Shared Telepresence Robots at ACM CSCW 2016

Carman Neustaedter[1], John Tang[2], Samarth Singhal[1], Rui Pan[1], Yasamin Heshmat[1] and Azadeh Forghani[1]
[1]School of Interactive Arts and Technology, Simon Fraser University, [2]Microsoft Research
carman@sfu.ca, johntang@microsoft.com, samarths@sfu.ca, ruip@sfu.ca, yheshmat@sfu.ca, azadehf@sfu.ca

ABSTRACT
Attendance at academic conferences via telepresence robots is now a reality. Yet it is not clear how to design for shared usage of telepresence robots. We conducted a study of telepresence attendance at the CSCW 2016 conference where we explored the use of shared telepresence robots, asynchronously by different attendees based on varying time slots, and synchronously by multiple attendees using a single robot concurrently. We compared this experience to attendees using robots dedicated to just them. Shared telepresence robots allowed users to overcome feelings of obligation, commitment, and responsibility. However, shared robots also created feelings of disappointment and awkwardness when it came to visually personalizing the robot. Shared robots caused remote attendees to lose out on the autonomy that comes with a dedicated robot along with the ability to regulate one's solitude and choose when and how to interact with others. These results suggest alternative telepresence robot designs that better support autonomy and control for remote attendees where designs move beyond single-user models to better suit shared usage and methods to gain support from local attendees.

ACM Classification Keywords
H.5.3 [Information interfaces and presentation]: Group and Organization Interfaces - *CSCW*

INTRODUCTION
Telepresence attendance at academic conferences is now a reality where we are seeing conferences utilize telepresence robots [24,28,31]: video conferencing systems that users can control and drive through remote environments. Studies have found telepresence attendance at confereces to be valuable in situations where it is not easy for a person to attend in person, e.g., accessibility challenges, cost restrictions, visa problems, lack of time for travel [28]. The video conferencing capabilities of the telepresence robots along with their mobility make it so that remote attendees can see presentations and mingle during social events [28].

To date, explorations of the use of telepresence robots has often focused on the use of dedicated robots where one person uses a single robot (e.g., conference attendance [12,28], home usage [46], work usage [25]). In contrast, shared telepresence robots have the potential to reduce a conference's cost for remote attendance by having multiple attendees use a robot at different points in time [28,31]. This is possible because remote attendees often want to attend at different times due to time zone differences and work schedules [31]. A study of telepresence at ACM CHI 2016 involved the use of shared telepresence robots, though the research focus was on understanding how the large-scale of the conference affected telepresence attendance, rather than shared robots [31].

Figure 1: Beams at CSCW 2016.

In contrast, our focus was on studying the specifics of shared telepresence robot usage at a medium-sized conference: the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW) 2016, which had approximately 750 local attendees, took place on two floors of a hotel, and had five parallel tracks. This creates a different setting than the aforementioned large-scale conference study. With CSCW, because the venue is smaller, especially compared to conferences like CHI, social mingling is more readily available as an activity for remote attendees due to shorter driving distances [31].

At CSCW, 19 people remotely joined the conference via eight telepresence robots (Figure 1). A portion of remote attendees shared robots while others used robots dedicated to just them as a point of comparison. Our study was exploratory where we were interested in seeing what emergent behaviors might occur. We wanted to learn how shared robots would benefit remote attendees and generally be used; how and why the use of shared robots might create challenges for remote attendance; and, if and how the design of telepresence robots may or may not map to the needs of remote attendees when robots are shared.

Cite As:

Neustaedter, C., Tang, J., Singhal, S., Pan, R., Heshmat, Y., and Forghani, A. (2017) Shared Telepresence Robots at ACM CSCW 2016, Connections Lab Technical Report 2017-0823-01, Simon Fraser University.

Our results show that shared telepresence robots offer benefits for remote attendees as they can help alleviate feelings of obligation, commitment, stress, and anxiety. We also observed unexpected sharing practices emerge as a result of participants being given flexibility in terms of how they used their telepresence robot. Shared telepresence robots supported *both* asynchronous and synchronous remote attendance. By asynchronous, we refer to people using the same robots but at different times, such as a group of graduate students time-sharing a single robot. By synchronous, we refer to behaviors where a robot was used by multiple people concurrently to support activities such as remote teaching and simultaneous multi-person viewing of talks and mingling. Such sharing practices were not reported for the study at CHI 2016 [31] and nor the study at Ubicomp/ISWC 2014 [28]. The learning of and adherence to social norms also became a shared responsibility where local attendees aided remote users.

Yet shared telepresence robots were not used without their challenges. The robots were ostensibly designed for single users making synchronous shared usage awkward in social interactions as people had to give-up feelings of control and autonomy. Unlike telepresence robot usage at CHI 2016 [31], robot personalization was very important for remote attendees; however, personalization was difficult because people were not always in complete control of how they appeared. These results point to the need to consider methods to support robot appropriation for individuals and small groups; forms of digital personalization; and, ways to support the creation and evolution of social norms around telepresence attendance. This extends the ways in which we think about identity, autonomy, and control when using video-mediated communication systems.

RELATED WORK

Remote Conference Attendance

Conference attendees have been able to remotely attend academic and industry conferences to varying degrees over the past two decades, including through avatars [8,16], mixed environments with both local attendees and avatar-based remote attendees [39], and telepresence robots [28,31]. Studies of these events show that remote attendance enabled a broader level of participation [8,16, 28,31,39]. People especially enjoyed social mingling through avatars and telepresence robots at mid-sized conferences [8,28,39] despite challenges with recognizing people, body language [8,16,28,36,39] and moving into and out of group conversations [8]. On the other hand, socializing was much more difficult and of secondary importance at large-scale conferences through telepresence robots [31]. In the case of CHI 2016, remote attendees said that the size of the CHI 2016 venue (~143,000 sq. ft) and large volume of local attendees (~3700) made them less concerned about activities such as presenting their identity and social mingling [31]. On the other hand, personalization of oneself (either via an avatar or decorations on a robot) was seen as being valuable at mid-sized conferences [8,28].

Remote attendees at both avatar-based conferences and conferences with telepresence robots also faced challenges related to automony [5]—how, when, and in what ways people were able to remotely attend the conferences. For example, time zone differences created issues for attending large portions of the events [8,16,28,31,39]. Researchers found that people had difficulties in being virtually present at the conference while dealing with life at their physical location, including work and personal obligations [8,28,31].

In contrast to this research, our work directly explores the use of *shared* telepresence robots. Here we investigate both asynchronous sharing of telepresence robots (where people pick time slots) and synchronous sharing (where people concurrently use the same robot). Only one telepresence robot was shared at the Ubicomp/ISWC 2014 conference [28]; thus, this paper does not explore shared robots in detail. The study at CHI 2016 explored 33 people attending using 10 asynchronously shared telepresence robots [28]. The shared use of robots caused some attendees to pre-position Beams in sessions to avoid long driving times across the venue and scheduling collisions occurred when people would try to sign into a Beam that was still in use [31]. Our study builds on this work to more deeply explore the use of shared telepresence robots in a medium-sized conference where fewer local attendees and shorter driving distances create an atmosphere more focused on social mingling than large-scale conferences [31]. We also explore synchronous sharing, in addition to CHI 2016's focus on asynchronous sharing only.

Workplace Settings

Telepresence robots have been studied in work settings and results have pointed to findings that are similar to conference attendance, e.g., issues with the presentation of self, social interactions, and telepresence robot design. However, these studies all focus on situations where a single person is using a telepresence robot and do not consider design needs when robots are shared amongst multiple people. Telepresence robots dedicated to a single person have been found to support group tasks [30] and promote knowledge of availability and social presence amongst co-workers [25], though new social norms have to be established for interactions [25,42]. Co-workers local to the robots have been found to help them out with elevators and charging [25], which is similar to what happens at conferences [28,31]. In some instances, people treat a remote person like a robot or object rather than a person [25], again found at conferences [28,31]. Wide field or panoramic views are valuable [17,18,22] for supporting peripheral awareness [17]. Similarly, mobility is important for task completion [18,33] and varying audio levels for conversations [18]. Some have suggested setting volume levels based on ambient noise [14]. In addition, there is a need to show a robot's location on a venue map [43].

Researchers have distilled the importance of appearance and how one looks in a telepresence robot [25]. This is similar to how one creates and shares their identity with others when in person [11]. Goffman describes such acts of human self-presentation as being similar to how a person presents himself or herself as an actor on stage with a certain appearance, attitude, and corresponding sequence of actions and interactions with others [11]. These interactions reinforce and sometimes extend one's conceptual model of identity [11]. Goffman's framework has been shown to be useful for understanding identity presentation in a telepresence robot [28]. With telepresence robots, researchers point to the need to present oneself in an appropriate way where one has both visual and audio feedback since it is hard to know how loud one sounds in the remote location and what one looks like [25,29,41,43]. Robot height has even been shown to affect persuasiveness [33]. While beneficial, this research does not explore personalization needs when multiple people share telepresence robots. This is a focal area of our work.

Video-Mediated Communication

More broadly, there is a vast amount of literature and theory on video-mediated communication, which our study explores and extends in terms of its relationship to shared telepresence robot usage and design. In workplaces, video communication systems have been shown to be valuable for casual interactions [45] and maintaining workplace awareness of others [2,4,9]. Here it is common for multiple people to share video connections as a part of meetings (e.g., one or more colleague uses Skype to connect in). Yet there is a tendency to design new systems with a focus on assigning one-person to each device or video display (e.g., [27,44]). This trend continues in the design of commercial telepresence robots that often look like a single person, e.g., with 'legs' and a single 'body' and 'head' (e.g., [38]).

In domestic settings, we see video chat systems being valuable for sharing conversations and joint activities [6,16,19,23]. Video communication systems tend to be used and designed in a way that turns the device (e.g., a tablet) into a 'portal' showing the remote location (e.g., [20,21,34]. This allows multiple people to gather around the display, such as groups of children or other family members, so that the systems support synchronous shared usage [1,10,23]. Synchronous usage sometimes involves scaffolding where parents help children use the video system [10]. Potentially negative social norms have emerged around domestic video chat usage where people might feel a sense of obligation to stay in view [13,23], conversations become constrained [13], or people do not know what to share [35]. Our work builds on this research to investigate synchronous and asynchronous shared usage of a video communication system, albeit one built into a telepresence robot, as well as the development of social norms around such usage.

There is also a longstanding history of investigating privacy and the ways in which people choose to participate in a video-mediated communication system (or not) and how designs support or do not support people's behaviors and needs [2,5]. This often relates to interrelated components of privacy, including solitude, confidentiality, and autonomy [5]. Solitude refers to people's ability to choose when to interact; confidentiality relates to what others know of oneself; and, autonomy refers to people choosing when and how they participate in a space [5]. Reciprocity has been shown to be important in video communication; this is the idea that if a person can see someone, that person can also see him or her [3]. When people can see or hear others but not be seen or heard themselves, there may be issues around disembodiment [2]. In order to design systems in a way that mitigates privacy issues, researchers have suggested designs that give users control over how they participate, such as flexibility to choose when, where, and how a connection is made [2,3,7,20]. Our research explores the aforementioned theories of privacy in video-mediated communication systems to understand and suggest design direction for shared telepresence robots.

STUDY METHOD

We studied telepresence attendance at CSCW 2016 with the goal of understanding the effects of shared telepresence robots on the remote user experience. We compared this experience to the use of dedicated telepresence robots. The study was approved by our university research ethics board.

Remote Attendees

We advertised telepresence attendance at CSCW via the conference web site and social media. We selected three attendees for workshops (based on whether they had a workshop paper accepted) and eight applicants for the main conference (only eight applied). Workshop attendees were assigned a dedicated robot for their workshop. Prior to the conference, we had attendees for the main conference (Monday to Wednesday) select which sessions they thought they would attend. Those with more than five sessions (five people) were assigned a dedicated telepresence robot that they could use at any point throughout the conference, while the remaining participants (three people) were assigned to shared robots with other remote attendees. Those sharing robots were told they could connect into the conference during only the sessions they pre-selected. In addition to our predefined sharing, two of the individuals selected for the main conference program decided to appropriate their dedicated robot in a manner that also made it shared: they shared their robot amongst graduate students in their university research group. In the first case, a robot was shared between one faculty member and two graduate students. In the second case, a robot was shared amongst seven grad students while the faculty member was at the conference in person. Overall, our participants consisted of:

- 3 workshop attendees using dedicated robots (D1-D3)
- 3 conference attendees using dedicated robots (D4-D6)
- 13 conference attendees using shared robots (S1-S13)

Remote attendees comprised of 11 graduate students and 8 faculty members. 10 participants were from the USA, 5 from Canada, 2 from Brazil, and 2 from China. One participant faced accessibility challenges that made it difficult to travel, while the rest faced either time or cost restrictions, or wanted to participate simply to try out a novel experience. Nobody was located in the same time zone as the conference; time zones for those at the main conference varied from 1 to 5 hours time difference. Two workshop attendees faced a time zone difference of 16 hours and one was 3 hours different. Remote attendees paid half of the cost of regular attendance given that the setup was still somewhat experimental. Participants completed an informed consent process prior to the conference.

Telepresence Setup

We used eight BeamPros (hereafter called Beams) [38] as the telepresence robots (Figure 1). Beams have two cameras, one facing forward and one facing the floor for navigation. Camera resolution is 480p with a field of view of 105° and 3x digital zoom. Users control movement with a keyboard, mouse or XBox controller.

The conference spanned two floors of the hotel's meeting space and there was a small elevator that connected the locations. Beams were docked inside a meeting room on the second floor and we placed a large mirror at the door of the room so that remote attendees could see themselves as they left the room; for workshops, we docked the Beams inside the workshop rooms. We created a group instant messaging backchannel in Skype and asked all remote attendees to connect to it while they were at the conference so we could provide them with technical support as needed and give them a channel to communicate with other remote attendees.

Prior to the conference, we asked remote attendees to mail us items to personalize their Beams so that each robot would be visually identifiable where knowledge of who each remote attendee/robot was might be built up over time. Such personalization was found to be important in past avatar-based conferences [8] and remote attendance via Beams, although personalization was only studied with dedicated robots and not *shared* ones [28].

We tested various items prior to the conference and found that items that could be affixed to the top of the Beam worked best since they were visible from a greater distance away. For these reasons, we asked participants to send us a hat, scarf, nice wig, or a bandana to decorate their Beam. Participants were told that if they did not, we would perform our own customizations using items we had on hand. We received items from seven people, including 3 hats, 2 scarves, 1 hat/scarf/pennant combination, and a shirt (which we affixed like a scarf). Both Beams that were shared amongst their university's research lab sent a single set of personalization items to be used by all people from that institution. Prior to each Beam's usage, we affixed the appropriate person/institution's item to the Beam. We attached a colored balloon to the top of Beams in cases where the person did not send us an item; the balloons bobbed slightly as the Beam moved which was purposefully chosen to help remote attendees be seen from far away. Thus, across all 19 remote attendees, 14 had personalized robots (albeit two were personalized for their research groups and not individuals) and 5 had colored balloons.

Data Collection and Analysis

Semi-Structured Interviews

We conducted interviews with each remote attendee within several days of the end of the conference over Skype [15]. Interview times varied between 15 and 60 minutes, with most taking longer than a half hour. One participant was not responsive to our emails and so we did not interview this person. Another had a very busy schedule so we exchanged emails asking a subset of the interview questions. Interview questions focused on the experiences around sharing a robot, robot personalization, navigation, social interactions, privacy concerns, and any additional topics that emerged from participants' experiences. All interviews were audio-recorded and transcribed, except for one where recording failed. We used thematic analysis to understand the main themes within our data. This involved multiple reviews of interview transcriptions and meetings amongst the paper authors to discuss their observations.

Observations

We observed the remote attendees throughout the conference, during sessions, breaks, and at the demo reception. We could not feasibly observe all of the sessions that the remote attendees went to, however, we managed to have at least one observer attend 19 of the 27 sessions that remote attendees were at along with breaks between sessions. Observers took notes about where the Beam parked, how they moved through the session or space, whether or not they asked questions, and if and how they interacted with people. We recorded images and videos at various points to capture the different types of interactions that occurred and the placement of Beams in sessions.

We did not collect observations of a class who shared a Beam with their course instructor (described more later) since the complexity of obtaining informed consent from all students in the class was high. Instead, we collected interview data from the course instructor as well as two students in the class, C1 and C2 (they responded to our call for interviews with students). These two students completed consent forms prior to being interviewed.

The amount of time that remote participants attended CSCW varied broadly. We tracked this time using their Beam login accounts. Two accounts were shared amongst multiple people at the institution, so we were only able to track usage for the entire institution. The mean time spent connected to a Beam with a single account was 9.9±6.7 hours (med=8.3, min=1.9, max=20.4 hours). Five accounts used a Beam the full 3 days of the main conference, two accounts used a Beam 2 days, and three accounts 1 day. Remote attendees primarily attended conference talks. Two remote attendees attended the poster and demo reception. Workshop attendees spent the majority of their time within their workshop room, though they were exposed to the larger conference audience (beyond workshop attendees) during break times which occurred in the hotel's foyer.

Backchannel Messaging

All remote attendees joined the Skype backchannel along with five support staff and one local attendee (the faculty supervisor for one of the institutionally-shared robots). Attendees made a total of 267 posts with an average of 22.8±26.1 posts each (median=16, min=1, max=104). Support staff made a total of 168 posts with an average of 44.6±37.2 posts each (med=39, min=11, max=105). The local attendee made 18 posts. We kept logs of the messages in the Skype backchannel and performed coding on a per-message basis to identify the type of content in the message. Codes were preselected as technical questions, answers, and comments and social questions, answers, and comments. We also coded each message according to who was the sender and the most likely intended recipient. Our goal was to understand the main uses of the backchannel. Two people independently coded the backchannel messages to understand who messages were being sent between and whether messages were about technical or social issues. Interrater reliability was 0.92 for recipients and 0.72 for the content of the message. Raters discussed the divergent codes and came to a consensus (reported in our results).

Next we describe our main themes where first talk about the ways that robots were shared asynchronously, followed by synchronous usage, personalization, and the ways that knowledge was shared between remote and local attendees.

ASYNCHRONOUS SHARING OF ROBOTS

Beams were shared asynchronously between participants from different institutions and, as mentioned, Beams were also shared asynchronously in an opportunistic fashion where people from the same research group shared a single Beam. In the following sections, we talk about the benefits and challenges of both types of asynchronous sharing.

Obligation, Commitment, and Responsibility

First, there were clear distinctions between the ways that participants thought about and used the Beams based on whether they were using a shared or dedicated Beam. Those who attended CSCW using a dedicated robot very carefully planned their time at the conference. Some even went to great lengths to block their calendars or let others at their local location know that they were unavailable for the duration of the conference. Thus, they worked hard to ensure their time was committed to what was happening at the *remote* location (the conference) and not at their *home* location. They also had a sense of responsibility for the robot in that they needed to occupy it on a daily basis and be responsible for its placement throughout the conference venue (e.g., returning it to its dock for charging). There was also a strong sense of the robot representing them and being "there's."

"I kind of warned people I'm going to be at the conference this week… we don't have classes yet this week. So I didn't have to teach or anything. Otherwise I would have had to make arrangements. "- D3

Being a dedicated Beam user was not always easy. Piloting the Beam and using it for long periods of time was challenging; participants said it required one to pay a lot of attention to drive around and also converse at the same time. Those in workshops faced challenges in trying to be 'present' for an entire workshop (~8 hours). Many dedicated robot users reported being fatigued and having to take breaks for exercise or to rejuvenate. Thus, while participants had complete autonomy over when and how they participated in the conference, such flexibility and the ability to participate extensively in the remote location actually made the experience quite challenging at times.

In contrast, those participants who were using shared Beams exhibited behaviors that would suggest much less attachment to their Beam and commitment to being in attendance, despite signing up to attend the conference at specific points in time. They described not feeling responsible for their robot and its whereabouts in the same way that dedicated robot users described their responsibility. Participants told us that they did not feel guilty or obligated to be at the conference the entire time and that they could balance conference attendance with other activities happening at their home location. Some were even not bothered if they missed one of their time slots. Thus, the shared nature of the robots meant participants were able to regulate their solitude in terms of their time away from the conference for themselves, as well as their autonomy, the ways in which they wanted to participate or not in the remote location.

For example, S11 found that family activities, last-minute work meetings, and casual interactions with co-workers would disrupt his attendance at CSCW and became a higher priority since they were happening in-person. Thus, his commitment to activities happening locally superseded any remote commitments he might have. A similar finding was found for avatar-based conferences [8]. S11 did not feel responsible for his robot and, instead, just wanted to see particular talks. As a result, he began to have his graduate students (who were sharing the Beam with him) move the Beam to locations that he may want to connect into.

"I was more obligated to attend to my real life meetings. I told my students, why don't you drive the Beam and just leave it there and when we get a chance between our schedule, we can come in and watch the session…Beam activities were secondary and my real life was more important."- S11

Like past studies [28,31], remote attendees received a lot of attention because they were in a Beam and it was still a novel technology. Not everyone was comfortable with such attention and the fact that they used a shared robot meant that they could try to avoid situations that might produce the most amount of anxiety or stress, which, for many, was social events. Instead, those with anxiety could let their colleagues who shared the Beam with them attend such activities and not feel guilty for missing out since someone else was making use of the Beam.

"I didn't apply for any of the social events because I thought that'd be so awkward…I have never been in the Beam before…I'm very junior in the field so I don't know many people there. Most would be strangers." – S9

"I don't have to stay there all the time. Using the Beam is pretty stressful, especially when I was listening to talk. I didn't know I was so obvious in the back of the conference room…. If you are on the Beam, you must actually look like you are paying attention. That's stressful.." – S5

Scheduling and Disappointment

One of the institutionally-sharing groups found it easy to coordinate the use of the Beam because there were only three of them sharing it. The second group shared a single robot amongst seven graduate students and faced scheduling challenges and disappointment as a result of the 'scarce' Beam resource. Students recorded what sessions they wanted to attend and then a senior student made the decision about who to let use the time slot. This was based around people's research topic and a desire to balance out the number of time slots that each person received. Despite efforts to ensure fairness, not everyone was able to see their preferred events. In this case, one's ability to exercise autonomy and participate in the remote location as desired was constrained by the pragmatics of having only one Beam available across several people.

"It was good because maybe some prefer morning times, some prefer evening times. Always there was someone on the Beam. As a lab we didn't miss anything. But as a person… personally I miss many, like there was an award session, so I missed that. And I really wanted that…I would have liked to see how the Beam could see the posters." – S6

Changing Personalization Items

For asynchronously shared Beams, we had to ensure that each Beam had the proper personalization items on before the remote attendee connected. This involved changing items during breaks for the Beams not shared within the same institution. This was challenging to time and meant that remote attendees had to return to the docking room at the breaks between sessions, which reduced mingling during breaks. In only one case was there ever an accidental situation where a Beam was not properly personalized for the remote attendee and the remote attendee connected into a Beam with the wrong appearance items on.

"I beamed into the one that was on the schedule that was Beam #6 and that ended up being [another person's] Beam and there was also that feeling that I was in someone else's Beam. Not that I'm in that person's body, but someone else's possession." – S3

While this was an isolated incident, it illustrates the potential for coordination issues around the personalization of shared robots. It also raises larger implications around the practical challenges of relying on others to help one construct attributes of his or her identity in a remote space when using a telepresence robot. When at a conference in person, people are able to construct the way they look as a part of their identity on their own, where this identity is then reinforced by local people who recognize them, interact with them, etc. [11]. Yet with physical embodiments such as telepresence robots, if physical objects are used to augment or present one's appearance, remote attendees lose their autonomy as a part of identity presentation.

SYNCHRONOUS SHARING OF ROBOTS

Beyond the asynchronous, turn-taking form of sharing, we saw sharing emerge as a simultaneous group activity. This was despite Beams primarily being designed for use by a single person at any given time. This occurred in two main ways: large group and small group sharing.

Large Group Sharing

One remote attendee who was a professor used his Beam to attend CSCW while teaching his undergraduate class of 32 students. He situated himself in the middle of the camera's view and the rest of the class was positioned behind him in a lecture theatre. They attended a CSCW session on education. While a potentially valuable learning experience, this setup juxtaposed varying needs for solitude and autonomy within the video-mediated environment. First, this setup gave the professor the opportunity to show his students the latest work in the field they were studying; thus, it allowed him to fulfill his goal of enhanced teaching. For example, the professor used naturally occurring social interactions with local attendees as opportunistic teaching moments about the technology.

"I found people treated me different ways. I pointed this out to the class to watch for… Some people go out of their way to interact with the Beam…Coming into a room, everyone just looks over and smiles at you. In person, I'd wonder why. Do I have something in my teeth? Do I have my fly open?" – S3, Professor

Yet because the robot was ostensibly designed for only a single user to operate, this shared usage did not allow the students to exercise their own autonomy to participate in the remote space as they might see fit. For example, we interviewed two students from the class and one of them wished he could have been able to experience firsthand what it was like to interact with people using the Beam. The other student did not want to directly interact with people through the Beam because of perceived social pressures and so he was content on only watching. In this way, the setup afforded him the opportunity to see what interactions might be like without infringing on his desire to not have to actually interact himself.

We also learned that both students felt a lack of social responsibility for the Beam, including where it was positioned in sessions and how it moved throughout the space because they were 'just watching' rather than driving it. They felt the responsibility for the Beam lay solely on the professor. In addition, the setup raised questions around confidentiality and who would be viewed on camera when varying, and possibly competing, privacy needs may be present amongst the students.

"It certainly feels like I'm less of a head and more of a portal. And people are seeing into my world…I got a webcam, not just the laptop, and kind of propped it up so we could get a wider view of the class. So people could see we were the class. I also reminded people that if you don't want to be seen then don't sit in this general area. Sit on the sides…. We did have one student who sat in that area because they didn't want to be seen." - S3, Professor

Both of the students interviewed from the class described purposely sitting in locations in the classroom where they could not be easily seen on camera. This allowed them to meet their own social needs of not wanting to be seen, though it meant that people local to the conference did not know they were watching. Neither student was concerned about this lack of reciprocity as they felt that only the active driver of the Beam needed to be visible on camera.

"I was in the row behind him but outside of the camera range…I didn't really want to be seen by people…I have had this phobia of not liking having my picture taken…It was interesting to see how people interacted with [my professor]." – C2

While the setup benefitted the students as it allowed them to remain anonymous, they were in fact disembodied in the remote environment yet still able to see and hear what was going on. This lack of reciprocity could easily raise concerns for local attendees in the conference location.

Small Group Sharing

Beams were also shared by small groups of two to three remote attendees at the same time where one person would operate the Beam and the other attendees would sit next to this person so they could be on-screen and see what was happening at the conference. Remote attendees used this setup to attend talks that they were both interested in. However, like large group sharing, because the robot was ostensibly designed for a single user rather than side-by-side pairs, the experience was challenging.

First, the view of the robot's cameras on the display were challenging to see given that two or three people were sitting next to each other. When viewing on a laptop screen, this meant that the users had to sit very close to one another and often inside one's personal space. Second, the setup reduced one's ability to multi-task. That is, it interfered with one's solitude and ability to 'do things on my own' if desired. It also meant that possibly confidential information might be visible on the screen of an individual, which was now visible to the person sharing the computer. Participants were also constrained in that they had to watch from their research lab, as opposed to being at home; this raises questions around their autonomy and where they participate in the mediated environment from.

"I attended two sessions and in one of them one of my student was interested, so he came and sat next with me in that session...We sat on the couch, and perched the laptop on the coffee table. Easily, the biggest challenge was that the view of the screen just wasn't as good. I think the other thing was that it prevented me from doing the normal sorts of a "side channel" activity (check email, etc.) that I would normally do." – S11When social interactions occurred, the shared robot configuration created confusion over who a local person was talking to. It also had the potential to interfere with people's solitude and desires to engage in conversation or not. For example, participants talked about one person wanting to engage in conversation, while the other did not.

"I think it was also a bit more confusing for others to interact with us, as it wasn't always clear who they were saying "Hi" to if the camera feed was choppy/block" - S11

"We both thought we should go to the demo night. But it was a bit awkward to have two people sharing a single Beam, one driving and two people sitting in the frame. It's kind of awkward when there are two people within the representation of what is one person...During paper sessions, however, this isn't really a problem, since people are generally paying attention to the talks and not to the other people in the audience." – S12

PERSONALIZATION

When in person at a conference, people can select and present their appearance on their own and they can change it as much as desired. Yet with telepresence robots, given the practical challenges of augmenting a telepresence robot's appearance, remote attendees were faced with constraints around how to present their appearance and it relied on remote people, the telepresence organizers, to facilitate the process. We describe how participants worked within these constraints to present their identity and what problems it created for shared Beams next.

Institutional Branding

First, participants who picked their personalization items very carefully thought about what facet of themselves should be the focal point of their Beam's customization. In every case, the Beams' appearances became focused around one's institution or location as a form of branding or advertisement, rather than a depiction of other attributes about the participant (e.g., gender, one's own appearance); this contrasts studies of avatar-based conferences [8,16,39]. These findings were universal across all participants who chose personalization items for their own Beam or their institution's Beam. Thus, we did not see differences between dedicated and shared Beams.

"At first when you told us to get the identity, what I thought I'd do was have a bandana with our university logo and a [country] flag…. So I looked in Amazon." – D3

"My school's color is red so and then also there was a flag with the school name. I feel like they will, the people will recognize that I'm, where I'm from." – S9

The selection of items to represent a location or institution was not always easy and participants said they struggled with the need to pick items that could be easily attached to the Beams, spend only a small amount of money (because they had to ship the items and may not get them back), and represent some facet of themselves. For example, S11 wanted to use a hockey jersey because he felt it represented his country, yet it would not have been possible to easily attach it to the Beams and it was very expensive.

The Environment and Setting

Second, participants thought about the environment that they would be a part of at the conference and how they should best match it while also 'standing out' and not looking generic. This was the case for both dedicated and shared Beam users and included trying to "look good" aesthetically or appear "cool". Such considerations were also found in avatar-based conferences [8,36].

"I was really worried what other Beams would look like… I didn't have the context of how cool other Beams would be. … Where's Waldo came up in kind of a jokester way. Where's the Beam… When I saw the whole lineup of all of the Beams, the initial concern about being cool, I thought we actually did a good job. We looked cool." – S8

Some participants (both dedicated and shared Beam users) were concerned about the seriousness of an academic conference given that it is a work setting, even though dress and demeanor at CSCW is often fairly casual. This was described by a participant who was in an institutionally-shared Beam as well as a dedicated Beam users who did not send in a personalization item and, instead, had a balloon attached to the top of her Beam. In the latter case, the participant wished she had an item that was more serious in nature than a balloon.

Self-Representation with Shared Beams

Despite the careful thought that went into personalizing the Beams' appearances, they were not satisfactory for everyone. In cases where the Beam setup was one of *synchronous* sharing, a single appearance did not raise issues. For example, the two students we interviewed from the class were fine with their Beam's personalization (with university colors) since they felt it should match their professor's desires for customizing the Beam and not their own. This could be attributed to the way in which the Beam was setup and used, with the professor as the focal point on camera and main driver. However, in cases where the Beams were *asynchronously* shared within a research group, dissatisfaction arose from five of seven participants who were not a part of the selection process for the Beam's personalization. In these cases, they had to give up control of their appearance to a designated person from their research group. That is, they lost autonomy over how they would be presented in the remote space. Rather than want institutional branding, these participants desired self-representation where they could present aspects of their own appearance and interests. For example, one shared Beam participant wanted her Beam to look very plain so that she would not stand out; she felt that the Beams were already attracting too much attention and she wanted to be less noticeable. One participant wanted to show more of her personal attributes, along with knowledge of her interests. Another participant was okay with representing her institution, but thought carefully about how her gender could be presented at the same time. She was also concerned about how the Beam's personalization was not what one would normally see indoors.

"Who wears a hat in somewhere? It's for outdoors. The outfit was more for the outdoors than the indoors. Who wears a hat in a conference? Our robot was more like a boy. The hat and scarf, it was more like a boy. Maybe I would be more girlish, feminine things…. There's some scarfs in the [school] store that are so nice and silky. They're like fancy…. The way that I myself wear a better dress for a more formal place." - S6

Overall, the challenge with the findings around self-representation is that we do not know if these participants would have *actually* made these selections if they were the ones choosing the personalizations, or if the concern was actually because of a loss of control over their autonomy and choice of self-representation at the remote location. Reactions from Local Attendees

For the most part, remote attendees received positive responses about their appearance when interacting with local attendees. Remote attendees were complemented for their personalization and many had pictures taken of them, likely because of the novelty of the technology and personalization. Only one participant (dedicated Beam) commented that she felt her appearance had received negative reactions from local attendees. She had not sent in a personalization item and, instead, had a balloon attached to her Beam.

"I wish I had sent/asked in advance for a more personal and maybe more 'serious' and not disposable object but I didn't get around to that, the balloon was good enough as far as bringing some color and personalization although it might have played into whatever enabled the sometime condescending/paternalistic/not taking seriously approach that I have felt from some attendants." – D4

KNOWLEDGE SHARING WITH LOCAL PROXIES

Remote attendance also involved local attendees helping remote attendees with social and contextual information. In this way, the concept of 'shared robots' also involved shared knowledge between *remote* and *local* attendees.

Beam Buddy

Like other conferences [28,31], using the Beam was not always easy in terms of navigation, parking, and social interactions. One of the institutions sharing the Beam recognized the value in having a local attendee available to help guide the Beam driver. They assigned a local attendee as a 'Beam Buddy' to help each remote. The Beam Buddy stayed with the robot at all times.

"We had a Beam buddy on the other side to help… It was more than navigating, it was helping us talk to people, introducing us." – S6

What worked especially well about the 'Beam Buddy,' as told to us by several participants, was a shared sense of trust. Because they had an established relationship with the person, they trusted the help that they received. At times a local attendee who already had a strong connection to the CSCW community played the role of Beam Buddy. This helped remote attendees who were graduate students meet people and feel like they were becoming more connected to the community. On the other hand, when the Beam Buddy was relatively new to the community, such connections were felt to be harder to create.

Backchannel

Help from local attendees occurred opportunistically for remote attendees who did not have a 'Beam Buddy.' This was uniform across attendees, regardless of whether they were in a shared or dedicated robot. This went beyond help with elevators and doors, as reported in past studies [25,28] to helping understand the local context and social interactions. In one of the workshops, this occurred through WeChat between a remote attendee who was from China and the workshop organizer.

"Because one of the workshop organizers was from China. So we had a WeChat group and the organizer from China guided me a lot (in WeChat)." – D3

More broadly, we found that despite the Skype backchannel being setup to offer technical support to remote attendees, remote attendees used it to gain *social* support from local attendees. 233 posts (39.7%) were about technical topics, such as questions about the Beams' features, comments on the telepresence setup, questions about Wi-Fi, etc. Yet more posts, 354 (60.3%), were about social topics such as social etiquette, social norms, etc. Posts between remote attendees and support staff accounted for 77.5% of posts, while only 12.5% of posts were between remote attendees.

Because support staff members were physically present in many of the sessions that the remote attendees were present in (for observation purposes), remote attendees would frequently ask questions of them in the backchannel. For example, they would ask about their positioning in the room, whether there was a better viewing location that was empty, if it was socially a good time to leave, if they were too loud, etc.; this contrasts prior studies of a backchannel at the CSCW 2004 conference where posts focused on presentation content [26]. In this way, support staff became local proxies for the remote attendees by providing them with information that was difficult to ascertain through the Beam. While the remote attendees did not necessarily have a long-term relationship with those supplying the advice and help, participants said that because support staff had experience with Beams, their judgment was trusted.

"I used it if it was something that I was not too sure of, what's the proper way to use the Beam, what's the proper social norm, if you are not allowed to do something." – S12

What seemed particularly important about the backchannel support that remote attendees were receiving was its perceived invisibility. That is, support could be given discreetly through the backchannel, rather than talking aloud in the conference sessions. This aided the remote attendees because it left the control of how they interacted with the Beam in their own hands, rather than giving it up to others. Remote attendees felt empowered as a result. Despite the benefit of the backchannel, participants experienced challenges typical of backchannels and threaded chat, e.g., information overload, distraction [37].

DISCUSSION AND CONCLUSIONS

We now discuss our results and their implications for the design of telepresence robots. First, we had imagined that shared robot usage would follow a model of participation that mapped to the design of the telepresence robots, e.g., one user at a time, fixed user changes according to a schedule. However, it was evident very quickly that telepresence attendance was much more about appropriation and allowing remote attendees to decide how to best make use of the opportunity that telepresence robots afforded them—e.g., shared by researchers at different institutions, time-shared amongst groups at the same institutions, simultaneously used by multiple users, used by an entire class. Thus, remote attendance was not always a *per person* activity like we expected. This was not found in previous studies of remote conference attendance via Beams [28,31]. Such synchronous usage is more akin to the way that a family might gather around a video chat device when at home [1,19,23] rather than how video might be used in the workplace within meetings or for one-on-one conversations [6,44].

Clearly social norms are developing and the model for remote attendance is not like the model for local attendance. In the local case, a person registers and attends the conference as him or herself. There is no other type of attendance. Yet remotely, there is a whole host of ways to utilize one's "attendance." The implication is that telepresence robots and associated video communication systems need to be designed to support a multitude of different sharing and telepresence needs. Certainly a 'one size fits all' model is not good enough when it comes to the design of telepresence robots for conferences or events similar in size and activities to CSCW. Differing attendance models require the design of different types of robots or telepresence technologies where users are able to exercise their autonomy to regulate how they participate in the remote space and regulate their solitude by choosing when to engage in interactions. Prior research does not suggest telepresence robot design directions based on changing needs around privacy, including autonomy and solitude.

Naturally, designing to balance the needs of multiple users within shared robots such that autonomy and solitude can be regulated is not an easy design task. Two people attending a conference through a single robot may want to choose a 'two-headed' telepresence robot design rather than one with a single display, or the ability to tile video feeds within a single robot display. The latter could work for viewing talks, but possibly not for social mingling given the potential social awkwardness associated with seeing two or more people on the robot's display. A person using a telepresence robot as part of a class may want a wider screen to show the entire class (similar to [17,22]), or the ability to visually obscure people who do not want to be on camera. Yet this may raise concerns around reciprocity and which remote users can be seen by local attendees. Disembodiment of remote users has been a longstanding concern with video communication systems [3] only now the system is moving throughout the environment, possibly raising more questions about who can be seen or heard without reciprocal viewing of the remote users. This is different than home environments using video chat as well as video conferencing in the workplace where the remote users are typically a small group of known individuals and cameras are constrained to a single general location [1,6,19,23]. Thus, our study extends the challenges previously raised in the video-mediated communication literature by illustrating how privacy challenges (e.g., disembodiment, solitude, autonomy) are dynamic based on movement in the remote space via a telepresence robot.

Second, personalization of the telepresence robots was clearly valuable for both shared and dedicated robots and it allowed participants to fulfill a range of social needs; this was similar to findings on avatar-based conferences [8,16,39], though we found a larger need for institutionally-focused personalization. Our results also contrast the study of CHI 2016 where personalization was said to not be a high priority because of the large venue size [31]. With a conference like CSCW, people did not want to look generic, as telepresence robots are currently designed to be. Yet what is also important is how people maintain control over their own identity and who makes identity and self-presentation decisions for them, in particular if a robot is shared by multiple people. Remote attendees could change what they looked like on camera, but unlike appearance options for avatar-based conferences [8,16,39], robot customization, and identity formation in turn, was not individually-based for shared Beams. Previous research on telepresence robots has found robot appearance to be important [25,28,33] and our work extends this knowledge to illustrate the issues around autonomy, as it relates to robot appearance, when robots are shared.

The challenge is that it can be difficult to address these needs with physical objects attached to robots, especially with large numbers of remote attendees. This challenge is exacerbated with shared robots. One possible solution could be digital augmentations where users digitally change the appearance of the robot from a remote location (e.g., change its color, adjust its height, display a picture on the back). Solutions such as this could give remote operators agency and control over their own appearance. Yet one must also be cautious with digital approaches. With digital personalization, there is risk that a person may try to change their personalization often since it easy to do so if they have remote control. In contrast, personalization that changes less over time may be easier for local attendees to identify with, which could serve to help users construct a robot-based identity longer term.

Third, we see new ways to think about knowledge sharing and social support from local attendees who are able to help remote attendees overcome navigation and awareness challenges. Here the idea of "shared" robots extends to include knowledge sharing between local and remote attendees. It also relates to understanding what one ought to do in a telepresence robot. This behavior is somewhat akin to the way parents might provide scaffolding to children using video chat, where they suggest how they should be interacting with remote family [10]. Certainly the hope is that the design of telepresence robots improves so that remote conference attendees do not experience limitations during navigation and social interactions [28,31]. Yet it is likely the case that there will always be design aspects that could be improved as social norms and needs evolve over time and so the need for contextual knowledge about one's remote environment may persist. Past studies of telepresence robots have described social support from people who are local with the robot (e.g., helping with elevators) [25,28], however, our results point to a more nuanced type of social support through local proxies. It also points to the sharing of cultural knowledge about how robots should be used in a conference setting. This is similar to how an understanding of how video-mediated communication systems should be used has evolved over time as part of workplace [7] and domestic culture [13,35].

One caveat with the approach of knowledge sharing through a backchannel relates to who is providing the social advice on what one ought to do in a telepresence robot. In our case, it was support staff. However, this may not always be desirable and it has the potential to leave the suggestion of social norms in the hands of a few, rather than being organic in the hands of many. This suggests that future design work for telepresence attendance could explore additional ways for remote attendees to communicate with a broader set of local attendees. This may provide them with the types of contextual information that our participants found valuable as well as suggestions around what may be considered socially acceptable behavior. This could then directly help to establish and evolve social norms around telepresence attendance at conferences. Naturally, having *all* local attendees participate in a backchannel would seem to be too much. Instead, a small, but diverse group, may be ideal. There are also likely other design opportunities to think about how social norms around robot usage should develop, evolve, and be shared.

Our study is limited in that we do not have data that explores the reactions from local attendees towards telepresence attendance. In contrast, our study focuses on the experience of the remote user. It is also the case that our study results do not likely generalize to larger conferences such as CHI or other events that might include several thousand attendees. We found many differences between a medium-sized conference, such as CSCW, when compared to past studies of the much-larger conference, CHI 2016 [31]. Future work should continue to explore design options for better supporting shared telepresence robots where our suggested directions show promise for such explorations.

REFERENCES

1. Morgan G. Ames, Janet Go, Joseph 'Jofish' Kaye, and Mirjana Spasojevic. 2010. Making love in the network closet: the benefits and work of family videochat. In Proceedings of the 2010 ACM Conference on Computer-Supported Cooperative Work (CSCW '10). ACM, New York, NY, USA, 145-154. DOI=10.1145/1718918.1718946
2. Victoria Bellotti and Abigail Sellen. 1993. Design for Privacy in Ubiquitous Computing Environments, Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93), Kluwer Academic Publishers, Milan, pp. 77-92.
3. Victoria Bellotti. 1998. Design for Privacy in Multimedia Computing and Communications Environments, Technology and Privacy: The New Landscape, Agre and Rotenberg eds., MIT Press, pp. 63-98.
4. Sara A. Bly, Steve R. Harrison, and Susan Irwin. 1993. Media spaces: bringing people together in a video, audio, and computing environment. Commun. ACM 36, 1 (January 1993), 28-46. DOI=http://dx.doi.org/10.1145/151233.151235
5. Michael Boyle and Saul Greenberg. 2005. The language of privacy: Learning from video media space analysis and design. ACM Trans. Comput.-Hum. Interact. 12, 2 (June 2005), 328-370. DOI=http://dx.doi.org/10.1145/1067860.1067868
6. Jed R. Brubaker, Gina Venolia, and John C. Tang. 2012. Focusing on shared experiences: moving beyond the camera in video communication. In Proceedings of the Designing Interactive Systems Conference (DIS '12). ACM, New York, NY, USA, 96-105. DOI=10.1145/2317956.2317973
7. Paul Dourish. 1993. Culture and Control in a Media Space, Proc. ECSCW '93, Springer, 125-137.
8. 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
9. Robert S. Fish, Robert E. Kraut, and Barbara L. Chalfonte. 1990. The VideoWindow system in informal communication. In Proceedings of the 1990 ACM conference on Computer-supported cooperative work (CSCW '90). ACM, New York, NY, USA, 1-11. DOI=http://dx.doi.org/10.1145/99332.99335
10. Azadeh Forghani and Carman Neustaedter. 2014. The routines and needs of grandparents and parents for grandparent-grandchild conversations over distance. In Proceedings of the 32nd annual ACM conference on Human factors in computing systems (CHI '14). ACM, New York, NY, USA, 4177-4186. DOI=10.1145/2556288.2557255
11. Erving Goffman. 1959. The Presentation of Self in Everyday Life, Doubleday.
12. Evan Golub, Brenna McNally, Becky Lewittes, Alazandra Shorter, and The Kids of Kidsteam. 2017. Life as a Robot (at CHI): Challenges, Benefits, and Prospects for Attending Conferences via Telepresence. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '17). ACM, New York, NY, USA, 758-769. DOI: https://doi.org/10.1145/3027063.3052761
13. Richard Harper, Sean Rintel, Rod Watson, and Kenton O'Hara. 2017. The 'Interrogative Gaze': making video calling and messaging 'accountable', in Harper, R., Licoppe, C. & Watson, D., Pragmatics, Special Issue: Interpersonal video communication as a site of human sociality, (2017) Vol 27: 3. September.
14. 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
15. 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.
16. 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

17. 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
18. 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
19. Tejinder K. Judge and Carman Neustaedter. 2010. Sharing conversation and sharing life: Video conferencing in the home. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10). ACM, New York, NY, USA, 655-658. DOI=10.1145/1753326.1753422
20. Tejinder K. Judge, Carman Neustaedter, and Andrew F. Kurtz. 2010. The Family Window: The design and evaluation of a domestic media space. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10). ACM, New York, NY, USA, 2361-2370. DOI=10.1145/1753326.1753682
21. Tejinder K. Judge, Carman Neustaedter, Steve Harrison, and Andrew Blose. 2011. Family Portals: Connecting families through a multifamily media space. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11). ACM, New York, NY, USA, 1205-1214. DOI=10.1145/1978942.1979122
22. 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.
23. David S. Kirk, Abigail Sellen, and Xiang Cao. 2010. Home video communication: mediating 'closeness'. In Proceedings of the 2010 ACM conference on Computer supported cooperative work (CSCW '10). ACM, New York, NY, USA, 135-144. DOI=http://dx.doi.org/10.1145/1718918.1718945
24. Kavita Krishnaswamy, Attending Conferences via Robots, http://www.washington.edu/doit/attending-conferences-robots
25. 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
26. 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
27. Gregor McEwan and Saul Greenberg. 2005. Supporting social worlds with the community bar. In Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work (GROUP '05). ACM, New York, NY, USA, 21-30. DOI=http://dx.doi.org/10.1145/1099203.1099207
28. Carman Neustaedter, Gina Venolia, Jason Procyk, and Daniel Hawkins. 2016. To Beam or Not to Beam: A Study of Remote Telepresence Attendance at an Academic Conference. In Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16). ACM, New York, NY, USA, 418-431. DOI=http://dx.doi.org/10.1145/2818048.2819922
29. 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
30. 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
31. Irene Rae and Carman Neustaedter. 2017. Robotic Telepresence at Scale. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17). ACM, New York, NY, USA, 313-324. DOI: https://doi.org/10.1145/3025453.3025855.
32. 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
33. 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.
34. Hayes Raffle, Rafael Ballagas, Glenda Revelle, Hiroshi Horii, Sean Follmer, Janet Go, Emily Reardon, Koichi Mori, Joseph Kaye, and Mirjana Spasojevic. 2010. Family story play: Reading with young children (and elmo) over a distance. In Proceedings of the SIGCHI

Conference on Human Factors in Computing Systems (CHI '10). ACM, New York, NY, USA, 1583-1592. DOI=10.1145/1753326.1753563

35. Sean Rintel, Richard Harper, and Kenton O'Hara. 2016. The Tyranny of the Everyday in Mobile Video Messaging. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). ACM, New York, NY, USA, 4781-4792.
36. N. Sadat Shami, Thomas Erickson, and Wendy Kellogg. 2011. Common Ground and Small Group Interaction in Large Virtual World Gatherings. In Proc. ECSCW.
37. Marc Smith, J. J. Cadiz, and Byron Burkhalter. 2000. Conversation trees and threaded chats. In Proceedings of the 2000 ACM conference on Computer supported cooperative work (CSCW '00). ACM, New York, NY, USA, 97-105. DOI=http://dx.doi.org/10.1145/358916.358980
38. Suitable Technologies, https://suitabletech.com/products/beam-pro
39. 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*.
40. 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
41. 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.
42. 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
43. 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
44. 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: https://doi.org/10.1145/1753326.1753482
45. Steve Whittaker, David Frohlich, and Owen Daly-Jones. 1994. Informal workplace communication: what is it like and how might we support it?. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94), Beth Adelson, Susan Dumais, and Judith Olson (Eds.). ACM, New York, NY, USA, 131-137. DOI: https://doi.org/10.1145/191666.191726
46. Lillian Yang, Carman Neustaedter, and Thecla Schiphorst. 2017. Communicating Through A Telepresence Robot: A Study of Long Distance Relationships. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17). ACM, New York, NY, USA, 3027-3033. DOI: https://doi.org/10.1145/3027063.3053240