

**Presenting to Local and Remote
Audiences:**

Design and Use of the TELEP System

Gavin Jancke
Jonathan Grudin
Anoop Gupta

September 13, 1999

Technical Report
MSR-TR-99-71

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

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 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, many presentations are now broadcast “live” over the corporate intranet. Microsoft Research broadcasts 5-10 presentations every week, and Microsoft Technical Education broadcasts a comparable number. A broadcast consists of the audio-video and the slides of the speaker. By clicking on a web page that lists the talk, employees can attend remotely from their desktop, 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 become less motivated and not deliver as good a talk, or in extreme cases get offended. It is not uncommon to hear a host say to a speaker (e.g., when only 5 people are present for the lecture), words to the effect of “Don’t be deceived by the

small audience in the room. There truly are lots of people watching remotely.”

Second, from the remote-viewer’s perspective, they do not experience the ambience and subtlety of the live talk and audience. For example, they cannot watch the expressions of other audience members or whisper a question to a colleague. With this system they cannot interact with or direct questions to a speaker. Given the microphone setup in many such lecture rooms, unless a speaker repeats live audience members’ questions, they are often inaudible to remote users.

Finally, consider the live audience perspective. It too is unaware of the remote audience and may infer from a small live audience a lack of interest in the topic (generally of greater interest to those who traveled to the lecture room). Their experience is also diminished by the reduction in interaction due to the lack of remote viewer questions.

Although one obvious way to eliminate these disadvantages is to disallow broadcast of talks (this has been considered at Microsoft Research, and at Stanford University for classes), in this paper we explore how we may leverage technology to enhance the benefits and minimize the disadvantages. In particular, 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 other remote viewers, 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 design tradeoffs we faced. We also report on audience behavior before and after the deployment of TELEP, and what we have learned so far.

The paper is organized as follows. The next section presents related work. We then present design goals and a system overview of TELEP. The next section presents a detailed description of the TELEP interface and design tradeoffs. The following two sections present experience with broadcast presentations before and after TELEP deployment. The final two sections focus on lessons learned and concluding remarks.

RELATED WORK

Videoconferencing systems (e.g., PictureTel [13]) linking two or three sites with audio-video have been in use for decades. They allow interaction via bi-directional audio-video channels and remote audience awareness via split-screen displays or multiple television monitors. The design focus for our system 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 than traditional videoconferencing, and consequently the tradeoffs differ.

Distance education programs at universities have long faced a similar challenge. For example, Stanford University's SITN program has offered courses to students at Bay Area companies for over 25 years [15]. SITN broadcasts the audio-video of a classroom to students via a microwave channel, with a camera crew cutting between the lecturer and the blackboard or slides. The students sit at designated conference rooms within their companies to watch the lecture. Students can ask questions by a telephone call patched into the audio system of the classroom.

As is probably evident, and as we can confirm from personal experience teaching at Stanford, lecturer awareness of remote students is minimal. He or she has no idea how many are attending "live" remotely, or how many have a VCR turned on to record for later viewing. The remote students' interactions occur as "crackling voices" in the middle of a lecturer's sentences (as remote students have no precise control over when to interrupt).

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, and instructor and students have more time and a greater incentive to interact and establish a relationship. Remote students have a comparable investment in understanding the material, 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 is research and commercial product development in systems targeted for desktop-to-desktop presentations (i.e., all the viewers are remote and the speaker is without a local audience, in an office or recording studio). Examples include Forum from Sun [3-5], Flatland from MSR [10], and commercial products such as Centra [11], NetPodium [12], and PlaceWare [14]. They provide a speaker's audio-video and slides, plus additional capabilities for asking and responding to multiple-choice questions. Viewers can raise hands, ask questions via audio-channel or chat, and vote. A textual list of attendees is available to the speaker and viewers. The restriction to text is common, as some of the systems are designed to

support very large audiences and make minimal assumptions about the interconnection bandwidth.

The TELEP system also provides awareness and interactivity, but the circumstances and features differ. The systems above were built for speakers who had no local audience and could devote more attention to the complex software interfaces. Rich back-channels and awareness were particularly important because the speakers had no live audiences. Some experiments showed that although remote viewers liked the systems, speakers were unsettled by the lack of feedback they would get from a local audience; the software interaction channels did not fully compensate.

In contrast, TELEP focuses on mixed live (local) and remote audiences, a very common scenario today. Because the speaker has to devote considerable attention to the live audience, we have kept the interface simple, requiring no keyboard use by the speaker. Presence of a live audience also affects how the remote audiences are displayed in the lecture room. By assuming higher bandwidth connectivity, we can evaluate the use of visual representations of remote viewers (image or video) for the first time in this context. The fact that there is a live audience may put less pressure on the software 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 Levinson, 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.

In addition to not requiring speakers to use technology, our situation differs in that the remote viewers have had up to two years experience passively attending lectures. TELEP can only increase or hold constant their sense of inclusion.

Finally, there has been considerable research on supporting informal interaction (e.g., Bellcore Cruiser[7], Xerox PARC PortHoles [1], Sun Montage [8], University of Toronto [6], and University of Calgary [2]). These systems addressed different issues and contexts, but influenced aspects of current systems, including TELEP.

TELEP OVERVIEW

In building a system such as TELEP, there are many choices to be made. This section presents the high-level design goals and constraints we established for TELEP, followed by 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 with no prior experience with the system. 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 audio-video of speaker and slides. A TELEP failure should not prevent people from viewing talks in the familiar non-interactive fashion.

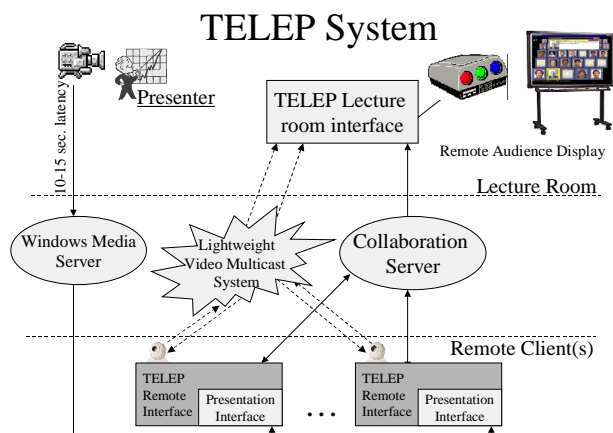


Figure 1: TELEP System Overview

TELEP System Overview

Figure 1 illustrates the TELEP system components and how they interrelate. There are two parallel systems. The first, shown on the left of the figure, is the system that has been used for two years to multicast presentations for passive viewing. Based on the Microsoft Windows Media Server, it broadcasts a speaker’s audio-video and slides to remote viewers. The display on a remote viewer’s screen appears as shown in the right window in Figure 2: a standard media player and slides that switch automatically as the speaker switches them. A key aspect of this part of the system is that the combined delay in the video-encoder, video-server, and client-side buffering introduce a delay of 10-15 seconds before the audio and video are received by the remote audience. This was not an issue for purely passive remote viewing, but will clearly constrain interaction

between the speaker and remote audience members using TELEP.

The second component is shown on the right in Figure 1. It produces the display of remote viewers and their questions in the lecture room (Figure 3) and also in a smaller window on remote viewers’ screens (to the left in Figure 2). We discuss these interfaces in detail in the next section.

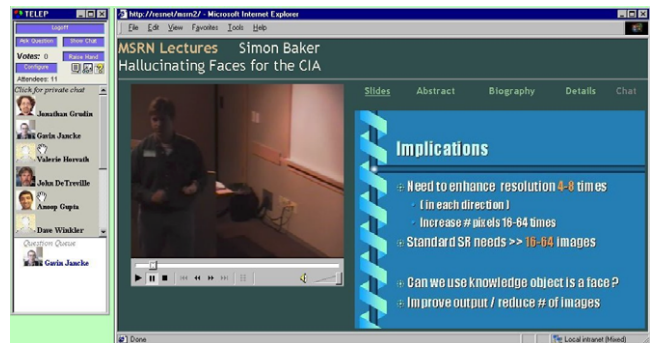


Figure 2: Remote User Layout: TELEP window on left, web page with speaker video and slides on right

Underlying the TELEP system is a collaboration server that communicates remote viewer actions (e.g., raising hand, voting, chat, etc.) to all other remote viewers and the lecture room display. The collaboration server is built on top of Microsoft Research’s Virtual World’s Server [9]. In addition, to give remote attendees with cameras the option of using streaming video for their representation we have built a custom lightweight video multicast system. This distributes the video (no audio) of remote viewers to all other remote viewers and the lecture room display.

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 is designed to read the multicast video frames, decompress them, and display them in real-time. It is written to be sufficiently lightweight that thirty or more videos can be played without saturating the processor. The decoder component is also adaptive: If processor usage exceeds a threshold the frame rate is decreased to avoid overwhelming the system.

DESIGN OF TELEP INTERFACE

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, together with the comparable study of post-deployment use of TELEP.

In this section we describe the initial design of the interface in the lecture room and the interface for remote viewers, along with the considerations that affected the design.

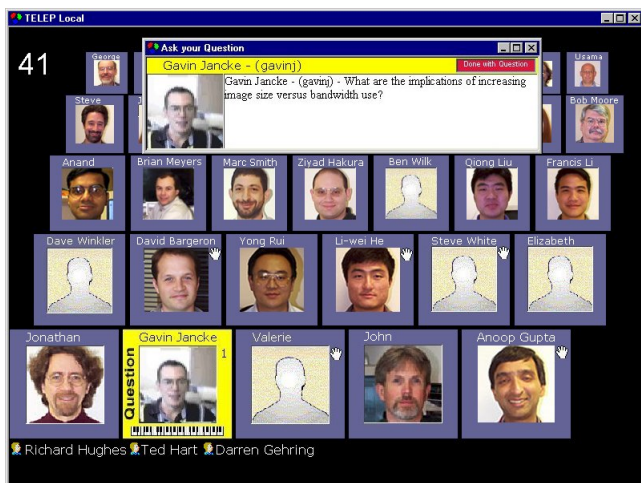


Figure 3: TELEP Lecture Room Display

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.

At individual discretion, viewers can 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 (currently representing users of the passive viewing system).

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 32x32 to 96x96 pixels, fonts from 8 to 11 pt Verdana. Currently, 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 affect their representations 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. A viewer can also change their form of representation (camera, still, generic) at any time, or close TELEP and disappear from view.

As a consequence of our goal of minimizing speaker training, speakers have no direct control of this interface. They can invite viewers to send a question or close a question box, but can only verbally manage the question queue should conflicts arise, as described below.

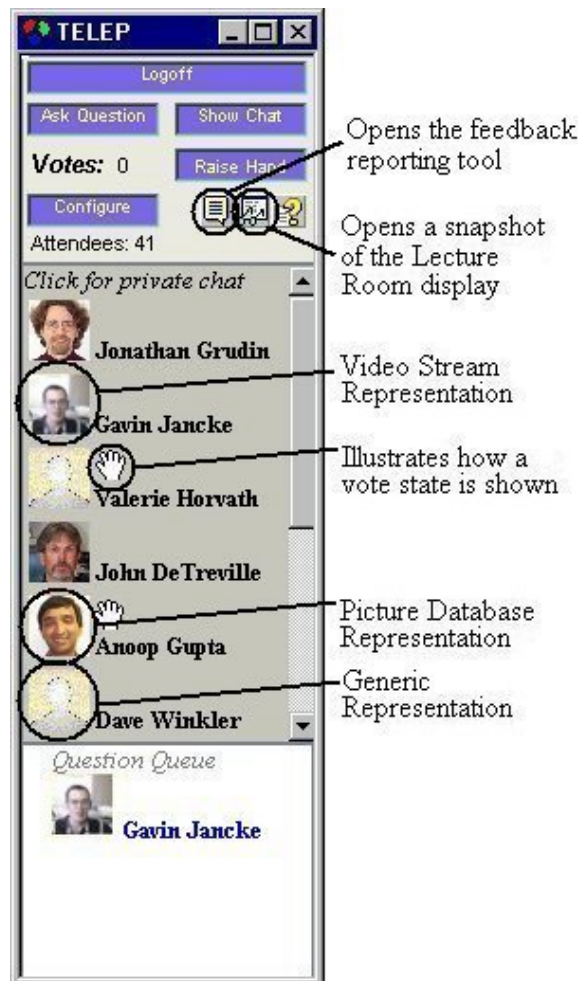


Figure 4: TELEP window for remote viewer

TELEP remote viewer interface

As noted above, TELEP currently runs alongside the pre-existing unidirectional application, a "presentation accessible web page" consisting of controls and two frames: one for the video of the speaker and one for slides. The slide frame can alternatively display other details: the host, 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 areas. The upper area has controls and indicators for the interactive features, system configuration and state information. The scrollable central area displays the representations chosen by the other remote attendees currently using the system. The lower area 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 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. In the case of photos or camera images, because many remote viewers are not acquainted but could easily cross paths in the future, it could also serve a minor community-building role.

The principal interaction features (asking questions, chatting, and 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; they are currently restricted to camera or generic images. A viewer sees a preview of their image before it is sent.

The icons to the right of the Configure button access a TELEP feedback window that invokes an email field, a window displaying a snapshot of the lecture hall display, and TELEP Help. The snapshot is used, rather than a live feed, to reduce the load on the processor and network.

Asking the Presenter a Question

When the Ask Question button is invoked, a window appears on the viewers display (Figure 5), a yellow border and question queue number 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.

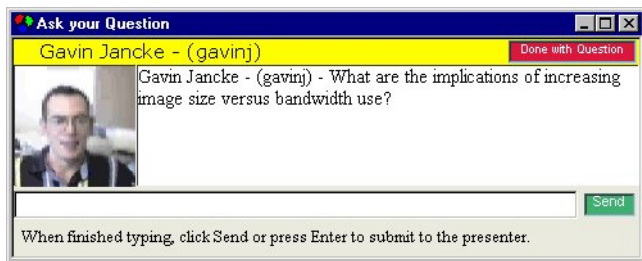


Figure 5: “Ask Question” window (questioner’s view).

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 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 the first test of the system indicated that speakers were not likely to read and use it, and it annoyed some viewers.

Remote Viewer Chat Feature

TELEP has a chat facility, not shown in the lecture room, for use among remote viewers. 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.

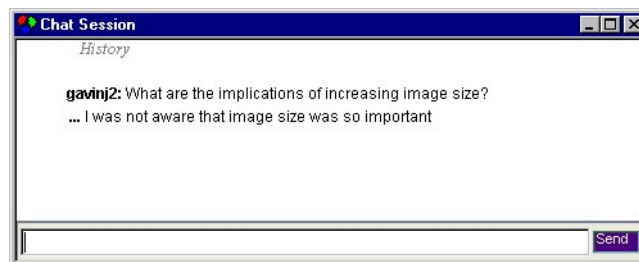


Figure 6: Chat window (remote only)

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 critical. Email talk announcements and a web calendar of televised talks provides links to TELEP (if installed) or the TELEP installation and user guide. Installation of TELEP requires one button click, and subsequent modifications automatically install when TELEP is launched.

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 passive remote viewing system. We do not know how many of the recipients had used the system. We received 182 replies. This is not a random sample, but it is a substantial number of people with an active interest in viewing presentations remotely.

The number of presentations they reported watching remotely was 9.7 on average. The median was 5, with two people estimating 100. They reported watching 54% percent of a presentation, on average.

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 presentation (many speakers do not make slides available in advance, in which case they alternate with the speaker in the video window) 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 81.6%, 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 watching on average 37% to 67% for 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 underestimated 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.

- They reported focusing on talk 56% of the time. The speaker received 55.6%, 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 remarkable 70% response rate, including a few lengthy discussions.

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 to 0.8 from 0.1.
- 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 display. But the reported effect is small and may decline as familiarity with TELEP grows.

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

TELEP is in routine use, requires little maintenance, and is liked by its users with no strong opposition. Nevertheless, many of the features have not been used as expected; these

lessons will guide the design of the version to be integrated with the projection elements of the passive system.

The lecture room video representation has not been useful. It is distracting, and remote viewers with cameras often do not want to be seen multitasking, on the phone, and so forth. On the other hand, 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 should be provided, 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.

The signaling of a question on the way should probably be dropped. This interacts with question-queue handling and the hands-free speaker goal. Given the rarity of queued questions, we should simplify for the initial-question case. If we provide speakers with a prominent "Next Question" button we could also simplify the queue handling, but at the cost of increasing system hardware and speaker training.

Private chat messages should open their own windows to eliminate the embarrassment of inadvertent exposure.

It should be possible to reduce the 15-second delay in presentations reaching remote viewers to a few seconds. This would make it possible to give remote viewers with microphones an audio channel to speakers. This was a feature of the Sun Forum system. However, it is more complicated than it seems at first. Questioners usually prefer to catch a speaker's attention before speaking, or carefully gauge the moment to interrupt. Remote viewers are unlikely to use this without a more complex interface.

Use will scale up when the passive viewing system is integrated, and other issues will arise. The screen real estate taken by images of other remote viewers may seem a poor tradeoff against seeing more names. The presentation room view will have to handle more than 38, perhaps by scaling down all images or the late arrivals in the back rows.

CONCLUDING REMARKS

TELEP has enabled more interaction, but the larger purpose is to raise mutual awareness of local and remote participants, and perhaps among remote participants. Indications are that it has succeeded 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.

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?

More casual participation will be more common as the amount of available online talks grows. Just as television can now provide scores of channels, computers could enable us to access thousands of presentations, internal and external to our workplaces.

ACKNOWLEDGMENTS

Steve White helped plan this project, Jeremy Crawford and James Crawford of MSRN helped with deployment, and Harry Chesley and Lili Cheng of the MSR Virtual Worlds Group contributed. We thank all of them.

REFERENCES

- [1] Dourish, P. and Bly, S. (1992). Portholes: Supporting awareness in a distributed work group. *Proc. CHI'92*, 541-547. ACM.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] Root, R.W. (1998). Design of a multi-media vehicle for social browsing. *Proc. CSCW'88*, 25-38. ACM.
- [8] Tang, J.C and Rua, M. (1994). Montage: Providing teleproximity for distributed groups. *Proc. CHI'94*, 37-43.
- [9] 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.
- [10] 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.
- [11] Centra Symposium Software. <http://www.centra.com/>
- [12] NetPodium. <http://www.netpodium.com/>
- [13] PictureTel Systems. <http://www.picturetel.com/>
- [14] Placeware Conference Center. <http://www.placeware.com/>
- [15] Stanford Instructional Television Network. <http://www-sitn.stanford.edu/>