Presenting to Local and Remote Audiences: Design and Use of the TELEP System

Gavin Jancke, Jonathan Grudin, Anoop Gupta
Microsoft Research
One Microsoft Way
Redmond, WA 98052-6399
{gavinj; jgrudin; anoop}@microsoft.com

ABSTRACT
The current generation of desktop computers and networks are bringing streaming audio and video into widespread use. A small investment allows presentations or lectures to be multicast, enabling passive viewing from offices or rooms. We surveyed experienced viewers of multicast presentations and designed a lightweight system that creates greater awareness in the presentation room of remote viewers and allows remote viewers to interact with each other and the speaker. We report on the design, use, and modification of the system, and discuss design tradeoffs.

Keywords
Tele-Presentation, Streaming Media

INTRODUCTION

The well-publicized availability of audio and video over the Internet and intranets ushers in new uses for digital technology, ranging from entertainment to distance education. Desktop computers can handle real-time audio and video. Many networks (including the Internet) require upgrading, but the technology is available. If streaming media prove to be of value, they can be delivered.

At Microsoft, as at many large corporations, presentations are now broadcast "live" over the corporate intranet, 5-10 research lectures and an equal number of more general presentations every week. Broadcasts include audio, video and slides. By clicking on web pages that list talks, employees can attend from their offices or even from home.

Clearly, there are potential benefits for remote viewers. They do not have to travel to attend the talk; if the talk is uninteresting they can quit without wasting time or risking offending a speaker or host, and if parts of the talk are uninteresting, they can multitask with other work (e.g., read email). However, there are also potential disadvantages.

First, from a speaker's perspective, remote viewing can result in fewer people attending live in the lecture room. To the extent that speakers are unaware of the remote audience, they may perceive a small live audience as lack of interest in their work. They may be offended or less motivated to deliver a good talk. It is not uncommon to hear a host reassure a speaker facing a small live audience with words to the effect of "Don't be deceived by the small audience in the room, lots of people are watching remotely."

Second, remote-viewers do not experience the ambience and subtlety of the live talk and audience. They cannot see the expressions of other audience members, whisper a question to a colleague, or direct questions to a speaker. Live audience members' questions are often inaudible to remote users unless repeated by the speaker.

Finally, a live audience member, like the speaker, may infer from a small live audience a lack of interest in the topic (which is generally of greater interest to those who traveled to the lecture room). The live experience is also diminished by the inability of remote viewers to ask questions.

Although one obvious way to eliminate these disadvantages is to disallow broadcast of talks, in this paper we explore how technology could enhance the benefits and minimize the disadvantages. We report on TELEP (short for telepresence), a system designed to provide speakers and local audiences with greater awareness of remote viewers, to provide remote viewers with a means to interact with speakers and each other, and to do this in a lightweight manner that requires little of remote viewers and almost no additional work by speakers.

TELEP is a working system currently used for seminars. In this paper we report on its design—the system components, the user interface and interaction paradigm—and the design tradeoffs we faced. We also report on audience behavior before and after the deployment of TELEP.

The paper is organized as follows. Following a section on related work, we present design goals and a TELEP system overview. The next section presents the TELEP interface and design tradeoffs. We then describe experience with broadcast presentations before and after deployment. The final sections highlight lessons learned.

RELATED WORK

For decades, videoconferencing systems (e.g., PictureTel [16]) have linked two or more rooms with bi-directional audio-video channels and split-screen displays or multiple television monitors. Our design focus is different. There may be scores of people attending remotely, each from an office. An office may or may not have a camera or microphone. The situation is much more asymmetric; consequently, the tradeoffs differ.

Distance education programs face similar challenge. Stanford University's SITN program has offered courses to students at Bay Area companies for over 25 years [18], broadcasting from a classroom via a microwave channel. The students sit at designated conference rooms within their companies to watch lectures. They can ask questions by a telephone call to the classroom.

We can confirm from personal experience teaching at Stanford that lecturer awareness of remote students is minimal. One has no idea how many are attending "live" remotely or how many have a VCR turned on to record for later viewing. Remote students do not have precise control over when to interrupt, so their questions occur as "crackling voices" in the middle of a lecturer's sentences.

TELEP is designed for a different context. Research seminars are usually given by visitors who use the system only once. Classroom instructors will use a system repeatedly, giving instructor and students more time and greater incentive to interact and establish a relationship. Remote students often have an investment in understanding the material that is equal to live students, which is often not the case in the situation we target. TELEP also differs in assuming more technology infrastructure, through which it can provide significantly greater awareness of remote viewers.

Closest to our work are systems targeted for desktop-to-desktop presentations (i.e., all viewers are remote, the speaker is in an office or studio without a local audience). Examples include Forum from Sun [4-6], Flatland from MSR [13], and commercial products such as Centra [14], NetPodium [15], and PlaceWare [17]. They may broadcast a speaker's audio-video and slides, and include additional capabilities for asking and responding to multiple-choice questions. Viewers can raise hands, ask questions via audio-channel or chat, and vote. A list of attendees is available to the speaker and viewers. Restriction to text is common, as the systems are often designed to support very large audiences and do not assume high connection bandwidth.

TELEP also provides awareness and interactivity, but the circumstances and features differ. The systems above were built for speakers with no local audience—they could devote more attention to complex software interfaces. Rich back-channels and awareness were particularly important because the speakers had no live audiences. Experiments have shown that although remote viewers like the systems, speakers are unsettled by the lack of feedback they would get from a local audience; the software interaction channels that have been tried did not fully compensate [5, 13].

In contrast, TELEP focuses on mixed live (local) and remote audiences. Because speakers must devote considerable attention to the live audience, we have kept their interface simple, requiring no keyboard use. Presence of a live audience also affects how the remote audiences can be displayed to the speaker. By assuming higher bandwidth connectivity, we can evaluate graphical representations of remote viewers (image or video) for the first time in this context. The existence of a live audience may reduce the pressure on the technology and increase the chance of success. Consider, by analogy, early radio, which started without studio audiences but introduced them because performers preferred a live audience.

In an extension to their work on Forum, Sun researchers conducted unpublished studies of "Forum Studio" with mixed live and remote audiences (John Tang, Rick Levenson, Ellen Isaacs, personal communications, 1999). Speakers stood before a podium containing a recessed computer monitor and used the Forum software to interact with remote viewers. Preliminary results contrasting local-only, remote-only, and mixed audiences showed that mixed audiences may learn less. Engaging with two audiences can distract speakers. Distant audience members may feel excluded, and a live audience may be distracted by a speaker's efforts to deal with the technology.

TELEP differs in two important ways: it does not require speakers to use technology, and our initial remote viewers have had up to two years experience passively attending lectures. TELEP can only maintain or increase their sense of inclusion, making successful reception easier to achieve.

Finally, research on supporting informal interaction (e.g., Bellcore Cruiser[8], Xerox PARC and NYNEX Portholes [1, 2], Sun Montage [11], University of Toronto CAVECAT [9], NTT Clearboard [7], and University of Calgary prototypes [3]) has addressed quite different scenarios, yet it has influenced our work.

TELEP OVERVIEW

Prior to deploying TELEP, we examined the use of the preexisting passive viewing system through observation and surveys of speakers, live audiences, and remote viewers. These data are discussed later. In this section we cover the design goals and provide an overview of the system

Design Goals and Constraints

- Presentations with a "live" audience in the lecture room and a remote audience attending from desktops.
- The lecture room interface should benefit both the speaker and the live audience.
- Medium-sized (fewer than 100) remote audiences, with access to computer but not necessarily a microphone or camera.
- Support for one-time visiting speakers. They should not have to use a keyboard. Suitable protocols for interaction should arise as naturally as possible.
- Assumption of adequate network bandwidth and computation, so it is feasible to multicast and render low-resolution video of remote viewers.

- Until proven to be reliable and acceptable, TELEP should be decoupled from pre-existing software used to watch the video of speakers and slides. A TELEP failure should not prevent people from viewing talks in the familiar non-interactive fashion.

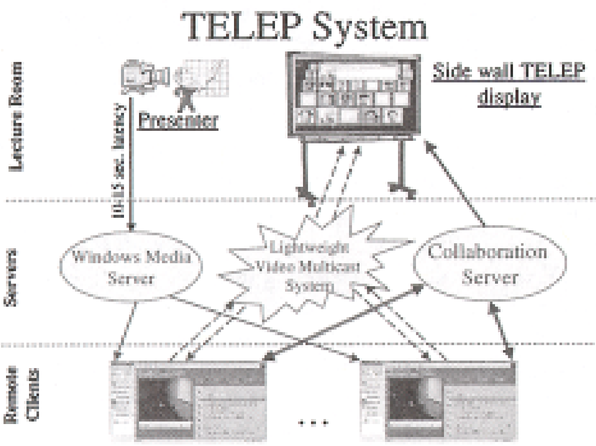

Figure 1: TELEP System Overview

TELEP System Overview

Figure 1 illustrates the interrelation of TELEP components. There are two parallel systems. The first, shown on the left, is the pre-existing system that multicasts presentations for passive viewing. Based on Microsoft Windows Media Server, it broadcasts a speaker's audio, video and slides. The corresponding display seen by remote viewers appears in the right window of Figure 2. The slides switch automatically as the speaker advances them. A key aspect is that the video-encoder, video-server, and client-side buffering, introduce a 10-15 second delay before the audio-video is received by the remote audience. This is not an issue for purely passive remote viewing, but it clearly constrains interaction between the speaker and remote audience members using TELEP.

Figure 2: Remote User Layout: TELEP window on left, web page with speaker video and slides on right

The second component is shown on the top-right in Figure 1. It produces the features in the left window of Figure 2 on remote users' desktops, and a lecture-room display as shown in Figure 3. We discuss these interfaces in detail in the next section.

Underlying the TELEP system is a collaboration server, built on MSR's Virtual World Server [12]. It communicates remote viewer actions (e.g., raising hand, voting, chat) to other remote viewers and in some cases to the lecture room display. To give remote attendees with cameras the option of using streaming video for their representation, we added a custom lightweight video multicast system. This distributes video (no audio) of remote viewers to other remote viewers and the lecture room.

The video encoder is designed to consume minimal processor cycles as it extracts and compresses live video frames from the video capture hardware. Multicast IP was chosen as an efficient network transport to distribute the video streams between remote clients and the lecture room client. The collaboration server manages the IP addresses and ports required for multiple concurrent streams.

The video stream decoder reads multicast video frames, decompresses them, and displays them in real-time. It is sufficiently lightweight that thirty or more videos can be played without saturating the processor. The decoder is adaptive: If processor use exceeds a threshold, the frame rate is decreased to avoid overwhelming the system.

DESIGN OF TELEP INTERFACE

In this section we present initial designs of interfaces for the lecture room and for remote viewers, along with considerations that affected the designs.

Figure 3: TELEP Lecture Room Display

TELEP Lecture Room Interface

In the lecture room, a dynamic, high-quality image is projected onto a large screen to the speaker's left (Figure 3). This TELEP display, visible to all in the room, is distinct from the normal projection of slides or overheads onto a screen behind the speaker. It constantly displays a representation of the remote audience.

Remote viewers can choose to appear as a live video feed from a desktop camera (for those who have one), a static digital image (for those with images in the system), a generic head and shoulders profile, or their logon alias at the bottom of the display (this is how users of the passive viewing system are represented).

An image is accompanied by a viewer's full name or first name if the name is long. The total number of remote viewers (including passive viewers) appears in the upper left. The images fill from the bottom and diminish in size in subsequent rows, giving a front-to-back impression. They range from 96x96 to 32x32 pixels, fonts from 11 to 8 pt Verdana. With this system, up to 38 images can be displayed; additional viewers can only watch. Overflow mechanisms are considered in the final section.

The black background was chosen to minimize increases in ambient light in the darkened lecture hall. However, a result is that the appearance or disappearance of images is quite noticeable to the live audience.

Remote viewers can control their representation on the lecture-room display in several ways. The border around the first author's image in the bottom row indicates that he has begun typing a question (it is yellow on the actual display). The number on the right indicates its position in the question queue. The animated keyboard beneath the image signals typing. When sent, a question appears in a large box, possibly overlaying other images until closed. Remote viewers can "raise a hand," as five viewers have done, enabling a speaker to verbally poll the entire audience. (The total of remote hands raised is not provided, though it could be if there is demand.) Viewers can change their representation (camera, still image, generic) at any time, or close TELEP and disappear from view. Of course, remote viewers can stop paying attention without visibly exiting by simply muting the sound.

Given our goal of minimizing speaker training, speakers do not directly manipulate the lecture-room interface. They use the standard forward audio-video channel to encourage questions, respond to questions when asked, and they can only verbally manage the question queue if conflicts arise.

TELEP Remote Viewer Interface

As noted above, TELEP currently runs alongside the pre-existing unidirectional presentation application, a web page consisting of two frames: one for the video of the speaker and one for slides. The slide frame can alternatively display other details: talk abstract, speaker biography, and so forth. Audio, video, and slide transitions are synchronized. Figure 2 is a typical arrangement, with these two frames in the center and right, and the TELEP window on the left.

The TELEP window, shown in detail in Figure 4, is divided into three main sections. The upper section has controls and indicators for the interactive features, system configuration and state information. The scrollable central section displays the representations chosen by the other remote attendees currently using the system. The lower section shows viewers who are preparing or waiting to send questions to the presenter. This question queue is intended to facilitate the development of social protocols to govern turn-taking.

The principal interaction features (asking questions, chatting, raising a hand) are described in the next subsection. The Configure button allows viewers to select or change representation forms. They can select live video if they have a camera. Most employees in the research division have photo images in a departmental database, which TELEP can locate. Many viewers are outside Microsoft Research, so we are developing a way for anyone to provide an image. A viewer sees a preview of their image before it is sent.

The icons to the right of the Configure button launch a TELEP feedback window, a dynamically captured snapshot of the lecture-room display, and TELEP Help respectively. The only information provided uniquely by the snapshot is the arrangement of the remote viewers.

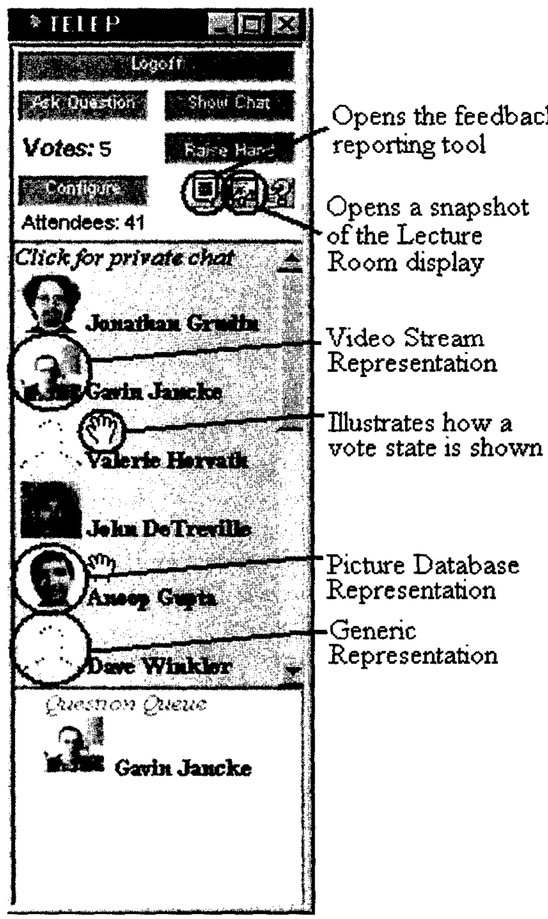

Figure 4: TELEP window for remote viewer

In the middle section, the number of remote viewers visible without scrolling would be greater if images were not displayed. The images could create more of a sense of co-presence. With viewers using photos they could also serve a minor community-building role–many remote viewers are not acquainted but could easily cross paths in the future.

Asking the Presenter a Question

When the Ask Question button is clicked, a window appears on the viewers display (Figure 5), a yellow border and question queue number immediately appears around the image in the lecture room (Figure 3), and an entry appears in the question queue on all remote displays. A prompt at the bottom of the window informs the viewer how to proceed based on their queue position and current state.

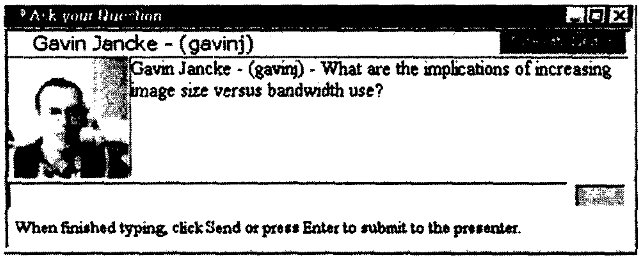

Figure 5: "Ask Question" window (questioner's view).

The remote viewer types text in the edit field at the bottom. If no other question is queued, the Send button is green and the prompt indicates that the question may be sent. Otherwise the Send button is red and the prompt indicates that another questioner is ahead in the queue.

When a question is sent, the text moves to the central area (as in Figure 5). At this point, a similar window appears on all other displays. (The lecture room display has no text entry field.) On remote displays the text entry field appears and the button to its right is labeled Reply, inviting others to respond to the question. The questioner may clarify or follow up the question, or thank the speaker, after hearing the response. Upon sending a question, a viewer is prompted to use the button in the upper right to close when done, to free the queue.

If a remote viewer sends a question when the Send button is red, it appears and the previously visible question is closed. This potentially anti-social queue-jumping feature is provided so that the discussion can move on if the previous questioner forgot to close and free the queue. This is a consequence of the minimal speaker-side interface.

We initially included more information about questioners in the window, drawn from the corporate personnel database. It was thought this might be useful for speakers, but early tests of the system indicated that speakers were not likely to use it, and it annoyed some viewers (privacy).

Remote Viewer Chat Feature

TELEP has a chat facility for use among remote viewers, not shown in the lecture room [cf. 9]. Invoked using the Chat button (Figure 4), a window appears (Figure 6). Clicking on a remote viewer's image opens another chat window for a private message. To reduce window clutter, when a message is typed and sent, the private chat window disappears and the message appears in the public chat window prefaced by "(person A to person B)" to signal that only the two can see it.

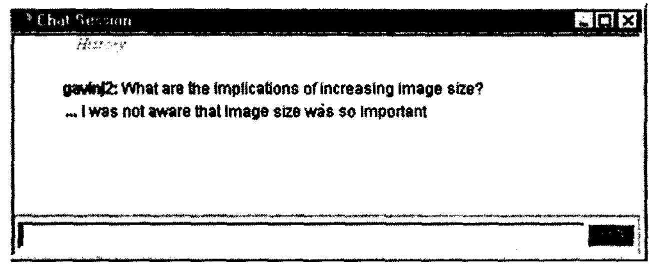

Figure 6: Chat window (remote only)

Hand Raising or Voting

A presenter may request a show of hands. As the local audience responds, remote viewers can click a button, causing hands to appear by their images (Figures 3 and 4). The vote tally is incremented. After thirty seconds, the hands disappear.

TELEP Installation, Invocation, and Maintenance

Ease of discovery and installation were considered to be critically important. For this reason, both email announcements of talks and the web calendar provide live links to the TELEP installation guide (or to TELEP once it is installed). Installation requires one button click, and subsequent modifications are automatically installed when TELEP is launched. This feature has been critical to adoption.

USER EXPERIENCE PRIOR TO TELEP

Within Microsoft Research, over 500 presentations to live audiences were multicast over the preceding two years. A distribution list of 1500 people receives talk announcements, which contain live links for viewing the presentation (and now for TELEP). Thus, many employees are fully familiar with viewing presentations live on their desktop, without interactivity. For them, the obvious comparison with TELEP is not attending in person, but between attending passively and attending with the interactivity TELEP affords. We were able to collect baseline data on how people attending in person (speakers and audience) regarded the remote viewers before and after TELEP was introduced, and how remote viewers assessed their experience before and after introduction of the system.

Initial Survey of Remote Viewing Experience

Prior to the release of TELEP, we prepared a web-based survey and emailed its URL to the presentation announcement distribution list.

This survey was designed to assess overall levels of satisfaction and problems with the pre-existing passive remote viewing system. We do not know how many of the recipients had used the system. We received 182 replies, primarily from people with an active interest in viewing presentations remotely.

The average number of presentations they reported watching remotely was 9.7; the median was 5. On average, they reported staying tuned through 54% of a presentation.

They were asked to indicate their satisfaction with the system on a 0 = Not at all to 6 = Extremely satisfied scale. The average was 3.65, slightly above the midpoint, with eight zeros and twelve 6's.

Respondents were asked how the system could be improved. The most frequent responses were requests for improved audio (in particular, for microphones that could capture live audience questions and comments), improved video, improved slide quality (if speakers do not make slides available in advance, they are simply shown in the video window and are not very readable) and greater system reliability. The most frequently requested software functionality was for remote viewer interaction with the speaker, requested by 18 respondents.

Baseline Survey of Local and Remote Experience

Next, prior to the announcement of TELEP, a paper survey was given to 11 speakers following their presentations to gauge their awareness of remote viewers and cameras, and to guess at the size of the remote audience. The local audiences ranged from 15 to 100, the remote audiences from 8 to 57, on average about 60% the local audience.

101 live audience members from eight of the talks filled out paper surveys that asked the same questions, as well as how much they had attended to the talk, daydreamed, did other work, and so forth. As noted in the introduction, remote viewing could increase multitasking or openness to distraction. We also measured live audience attrition. For four talks, we asked remote viewers to fill out a web survey that addressed the same issues. We received 31 responses.

Speakers

- Speakers were oblivious to the remote audience.

Although informed prior to talks of the ceiling-mounted cameras, nine of eleven speakers rated their awareness of remote viewers as 0 on a 0-to-6 scale, with one 1 and one 2. Ten rated the effect on their behavior at 0, with one 1. All speakers reported never looking at a camera.

- Speakers underestimated the remote audience size.

Might speakers imagine a large remote audience and be disturbed to have TELEP reveal its size? This concern appears to be unfounded: 9 of 11 speakers underestimated the remote audience size; only one greatly exaggerated it. (Actual average was 29, estimates averaged 27.)

Local Audience

- Local audiences are oblivious to remote audience.

Local audience members know that lectures are broadcast, but reported not being aware during a talk: their average rating was 0.5 on a 0 to 6 scale, with four in five rating it 0. They rated the effect on their behavior even lower at 0.2.

- They slightly underestimate remote audience size.

In only one case did an audience overestimate the remote audience size. The consensus was extremely close, but low.

- They report focusing on the talk 82% of the time.

The speaker had the viewers' attention 81.6% of the time, thinking or daydreaming 15.6%, reading or working 1.3%, and other (sleeping, looking at people, etc.) 1.6%.

Remote Audience

- They reported higher attrition than in the room.

For the live audiences measured, 65% to 90% of attendees stayed to the end. Remote viewers reported making it through only on average 37% to 67% of different talks.

- They reported greater awareness of local audience.

The average across presentations was 3.2, with behavior affected rated at 1.4. These are low, but much higher than the local audience awareness. Several specified the benefits of hearing audience questions when they were audible or repeated by the speaker, and frustration when not.

- They overestimated remote audience size and under-estimated live audience size.

Remote viewers were the only group to overestimate remote attendance. When averaged, they were close, but overestimated every talk. Their estimates of live audiences were low for all talks except one. They do see occasional camera shots of the audience, but not of the whole room.

- When they were watching, they reported focusing on talks 56% of the time.

The speaker received 55.6% of their attention, thinking or daydreaming 9.8%, reading or other work up to 32% and "other" 2.6%.

For several talks, one author attended and observed interaction. No speaker was seen to poll the audience. Many questions included clarification or follow-up, which TELEP supports but the audio delay makes difficult. Many questions or comments were longer than we would expect people to type. Occasionally a discussion broke out.

VIEWER EXPERIENCES WITH TELEP

The first formal use of TELEP was a presentation to introduce TELEP itself. It was treated as a pilot and to obtain feedback. Some of the design features described above were influenced by this feedback.

TELEP has since been in regular use. It is described briefly to speakers along with the usual A/V preparation, typically a few minutes before the presentation. The authors have not intervened appreciably, other than to observe and collect data. TELEP participation in talks has ranged from 2 to 40.

Survey data addressing awareness issues are discussed below. The interaction to date has consisted of spontaneous questions from remote viewers and a little chat among remote viewers: there has been no polling.

Questions have ranged from zero to three for a talk. To date, remote questions have not coincided or required queuing. The appearance of questions has not generally been noted by speakers, but audience members (not the authors) have pointed them out. The audience has explained the latency, but speakers have to decide how to handle it. The appearance of the "question being typed" indication forces speakers to decide whether to wait or continue—and questions have been longer than we anticipated, longer than our fixed-size window could handle on occasion.

To date, chat has been used more among remote viewers, the camera operator, and the author-observers to discuss TELEP than for content. Placing private chat (appropriately labeled) in the same window as public chat has resulted in replies to private messages almost invariably being made in that window, meaning that they were made public.

Speakers were surveyed immediately following nine talks. For 8 of these, paper surveys were distributed to the live audience; 82 were filled out. 15 remote TELEP viewers responded to a request to fill out a web survey.

During two recent talks, email was sent to 36 people using the passive system only, asking them to select among alternative explanations for why they were not using TELEP. This timely intrusion yielded a 70% response rate, including a few lengthy discussions. (Response rates to other surveys: 100% of speakers, over 50% of local audiences and TELEP users, and about 25% of remote audiences when polled after a talk had completed.)

Speaker Reactions to TELEP

- Speakers generally found TELEP interesting.

They did not seem bothered, although two wrote that some training would be useful, presumably for handling questions and the 15-second latency.

- Speakers became aware of the remote audience.

Awareness rose from 0.3 to 2.2 on the 0 to 6 scale, with no presenter indicating zero. 5 of 9 reported an effect on their behavior, but not much: the average rose from 0.1 from 0.8.

- Speakers equated the remote audience to images.

Speakers estimated the remote audience size to be roughly the maximum number of images at any one time. They overlooked the aliases of passive viewers, even when these had been explained, and did not consider remote viewer turnover. (The total number of remote viewers could be twice the number appearing at any one time.)

- Speakers equated the display with the camera.

Speakers reported looking at a camera 2.6 times (versus 0 pre-TELEP). They actually were looking at the display, which was not near a camera.

Local Audience Reactions to TELEP

- The audience generally found TELEP interesting.

Most comments were positive, but some reported being distracted by changing images, especially video.

- They became more aware of the remote audience.

Their awareness rose from 0.5 to 2.9 on the 0 to 6 scale. About half reported some effect on their behavior, with the average rising to 1.0 from 0.2.

- Their remote audience size estimates reflected the number watching at one time.

Their estimates reflected the total shown on the display when around its peak. Given the relatively high turnover, this is considerably less than the total present overall.

- Their focus on the talk may have dropped slightly.

They reported 77% of their attention on the speaker (down 5%), 14.8% daydreaming or thinking, 4.6% other work (up 4%), and 2.5% "other" (up 1%), with many attributing this last to the TELEP display. This is a possible negative effect, but it is small and may decline as familiarity rises.

Remote Viewer Reactions to TELEP

- Satisfaction reported for TELEP is quite high.

TELEP received 4.4 on the 0-6 scale, up from 3.6 for the passive viewing system. But there were few 6's and numerous suggestions for improvement.

- Their estimates of remote audience size dropped.

They appeared to base the estimate on the number of TELEP viewers, not considering the passive viewers.

- Attention to speakers dropped somewhat.

TELEP users reported attending to the speaker 44% of the time, down 12% from passive viewers. Most of this was a 350% increase in "Other" activity, which several identified as being TELEP experimentation. Future polling will determine whether or not this will drop with experience.

- Some remote viewers prefer anonymity.

Several of those still watching passively mentioned the desire to be invisible, particularly when attending in the background. "More often I'm watching it (a presentation) in the background, and so prefer to remain in the background. There's a certain symmetry to it." "I would use Telep, if my identity were only revealed when I asked a question."

LESSONS LEARNED: REDESIGNING TELEP

TELEP is in routine use, requires little maintenance, and is liked by its self-selected users. Nevertheless, many of the features were not used as expected; these lessons guided the design of the next version.

Use of live video by remote viewers for representation on the lecture-room display was not successful. The live audience found it distracting, and remote viewers with cameras did not want to be seen multitasking, on the phone, and so forth. It appears though, they may be willing to show this view to other remote viewers, and they may like to turn it on when directing a question to the speaker.

Anonymous representations are needed, perhaps as an unlabeled smaller image to the back of the display. All remote viewers should probably be represented to restore the relatively accurate estimates of remote attendance. Arguably, remote questioners should have to be identified.

A camera should be placed near the display, since speakers assume one is there. The arrival of a question should be signaled by a sound. Possibly the projection should be behind the audience rather than to its side, or in both places. Some remote viewers noted that just as local audiences can see them, they would like a camera view of the audience.

The signaling of a question being typed was disruptive: speakers did not know whether to wait or continue. We also found that remote viewers are willing to ask questions, but they do so rarely and almost never queue questions. Design should focus on simplifying the initial-question case.

Displaying public and private chat messages in one window was confusing and lead to inadvertent exposure.

For the lectures we observed, audience polling is extremely rare. More complex methods of presenting alternatives and tallying votes may be useful in other settings.

Reducing the 15-second delay in presentations reaching remote viewers to a few seconds would enable remote viewers with microphones to have an audio channel to speakers. This was a feature of the Sun Forum system. However, it is more complicated than it first seems. Questioners often prefer to catch a speaker's attention before speaking. Remote viewers may rarely use this without a more complex interface.

A New Version of TELEP

We have now released a version of TELEP designed to profit from the experiences described above. It is integrated into the viewing system, so all remote viewers participate.

Representation options have been expanded: Viewers can identify an image or use a camera snapshot as a still image. They can independently choose to show their name or be anonymous. The display can accommodate up to 60 images, moving anonymous images to the rear and moving images forward as vacancies occur to reduce visual disruption.

Questions are not foreshadowed in the lecture room until sent, at which point a sound signals their arrival. Dynamic messages guide viewers when multiple questions are being queued. Public and private chat are distinct windows. A header indicates that chat is not seen in the lecture room.

We are starting to collect data on the use of the new system. With users no longer self-selected and the basic problems addressed, we will observe how use of the system evolves.

CONCLUDING REMARKS

Although TELEP has enabled more interaction, a larger purpose was to raise mutual awareness of local and remote participants, and among remote participants. Indications are that it is succeeding in this. This could have important indirect consequences. Our initial survey found that major dissatisfactions of remote viewers included not having questions asked loudly enough or repeated by the speaker, not having slides delivered early enough to put online, not having legible overheads or whiteboard writing. As speakers and (equally importantly) their local hosts become more aware of the remote viewers, these problems will be more naturally addressed.

With use of TELEP, will more attention to remote viewers be at the expense of the local audience? Will it lead more people to attend remotely, where they are subject to more distractions? Will smaller live audiences demotivate speakers, or will more interaction with remote, often large audiences compensate? These are interesting questions that longer-term studies with TELEP will help answer.

ACKNOWLEDGMENTS

Steve White worked on early version of prortotype, Jeremy Crawford and James Crawford of MSRN helped with deployment, and Harry Chesley and Lili Cheng of the MSR Virtual Worlds Group contributed.

REFERENCES

[1] Dourish, P. and Bly, S. (1992). Portholes: Supporting awareness in a distributed work group. *Proc. CHI'92,* 541-547. ACM.

[2] Girgensohn, A., Lee, A. and Turner, T. (1999). Being in public and reciprocity: Design for Portholes and user preference. *Proc. INTERACT'99,* 459-465.

[3] Gutwin, C., Roseman, M. and Greenberg, S. (1996). A usability study of awareness widgets in a shared workspace groupware system. *Proc. CSCW'96,* 258-267. ACM.

[4] Isaacs, E.A., Morris, T., and Rodriguez, T.K. (1994). A forum for supporting interactive presentations to distributed audiences. *Proc. CSCW '94,* 405-416. ACM.

[5] Isaacs, E.A., Morris, T., Rodriguez, T.K., and Tang, J.C. (1995). A comparison of face-to-face and distributed presentations. *Proc. CHI '95,* 354-361, ACM.

[6] Isaacs, E.A., and Tang, J.C. (1997). Studying video-based collaboration in context: From small workgroups to large organizations. In K.E. Finn, A.J. Sellen & S.B. Wilbur (Eds.), *Video-Mediated Communication,* 173-197. Erlbaum.

[7] Ishii, H., Kobayashi, M., and Arita, K. (1994). Iterative design of seamless collaboration media: From Team-WorkStation to ClearBoard. *Comm. of ACM, 37,* 8, 83-97.

[8] Mantei, M.M., Baecker, R.M. Sellen, A.J., Buxton, W.A.S. and Milligan, T. (1991). Experiences in the use of a media space. *Proc. CHI'91,* 203-208.

[9] Rekimoto, J., Ayatsuka, Y., Uoi, H. & Arai, T. (1998). Adding another communication channel to reality: An experience with a chat-augmented conference. *CHI'98 Summary,* 271-272.

[10] Root, R.W. (1998). Design of a multi-media vehicle for social browsing. *Proc. CSCW'88,* 25-38. ACM.

[11] Tang, J.C and Rua, M. (1994). Montage: Providing teleproximity for distributed groups. *Proc. CHI'94,* 37-43.

[12] Vellon, M., Marple, K. Mitchell, D. and Drucker, S. 1998. The Architecture of a Distributed Virtual Worlds System. *Proc. of the 4th Conference on Object-Oriented Technologies and Systems (COOTS).* 1998.

[13] White, S.A., Gupta, A., Grudin, J., Chesley, H., Kimberly, G. and Sanocki, E. (2000). Evolving use of a system to support education at a distance. To appear in *Proc. HICSS-33.* IEEE.

[14] Centra Symposium Software. http://www.centra.com./

[15] NetPodium. http://www.netpodium.com/

[16] PictureTel Systems. http://www.picturetel.com/

[17] Placeware Conference Center. http://www.placeware.com/

[18] Stanford Instructional Television Network. http://www-sitn.stanford.edu