telepresence options such as video conferencing for presentations because of the ability to navigate and socialize at the conference.

Yet the experience was by no means perfect and we found lots of areas for improvement, as mentioned throughout the recommendations in our results. Compared to conference attendance via virtual environments, we see that similar issues emerged when using telepresence robots. This

included challenges in constructing one's identity, recognizing the identity of others, and moving into and out of conversations [2,13]. These all relate to the manner in which one is able to understand his or her identity and self-representation and customize it as desired to create an ideal self [10]. In our case, the remote attendees were limited in the ways they could do this customization. This suggests future work in ways to better manage and create one's self-representation.

We also influenced self-representation through some of the decisions we made as conference organizers (e.g., partitioning areas off for the Beams to park which affected their interactions with others, changing their robots when technical issues arose). In future situations, such alterations should be minimized or discussed with remote attendees ahead of time to give them the opportunity to exercise their autonomy as it relates to identity and self-presentation.

Telepresence robots brought with them the additional complexity of a mixed context where remote attendees were both present at home/work along with being at the conference. This also meant that they 'stood out' as being different from other attendees, unlike an avatar-only conference. These findings relate to privacy challenges in video-mediated environments where one must take additional steps to exercise autonomy, confidentiality, and solitude [1].

Design Suggestions and Training

We feel that many of the problems we uncovered for robotic telepresence at conferences can be overcome with design alterations to the telepresence robots, additional training or education on how to best interact while using a telepresence robot, or training for local attendees on how to best engage with remote attendees along with an understanding of the limitations of the telepresence robots. However, issues around social awkwardness and establishing social norms may represent especially difficult problems. Lee and Takayama [18] studied long-term deployments of telepresence robots in workplace settings and found that social norms were a problem even in that context. Takayama and Go [31] attributed the problem to one of mixed metaphors. They recommend resolving the confusion by making the identity of the remote person even more readily apparent. It is not clear exactly how to put this into practice in a conference setting, or if it would be sufficient to establish social norms and mitigate the social breakdowns. As said, our research was a form of breaching experiment [3,6,21] where the remote attendees 'disrupted' the normal flow of conference attendance. Such disruptions are likely to ease over time, yet social norms will need to evolve to accommodate and adjust to robotic presence.

Suggestions for Future Conference Telepresence

Reflecting on Ubicomp/ISWC's telepresence setup, we feel that ripe areas for exploration may include expanding out the number of remote attendees or considering the effects of time-sharing telepresence robots for a particular period of

time (e.g., a desired presentation session). The benefit of having one robot per attendee certainly allows people to customize their robot and develop their desired level of attachment to or identification with it. However, this could be a costly endeavor if there are a large number of remote attendees. It would also increase the amount of technical support that would be needed to manage and oversee the telepresence setup. Time-shared telepresence robots could reduce infrastructure costs, yet may detract from one's construction of identity, unless the robots could be easily altered in appearances to match new attendees. It is also not clear what is the right balance of telepresence robots to in-person attendees. It is likely that this would largely depend on the amount of open foyer space that is used for conference breaks, as this was the toughest area for remote attendees to navigate. Other options may include representing remote participants in a streamed lecture [9], or the use of mobile or wearable telepresence devices [8].

Lastly, we hope that other conferences will replicate and improve this approach, but even more we hope that researchers attempt, study, and report other approaches to remote attendance. There are many other approaches that could be attempted, either individually or in combination.

ACKNOWLEDGEMENTS

We thank the Ubicomp/ISWC conference chairs, A.J. Brush, Adrian Friday, and Lucy Dunne, for their help and support throughout the planning and conducting of the conference telepresence, and John Lateulere for his technical support throughout the conference. We are very gracious to Suitable Technologies for their conference sponsorship of the BeamPros. Lastly, we thank NSERC and the GRAND NCE for funding portions of this research.

REFERENCES

1. Michael Boyle, Carman Neustaedter, and Saul Greenberg. 2009. Privacy Factors in Video-Based Media Spaces, *Media Space 20+ Years of Mediated Life*, Springer.
2. Juliet M. Corbin and Anselm Strauss, 2008. Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory, Sage Publications.
3. Andy Crabtree. 2004. Design in the absence of practice: breaching experiments. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques* (DIS '04). ACM, New York, NY, USA, 59-68. DOI=10.1145/1013115.1013125
4. Paul Dourish. 1993. Culture and control in a media space. In *Proceedings of the 3rd European Conference on Computer-Supported Cooperative Work* (ECSCW'93) (Milan, Italy). Kluwer Academic Publishers, Dordrecht, 125--138.
5. Thomas Erickson, N. Sadat Shami, Wendy A. Kellogg, and David W. Levine. 2011. Synchronous interaction among hundreds: an evaluation of a conference in an

- avatar-based virtual environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '11). ACM, New York, NY, USA, 503-512. DOI=10.1145/1978942.1979013
6. Harold Garfinkel, H. 1967. Studies of the routine grounds of everyday activities, *Studies in Ethnomethodology*, Englewood Cliffs, New Jersey: Prentice-Hall, 35-75
7. Akira Hayamizu, Michita Imai, Keisuke Nakamura, and Kazuhiro Nakadai. 2014. Volume adaptation and visualization by modeling the volume level in noisy environments for telepresence system. In *Proceedings of the second international conference on Human-agent interaction* (HAI '14). ACM, New York, NY, USA, 67-74. DOI=10.1145/2658861.2658875
8. Kori Inkpen, Brett Taylor, Sasa Junuzovic, John Tang, and Gina Venolia. 2013. Experiences2Go: sharing kids' activities outside the home with remote family members. In *Proceedings of the 2013 conference on Computer supported cooperative work* (CSCW '13). ACM, New York, NY, USA, 1329-1340. DOI=10.1145/2441776.2441926
9. Jia Chun, Popescu, V., Dark, M., York, C. 2008. Virtual Classroom Extension for Effective Distance Education, *IEEE Computer Graphics and Applications*, Volume 28 (1), 64-74.
10. Erving Goffman. 1959. *The Presentation of Self in Everyday Life*, Doubleday.
11. Serena Hillman, Azadeh Forghani, Carolyn Pang, Carman Neustaedter, and Tejinder Judge. 2014. Conducting Interviews with Remote Participants, In Tejinder K. Judge and Carman Neustaedter (Eds), *Studying and Designing Technology for Domestic Life: Lessons from Home*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
12. Scott E. Hudson and Ian Smith. 1996. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (CSCW '96), Mark S. Ackerman (Ed.). ACM, New York, NY, USA, 248-257. DOI=10.1145/240080.240295
13. Steven Johnson, Irene Rae, Bilge Mutlu, and Leila Takayama. 2015. Can You See Me Now?: How Field of View Affects Collaboration in Robotic Telepresence. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 2397-2406. DOI=10.1145/2702123.2702526
14. Michael L. W. Jones. 2000. Collaborative virtual conferences: using exemplars to shape future research questions. In *Proceedings of the third international conference on Collaborative virtual environments* (CVE '00), Elizabeth Churchill and Martin Reddy (Eds.). ACM, New York, NY, USA, 19-27. DOI=10.1145/351006.351009
15. Norman P. Jouppi. 2002. First steps towards mutually-immersive mobile telepresence. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work* (CSCW '02). ACM, New York, NY, USA, 354-363. DOI=10.1145/587078.587128
16. Atsunobu Kimura, Masayuki Ihara, Minoru Kobayashi, Yoshitsugu Manabe, and Kunihiro Chihara. 2007. Visual feedback: its effect on teleconferencing, *Proc of the Conference on Human Computer Interaction*, Springer-Verlag, LNCS 4553, 491-600.
17. Annica Kristoffersson, Silvia Coradeschi, and Amy Loutfi. 2013. A review of mobile robotic telepresence. *Adv. in Hum.-Comp. Int.* 2013, Article 3 (January 2013), 17 pages. DOI=10.1155/2013/902316
18. Min Kyung Lee and Leila Takayama. 2011. "Now, i have a body": uses and social norms for mobile remote presence in the workplace. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '11). ACM, New York, NY, USA, 33-42. DOI=10.1145/1978942.1978950
19. Tristan Lewis, Jill Drury, and Brandon Beltz. 2014. Evaluating Mobile Remote Presence (MRP) Robots. In *Proceedings of the 18th International Conference on Supporting Group Work* (GROUP '14). ACM, New York, NY, USA, 302-305. DOI=10.1145/2660398.2663777
20. Joseph F. McCarthy and danah m. boyd. 2005. Digital backchannels in shared physical spaces: experiences at an academic conference. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '05). ACM, New York, NY, USA, 1641-1644. DOI=10.1145/1056808.1056986
21. Stanley Milgram, Hilary Liberty, Raymond Toledo, and Joyce Wackenhut. 1986. Response to intrusion into waiting lines, *Journal of Personality and Social Psychology*, Vol. 51(4), Oct 1986, 683-689.
22. Carman Neustaedter and Saul Greenberg. 2003. The Design of a Context Aware Home Media Space, *Proceedings of the International Conference on Ubiquitous Computing*, Springer, 297-314.
23. Carman Neustaedter, Saul Greenberg, and Michael Boyle. 2006. Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Trans. Comput.-Hum. Interact.* 13, 1 (March 2006), 1-36. DOI=10.1145/1143518.1143519
24. Andreas Paepcke, Bianca Soto, Leila Takayama, Frank Koenig, and Blaise Gassend. 2011. Yelling in the hall: using sidetone to address a problem with mobile remote presence systems. In *Proceedings of the 24th annual ACM symposium on User interface software and technology* (UIST '11). ACM, New York, NY, USA, 107-116. DOI=10.1145/2047196.2047209

25. Eric Paulos and John Canny. 1998. PRoP: personal roving presence. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '98), Clare-Marie Karat, Arnold Lund, Joëlle Coutaz, and John Karat (Eds.). ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 296-303. DOI=10.1145/274644.274686
26. Irene Rae, Bilge Mutlu, and Leila Takayama. 2014. Bodies in motion: mobility, presence, and task awareness in telepresence. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems* (CHI '14). ACM, New York, NY, USA, 2153-2162. DOI=10.1145/2556288.2557047
27. Irene Rae, Leila Takayama, and Bilge Mutlu. 2012. One of the gang: supporting in-group behavior for embodied mediated communication. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12). ACM, New York, NY, USA, 3091-3100. DOI=10.1145/2207676.2208723
28. Irene Rae, Leila Takayama, and Bilge Mutlu. 2013. The influence of height in robot-mediated communication. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction* (HRI '13). IEEE Press, Piscataway, NJ, USA, 1-8.
29. Shervin Shirmohammadi, Shun-Yu Hu, Wei Tsang Ooi, Gregor Schiele, and Arno Wacker. 2012. Mixing Virtual and Physical Participation: The Future of Conference Attendance? *Proceedings of the Conference on Massively Multiplayer Virtual Environments*.
30. Strickland, E. Should I Attend a Conference Via a Telepresence Robot?, IEEE Spectrum, <http://spectrum.ieee.org/autoton/robotics/industrial-robots/attending-conference-via-telepresence-robot>
31. Leila Takayama and Janet Go. 2012. Mixing metaphors in mobile remote presence. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* (CSCW '12). ACM, New York, NY, USA, 495-504. DOI=10.1145/2145204.2145281
32. Leila Takayama and Helen Harris. 2013. Presentation of (telepresent) self: on the double-edged effects of mirrors. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction* (HRI '13). IEEE Press, Piscataway, NJ, USA, 381-388.
33. John C. Tang. 2007. Approaching and leave-taking: Negotiating contact in computer-mediated communication. *ACM Trans. Comput.-Hum. Interact.* 14, 1, Article 5 (May 2007). DOI=10.1145/1229855.1229860
34. Katherine M. Tsui, Munjal Desai, Holly A. Yanco, and Chris Uhlik. 2011. Exploring use cases for telepresence robots. In *Proceedings of the 6th international conference on Human-robot interaction* (HRI '11). ACM, New York, NY, USA, 11-18. DOI=10.1145/1957656.1957664
35. Gina Venolia, John Tang, Ruy Cervantes, Sara Bly, George Robertson, Bongshin Lee, and Kori Inkpen. 2010. Embodied social proxy: mediating interpersonal connection in hub-and-satellite teams. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '10). ACM, New York, NY, USA, 1049-1058. DOI=10.1145/1753326.1753482