

- Dourish, P. and S. Bly, Portholes: Supporting awareness in a distributed work group, *Proceedings of the Conference on Computer Human Interaction (CHI) '92*, Monterey, CA, May 1992, pp. 541-547.
- Gibbons, J.F., W.R. Kincheloe, and K.S. Down, Tutored videotape instruction: A new use of electronics media in education, *Science*, Vol 195, 1977, pp. 1139-1146.
- Minneman, S. L. and S. A. Bly, Managing a trois: a study of a multi-user drawing tool in distributed design work, *Proceedings of the Conference on Computer Human Interaction (CHI) '91*, New Orleans, LA, April/May 1991, pp. 217-224.
- Moore, D.S., The place of video in new styles of teaching and learning statistics, *The American Statistician*, Vol. 47(3), 1993, pp. 172-176.
- Root, R. W., Design of a multi-media vehicle for social browsing, *Proceedings of the Conference on Computer-Supported Cooperative Work*, Portland, OR, September 1988, pp. 25-38.
- Stefik, M., D. G. Bobrow, G. Foster, S. Lanning, and D. Tatar, WYSIWIS revised: Early experiences with multiuser interfaces, *Proceedings of the Conference on Computer-Supported Cooperative Work*, Austin, TX, December 1986, pp. 276-290. Reprinted: *ACM Transactions on Office Information Systems*, Vol. 5, No. 2, April 1987, pp. 147-167.
- Stefik, M., G. Foster, D. G. Bobrow, K. Kahn, S. Lanning, and L. Suchman, Beyond the chalkboard: Computer support for collaboration and problem solving in meetings, *Communications of the ACM*, Vol. 30, No. 1, January 1987, pp. 32-47. Reprinted: *Computer-Supported Cooperative Work: A Book of Readings*, Irene Greif (Ed.), San Mateo, CA: Morgan Kaufmann Publishers, Inc., 1988, pp. 335-366.
- Tang, J. C. and E. A. Isaacs, Why do users like video? Studies of multimedia-supported collaboration, *Computer Supported Cooperative Work: An International Journal*, Vol. 1, Issue 3, 1993, pp. 163-196.
- Tang, J.C. and L.J. Leifer, An Observational Methodology for Studying Group Design Activity, *Research in Engineering Design*, 1991, Vol. 2, No. 4., pp. 209-219.

is the most important factor, for example, reduced feedback may be an acceptable price to pay. And over time, we expect that new conventions will develop about proper behavior in this kind of setting, which may help enhance both speakers' and audiences' experiences with distributed presentations.

It is also worth exploring ways to expand the range of situations in which distributed presentations can be effective. Indeed, we have considered a number of approaches to improving the interactivity possible over Forum. One approach is to change the nature of the environment by actually including a local audience. Another is to improve upon the technology and design of the application.

Future directions

We are planning to build a specially equipped studio that will accommodate up to 35 local audience members. A remote audience would see video of the speaker and the local audience, who in turn would see, on a screen at the front of the room, the slides and the remote audience feedback. Attention will be paid to providing adequate studio audio capabilities so that remote audience members can hear local participants without requiring the latter to make a special effort to be heard. We are also exploring the use of input devices that would enable the speaker to walk around and interact with the screen rather than standing in front of a monitor using a mouse.

During the technical course experiment, we simulated such a studio and found that it raised a number of interesting issues. For one, we found that the students who received the talk in the "mixed environment" (half over Forum and half in the classroom) performed worse on a final exam than did either the classroom-only or Forum-only students. We believe this result was caused by an insufficient integration between the two audiences. The speaker found it difficult to attend to both audiences, partly because it was harder for the remote audience to make their presence felt and partly because certain aspects of the interface accentuated the difference between the groups. For example, when the speaker asked a yes/no question, the remote audience could answer yes or no at once, but the local audience had to be prompted for each response. Creating such a studio will force us to reconsider the design of Forum to match a more heterogeneous environment.

Aside from building a studio, we are considering other potential technology and design enhancements to Forum. Since the feedback from the audience to the speaker was relatively impoverished, one obvious enhancement would be to provide video of the audience. Although the large group size would preclude showing everyone in the audience, one could experiment with a matrix of small, slowly updating images, similar to Port-holes [Dourish and Bly, 1992]. Currently, most users do not have machines equipped with cameras, and even if

they did, transmitting that much video data would probably overload most networks.

However, it is feasible to experiment with audio without waiting for technology advancements. As noted, because of technical problems with echo and varying audio levels, we compromised by forcing audience members to hold down a button when speaking. During that time, the speaker's audio was reduced, which lowered the amount of sound feeding back into the microphone. It would be worthwhile to spend more time improving the audio so that spontaneous reactions from willing audience members (such as applause and laughter) could be mixed and transmitted to the group.

In the absence of audio improvements, we have also considered providing some interface widget to enable the remote audience to indicate amusement and appreciation. We are reluctant to simply put Applause and Laughter buttons on the interface because they would require users to take deliberate, artificial actions to indicate emotions that are usually conveyed spontaneously and naturally. However, we think it is worth experimenting with different techniques because this information is so valuable to the speaker.

Other options for increasing participant interaction and awareness are likely to become available once computer and network technology can better support the bandwidth requirements of distributed multimedia applications. Although we should not expect to eliminate the need for face-to-face talks, it should be possible to broaden the range of situations in which computers can support distributed presentations that are effective and satisfying to both speakers and audiences.

ACKNOWLEDGEMENTS

We would like to acknowledge the contributions of the other members of the Collaborative Computing (COCO) group within SunSoft. We would especially like to thank Darrell Sano, who helped design the user interface, Don Hoffman and Michael Speer, who provided the video program used by Forum, and David Gedye, who was instrumental in initiating the Forum project. We thank SERA Learning Technologies as well as Rob Harris of SunU and Amy Pearl for their efforts in running the experiment on training over Forum. We appreciate the support we received from Jim Kubon and Sun's system support staff as we experimented with network resources. And finally, we especially thank the many people within Sun who gave or attended weekly Forum broadcasts and gave us their candid feedback.

REFERENCES

Bly, S. A., S. R. Harrison, and S. Irwin, Media spaces: Bringing people together in a video, Audio, and Computing Environment, *Communications of the ACM*, Vol. 36, No. 1, January 1993, pp. 28-47.

audience member could send suggestions to that person, rather than bothering the rest of the audience with their advice. Ultimately, however, the best thing we could do was to modify the interface to make it easier to use, which we did. The fewer problems users had learning to use the interface, the less they needed troubleshooting help.

CONCLUSION

Design recommendations

The observations discussed so far emerged during our highly iterative process of developing Forum. Each time we tested a revised design with users, we learned more about how certain design decisions affected participants' ability to collaborate during a talk. What emerged were the following principles, which we believe apply generally to distributed, one-to-many large group applications. Designing by these principles can improve the ability of participants to collaborate in the joint activity being supported.

Because participants are distributed, it is important to:

- enable them to express as much feedback as possible,
- make feedback mechanisms as natural as possible.

It is helpful for tools for large group interaction to:

- provide as many cues as possible about the size of the group,
- provide capabilities to enable a sense of awareness among group members,
- focus on the application's ease of use to minimize the need for joint troubleshooting, which can be difficult for large groups to manage.

Designs of tools for participants with asymmetric roles should:

- have UIs for the different roles that are as similar as possible,
- provide cues to indicate which aspects of participants' interfaces are similar as well as different,
- indicate what other participants see whenever they don't see the same thing.

A role for distributed presentations

When we began designing a tool for distributed presentations, we recognized the importance of speaker-audience and audience-audience interactivity. However, the experience of building and testing Forum has made it much more vivid to us how critically the success of a presentation can depend on the speaker's ability to gauge the audience's ongoing reactions and on the audience's ability to build upon each other's participation. Although Forum was quite successful in enabling inter-

esting and rewarding distributed presentations, its effectiveness seemed to be highly dependent on the particular topic and speaker.

In particular, Forum was most effective either when the topic sparked a great deal of audience participation or when the purpose of the talk was simply to provide information to a wide audience. When the topic was controversial, audience members tended to contribute frequently, which not only gave the speaker feedback about the audience's reaction but also prompted audience members to react to each other's contributions. These types of talks tended to be very satisfying for both the speaker and the audience. During lecture-style talks, speakers communicated a message to an audience that included people who might not be able or willing to attend a face-to-face talk. Forum enabled the speaker to reach a wider audience than they would at a face-to-face talk, and it enabled the audience to learn information they would have missed had they been required to invest more time and effort to attend.

From a speaker's point of view, Forum was clearly less effective for talks that depended on a certain degree of interaction but were not controversial enough to generate strong audience participation. The lack of feedback made speakers feel like they were "in a giant black box," as one speaker put it, and they were likely to blame themselves for not provoking interest. It is interesting that even in these cases, audience members were far less likely to be dissatisfied, however, because the cost of attending was so low and they usually learned at least some useful information even if it did not inspire them to ask questions.

On the other hand, from the audience's point of view, Forum's environment was in some ways less limited than a face-to-face setting, primarily because the politeness rules for the audience were more relaxed. Over Forum, it was no longer rude to do other work while listening, to send notes to other audience members, or even to leave in the middle of a presentation. In effect, Forum provided audiences with the opportunity to get the most out of the time spent during a presentation, especially when they were only mildly interested in the topic. Of course this change affected the atmosphere of presentations, again putting more pressure on speakers to reach out for interaction.

The fact that Forum is a more limited environment than face-to-face settings does not mean that this kind of technology should not be pursued. In fact, we think Forum represents a step beyond many of the distance learning technologies because it provides support for real-time, spontaneous interaction among the speaker and audience members. Still, our findings indicate that distributed presentations must be treated as a different kind of presentation environment. Speakers must consider whether their goals in giving a talk can be met in a less interactive setting. In cases where wide distribution

to make sure that his expressions and gestures came off as he intended. He said he even used it to assure himself that he looked presentable and that he didn't have smudges on his face.

Limitations of creating shared environment

Despite efforts in the design of Forum, there were ways in which the asymmetry of the situation and the user interfaces caused collaboration problems. In particular, the asymmetry made it more difficult for the participants to jointly learn to use Forum in a number of ways. And it was particularly difficult to help others troubleshoot problems.

More difficult to learn. The fact that there was any asymmetry at all between the speaker's and audience's UIs created some doubt in users' minds as to which aspects were shared and which were not. It was not uncommon for new users to wonder aloud whether the speaker or other audience members could see what they were doing.

For example, when first using the Poll, some users wondered whether the speaker could tell who was voting on which option. As you can see from Figures 1 and 4, the Poll meter is the same on everyone's interface, but audience members could not be certain of this. This uncertainty could affect their behavior, since they might vote differently (or not at all) if they thought their vote was not anonymous. Some users wondered whether just the speaker or the entire audience could see their written questions once they sent them in, even though they never saw anyone else's written questions. A few times, audience members sent written comments that were never acknowledged by the speaker. Users were disconcerted by this because they didn't know whether they had been ignored or the technology hadn't worked. As one person put it, "My comments have disappeared into a black hole as far as I can tell."

It was interesting that in most of the cases where audience members were confused about what others saw, their view was equivalent to everyone else's. This is an interesting problem because it indicates that, once the notion of WYSIWIS is broken, users doubt whether each aspect of the interface is shared among participants. It is not enough to provide similar views; one must also provide cues to indicate which aspects of users' views are shared.

On the other hand, there was at least one way in which the asymmetry was useful for learning. By participating as an audience member in Forum talks, users could potentially learn how to give talks that were well adapted to Forum's environment. Some speakers who had been audience members realized the effectiveness of the Poll, for instance, and used it regularly in their talks. One speaker said he had been annoyed by other speakers who had covered up parts of their slides with

the pointer, so he tried to avoid doing so. It is interesting, however, that even experienced audience members, who knew first hand how useful it is to split their attention while watching a talk, still found it disconcerting as speakers when the audience didn't appear to be paying full attention. Apparently, this knowledge didn't override their need for audience feedback.

More difficult to troubleshoot. One of the more frustrating aspects about using Forum was trying to help someone who was having trouble. Symmetric applications for small group use can manage this problem by enabling everyone to see what everyone else is doing. A shared drawing or editing tool, for instance, might show everyone's cursor when users are over UI controls as well as when they are actually drawing or writing [Minneman and Bly, 1991]. Doing so enables everyone to see when someone is having trouble manipulating the UI, and they can offer to step the person through the task. However, if not everyone shares the same UI, this is not possible. Nor is it practical to show the activities of a few hundred users. Also, it is less appropriate in a large group setting to stop focusing on the subject matter to help one person carry out a task.

The most common problem that arose in using Forum was that an audience member wasn't able to speak. An earlier version of the speaking interface turned out to be relatively difficult for audience members to use, and this design caused certain common errors. Other times, audience members' attempts to speak were foiled by technical problems (microphone wasn't plugged in, battery was dead) or by buggy software. (After all, these tests were designed in part to help debug the software.)

The speaking interface also happened to be the one part of the UI that differed most between the speaker and the audience. As mentioned, the speaker did not have to take any action to speak, but the audience members had to hold down a button to speak, and they were expected (but not required) to get in a queue to indicate they wanted the floor. When a user had been called on by the speaker and then didn't reply, it was very difficult for the speaker to help them. They couldn't step them through the process because they couldn't see the audience interface. Even if they could have seen it, they still wouldn't be able to see what that particular user was doing and what problems they were having. There was no way to know they were having trouble manipulating the UI or that their microphone wasn't turned on.

After observing participants struggling with troubleshooting problems, we felt there were only limited things we could do to reduce its occurrence. As a start, we familiarized speakers with the audience interface if they hadn't experienced it as an audience member. Once we added the message capability, we found that it was the most effective way for audience members to give each other help. If someone's audio level was set too high or they appeared to be trying to speak, another

tive for interacting with others. As one user put it, “I can tell another audience member that I agree/disagree with what they just said. Or point out to them where they can get more information, etc.” Still, the majority of audience members had yet to discover the messages. We expect that the message capability is likely to increase in use over time as more people discover it on their own or receive a message from someone who has used them before.

Limitations of awareness tools

Despite our attempts to create a sense of shared environment, it was clear that speakers and audiences did not feel as much like they had shared a common experience as they do after a face-to-face presentation. Several people commented that their understanding of other audience members largely depended on the extent to which the speaker involved the audience in their presentations. Talks in which speakers asked many poll questions and encouraged audience questions left the audience with a much greater sense of community than those who gave more lecture-style presentations. Of course, this is also true of face-to-face presentations, but audience members have somewhat more control over their degree of interaction with the people around them. They are also free to look around the room to see how people are reacting to particular parts of the talk and to chat with others as they head back to their offices.

In addition, we noticed that participants sometimes behaved in ways that were more appropriate for smaller groups, apparently because the audience list was only a subtle indication of the group size.¹ For example, sometimes people raised their hands in the middle of a point and when called on said, “Oh, I was just testing the interface.” Another time, when the speaker had already moved on after asking a poll question, an audience member asked the speaker to repeat the question because he had been distracted when it was asked. In fact, we found that the tone of the talk seemed to be influenced more by the type of talk than the size of the group. Lecture-style presentations with a small audience tended to be more formal than discussion-style presentations with a large audience.

Finally, we had not anticipated how frequently people were going to watch in groups around one person’s desktop. Initially, the prototype ran on only one version of the operating system, so people who were running a different OS had to share a workstation to see a presentation. As a result, when questions were asked by people who were not at their own workstation, the wrong name and picture appeared. When those people sent in written questions, their signature didn’t match the name of the sender. The fact that people could watch from another workstation also foiled our attempt to rule out any-

1. People may also have felt freer to experiment because some of the talks were considered tests.

mous participation. Although we are unaware of any such cases, it was possible for someone to unwittingly say something offensive to someone not listed as an attendee.

ASYMMETRIC ROLES

In many synchronous collaborative activities supported by CSCW tools, the participants play equivalent roles. However, the nature of presentations is that the speaker and the audience play different roles. This asymmetry means that the tool cannot be designed strictly according to the “What You See Is What I See” (WYSIWIS) principle [Stefik, 1986, 1987]. Tools designed using this principle help users interpret what is going on in an interaction and enhance participants’ ability to jointly learn how to use the application. Because tools designed to support presentations must provide different capabilities for the speaker and the audience, it is important to create environments that are as similar as possible.

Effectiveness of creating shared environment

In designing Forum, we attempted to create a shared environment by making the overall layout of the controls as similar as possible and using the same controls any time the speaker and the audience shared the same functionality. For example, the basic layout of the poll and audience’s speaking interface is almost identical, and the UI elements for these components are the same. Both the speaker and the audience see the same graph develop as the poll results come in. Everyone sees the same **Hand Raised** queue as well as the name and icon of the audience member who is speaking.

In those cases where the speaker and audience had different controls, we tried to suggest the other interface whenever we could. For example, the speaker does not have to hold down any controls to speak but the audience does.² Since we were using a metaphor of raising one’s hand to request the floor, we used the label **Hand Raised** for the question queue, and **Lower Hand** for the button that enables them to remove someone from the list. This encouraged them to use terminology that made sense to the audience when referring to these actions.

We also provided a video preview for speakers so that they would know what the audience was seeing. They used the preview to see when they were on the air, to help them frame themselves and to make sure their gestures were seen. One speaker said he also found that having the video helped him relax as he gave his presentation because, unlike at a face-to-face talk, he was able

2. We had originally hoped to provide continuous voice-activated audio of the audience, but we encountered certain technical problems. Not only were we unable to eliminate audio echo but it was not obvious how to determine which sounds from the audience were appropriate to transmit.

as uncomfortable for them to not have a smooth way of expressing amusement at a joke or appreciation at the end of the talk. In fact, we observed a number of cases when audience members spontaneously thanked the speaker at the end of a talk even though they weren't transmitting their voice to the speaker.

Although the audience's apparent unresponsiveness was uncomfortable for speakers, we found that many had received feedback from audience members during the week after their talk. Many said acquaintances and even strangers approached them to tell them they had enjoyed the talk or to ask follow up questions. These encounters were not as gratifying as getting an immediate response to a presentation, but most said it did enhance their satisfaction with the experience.

Another problem for collaboration was that, after audience members asked a question, they could not indicate their satisfaction with the speaker's response. We noticed occasions when the speaker spent a long time addressing an issue, long after the questioner felt their question had been answered. Some speakers handled this problem by asking, "Does that answer your question?" Others just said they had felt frustrated that they couldn't tell if they had covered the issue the questioner had in mind. At the other end, some audience members mentioned they wished they could have indicated in some subtle way when they were satisfied with the answer. Of course, this same problem can occur in very large presentations when a speaker may not be able to see who asked a question and therefore cannot tell when they have provided an adequate answer.

Behaviors interpreted differently. We noticed a few interesting cases in which audiences behaved over Forum similarly to the way they do at a face-to-face talk, but their behavior was interpreted differently nonetheless. The use of the poll was one such example.

While using Forum, speakers sometimes asked the audience to "vote" to indicate their opinion on a matter, just as they sometimes ask local audiences to raise their hands for the same purpose. In such situations at a face-to-face talk, it is common for not everyone to express an opinion, and for some people to wait until they see how others are voting before they respond. Although these same behaviors often happened over Forum, speakers reported that they found the audience's behavior unresponsive. Some explicitly prompted those not voting to express an opinion, and others repeated the question when the response wasn't immediate.

Part of this difference in interpretation may be due to capabilities provided by the Forum interface. Because it is trivial in Forum to add the number of votes for each option and compare it with the total attendance, it is obvious when not everyone votes. In a face-to-face talk, audiences must first vote on one option and then on the second, so it is not as clear when people abstain. Also,

speakers found it less comfortable waiting for the votes to come in during a Forum talk because they could not look around the room to see how people were reacting.

This effect was particularly interesting because it stemmed from what we considered to be added functionality compared with a face-to-face talk. It can be nice to know exactly how many people are attending and how many are voting for an option. It is interesting that when concrete information was provided, people were more likely to interpret rather ordinary behavior as uncooperative.

LARGE GROUP SIZE

Part of the appeal of technology to support distributed presentations is that it enables potentially large groups of people to participate. A particular audience member's perception of a presentation can vary little whether the audience has a dozen or several hundred members. This effect makes large group presentations feel more intimate, but it also reduces audience members' awareness of each other. This awareness helps audiences understand how to behave appropriately and it provides context to interpret the impact of a talk. It is important for tools designed to support large group participation to enable some sense of awareness of the participants.

Effectiveness of awareness tools

Our initial attempt to create this awareness was providing a list of attendee's names in the Audience window, which users could click on to see the person's picture and some basic information about them. Speakers and audiences browsed through the list before and during the talk to see who was attending. They often clicked on each name to see the pictures of friends and to connect a face with an unfamiliar name. A few times, audience members visited other participants' offices during a talk after seeing their name on the list. Some spoke about the talk to participants they encountered several days after it was over. The list was also important for helping speakers target their talk to the audience. We decided not to allow users to watch anonymously, because we felt participants should be given the opportunity to adjust their comments to the particular people in the audience.

Still, the list alone was fairly limited, so we added the ability to send short text messages to other participants, which enabled people to "pass notes" during the talk without being rude. During the talks that included this facility, about a quarter of those in an audience sent a message, sometimes to make a comment about the talk, other times to discuss other plans or even to crack jokes. Messages were quite short and the exchanges consisted of about four turns, sometimes spread out over an hour talk. Participants usually, but not always, alternated turns.

Audience members who had used the messages indicated in the surveys that they found the capability effective.

effective for helping them focus their attention, especially when they were first getting used to the interface. In addition, just displaying the slides on everyone's desktop eliminated visibility problems that often occur in face-to-face talks with large audiences. Audience members said that if they missed part of the talk, they looked back in the slides to fill themselves in. Some said that before asking a question, they checked previous slides to make sure their question hadn't already been covered.

Limitations of interaction tools

Although the mechanisms for interaction were effective, it is clear that they did not provide for the full range of communication. Interactions were reduced in part because of limitations in the technology and design, but also because the distributed environment changed the affordance of certain behaviors. In particular, audiences found it difficult to avoid splitting their attention between the talk and other work activities available in their offices.

Audience's gain is speaker's loss. The most dramatic way in which Forum talks differed from face-to-face ones was that the audience had more freedom to vary their level of participation. But that freedom often came at the expense of the speaker's satisfaction with the experience. In Forum's one-to-many setting, information about audience activity was much sparser than that about the speaker. As a result, the social pressures that regulate the behavior of co-located audiences were greatly diminished, which allowed audience members to behave in ways that are not as acceptable among a local audience.

For example, it was very common for audience members to pay attention to the talks during the more interesting periods, but then read email or focus on other work during slower periods. Nearly every user who responded to surveys mentioned that Forum's greatest advantage was that it enabled them to get other work done during the less interesting parts of a talk. Audience members also chose to take certain interruptions, such as phone calls or visitors, that they would have missed had they been attending the talk in a meeting room. This freedom could also be a burden, however; many users had trouble ignoring distractions in their office even when they were very interested in the talk.

Because the amount of effort required to attend a talk was so much lower, some people signed up for talks that were only potentially of interest, knowing they could decide to leave without calling attention to themselves.¹

1. Certain people who are well-known around the company said they particularly appreciated having the ability to leave discreetly because they are especially vulnerable to insulting speakers if they leave a talk. Some said they regret having to avoid talks of marginal interest for this reason.

In fact, because audience members could scroll through the slides independently of the speaker, some people scanned through the slides when the talk began to decide whether they wanted to stay for the rest of the presentation. And finally, sometimes people watched talks in small groups, which enabled them to make comments about the content of the talk, or even about the speaker, without being rude.

From the speakers' point of view, however, the fact that some of the audience was paying less attention than usual created situations that detracted from their own experience of giving the talk. For example, after asking the audience to vote on an issue, some speakers felt disappointed, if not insulted, when only a portion of the audience responded. They had no way of knowing whether some audience members were on the phone or away from their desks. They certainly had no cues indicating that some of the audience members were less interested in the topic than usual because the cost of attending was so low.

Some speakers found it disturbing when audience members asked them about an issue they had already addressed, clearly indicating they hadn't been paying attention. In fact, when the technical course was taught with a mixed audience, some local audience members said the people attending over Forum seemed less intelligent because they more often asked questions that had already been addressed.

On the other hand, we were interested to find that certain speakers who tended to get anxious giving talks in front of a crowd said they found it more relaxing to give a talk over Forum. Sitting in a comfortable office talking to just a camera and perhaps one other person in the room enabled them to treat the talk as a more informal presentation. They said they were able to concentrate on their material rather than worrying about how they were being perceived. Still, dynamic speakers who depend heavily on audience feedback felt quite constrained by the technology.

Feedback reduced. As expected, there were a number of ways in which the limited bandwidth reduced the amount of feedback in ways that detracted from both the audience and the speaker's experience.

Most disconcerting among these were the lack of laughter or applause. We noticed that speakers tended to pause after saying something amusing as if waiting for a response, but then moved on awkwardly when they got none. They seemed to find it particularly uncomfortable to end their talks without any applause to provide closure. One speaker ended his talk by thanking the audience members for their participation, expressing his hope that they had learned something from the talk, informing them how they could reach him after the talk, and even wishing them a good afternoon before finally signing off. Certain audience members said it was just

users, but over time more non-technical speakers and audience members used the system.

In addition, we experimented with an early version of Forum to give a course on a technical topic through SunU, the company's in-house training organization. That event was a structured experiment in which the same three-hour course was given three times, once to a live audience, once to an audience watching through Forum, and once to a mixed audience: some in the room with the speaker and some watching over Forum.

During every talk and during the technical course, we videotaped the speaker and at least one audience member, who was asked to watch the talk just as they normally would. One person reviewed the tapes in detail and the group met to look at highlights that brought out interesting aspects of Forum's usage. (See Tang, 1991, for a further description of this analysis methodology.) In addition, questionnaires were e-mailed to audience members four times over the 10 month period, yielding 97 responses. The speakers were usually interviewed after their talks to learn more about their reactions. Using these data, we tried to understand how effectively Forum overcame the obstacles introduced by the distributed nature of the audience, the large group size and the asymmetric roles of the participants.

DISTRIBUTED PARTICIPANTS

It is a great convenience to participate in a talk from one's own workstation. In fact, users' overwhelming response to Forum was that they loved the convenience. A typical survey response was as follows: "I appreciate the fact that I can listen to a talk while I continue to work in my office. It is nice that I don't have to track down a conference room in some foreign building. It is nice that I don't have to wait in the conference room for the talk to start." However, this convenience comes at the cost of reduced interactivity unless mechanisms are provided to overcome the barrier of distance. The interaction mechanisms we provided in Forum did enable a certain degree of audience participation, which most users said greatly enhanced the experience of giving and attending a talk. However, the distributed nature of the audience had certain detrimental effects that were not overcome by these interaction tools.

Effectiveness of interaction tools

Forum enabled audiences to see and talk to the speaker, which are a natural part of face-to-face presentations, and it also provided more technology-driven interaction tools, namely the poll and the written comments. The video and audio were the most effective elements for enabling two-way interaction, but the technology-based mechanisms played surprisingly important roles as well.

The audience asked spoken questions to varying degrees in every talk. Some speakers used Forum to present their prepared materials and then took questions from the

audience at the end, and others took questions as they arose during the presentation. Nearly every speaker said the voice interaction with the audience was very important to give them a sense of the audience's interest level and understanding, and to provide a sense that their efforts had been worthwhile. Speakers also said they liked seeing a picture of the person speaking, especially when they were not familiar with the person by name.

From the audience's point of view, the video was very important for creating a connection with the speaker. When the speaker looked into the camera, it appeared to each audience member as if she were speaking directly to them. As one user put it, "it feels more intimate than a live talk with a large audience — just me and the presenter and maybe my officemate on the other side of the room."

Written questions were a useful feature because they enabled both speakers and audience members more flexibility in handling questions. Some speakers used the more traditional method of waiting until the end to take written questions, but others used them as a chance to take a break from the prepared materials and regain a connection with the audience. One speaker said he looked at written questions when he started to lose his sense of the audience. From the audience's end, people were able to ask questions as they arose in their minds rather than waiting for a good time to break in. People who tend to be shy about talking in front of a group were able to contribute their comments, as were those who did not have working microphones or were experiencing other audio problems. And in general, written comments tended to be expressed more concisely than the spoken ones.

The poll meter was the most unconventional feedback mechanism provided and also the most enthusiastically received. Asking poll questions during the talk enabled speakers not only find out the answer to a question, but also to assure themselves that there were people "out there." One speaker used a five-choice poll in her talk about the upcoming Academy Awards. She presented the choices for each major award, held a discussion about each nominee and then asked the audience to vote for the winner. This use of the poll kept the audience highly engaged and gave the speaker a strong sense of the audience's attitudes. In general, audiences liked that the poll gave them the opportunity to indicate their presence as well as express their opinion. Some said that they especially liked having an anonymous method for expressing feedback. "It was nice to be able to anonymously raise my hand with, say, a wrong answer and not feel embarrassed," one user wrote.

Finally, the slides were also effective for maintaining a shared focus among the speaker and audience, just as they do in a traditional talk. We were surprised that only about one speaker in four used annotations, but when they were used, audience members said they were very

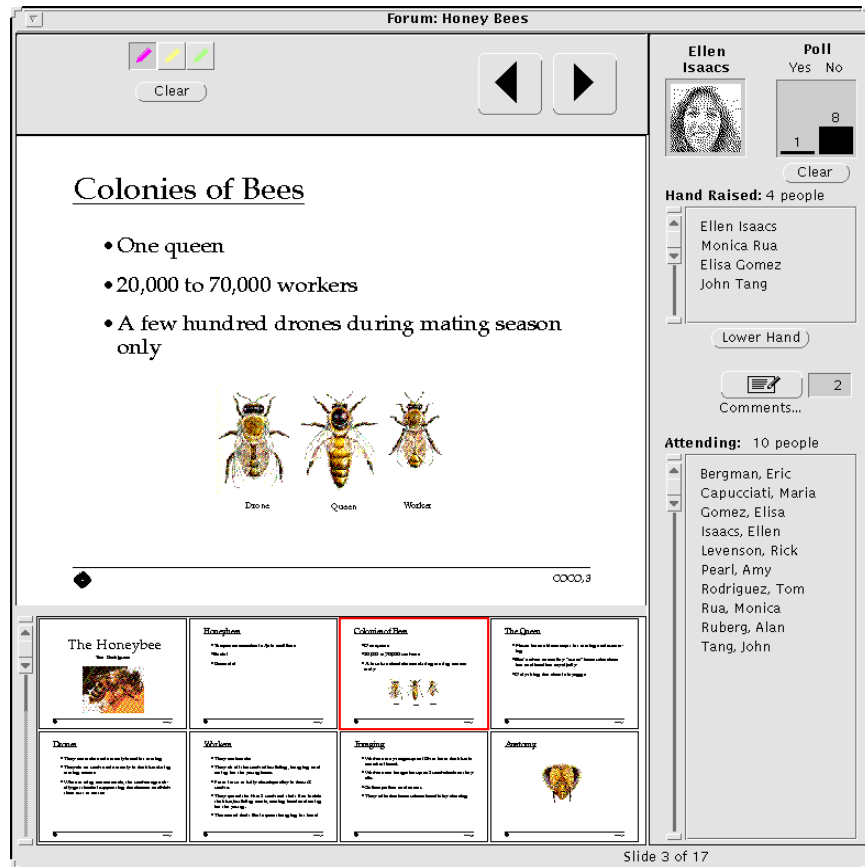


Figure 4. Forum speaker's interface. An audience member is asking a question.

Figure 4 shows the speaker's interface for giving presentations. The large slide area in the upper left shows the current slide being presented to the audience. Speakers can progress one by one through the slides by using the arrow buttons, or they can go directly to a different slide by clicking on one of the thumbnails in the region below the main slide. They can also use the colored pens in the upper left to annotate their slides, and their marks are seen by the audience, as are their gestures with the pointer. In the upper right panel are the interfaces to the three mechanisms for audience interaction.

The upper right region shows the results of an audience poll, which the speaker can clear by pressing the button beneath the graph. (Before a presentation starts, a speaker may elect to provide a poll with as many as five choices, which several speakers did during the tests.) To the left and below are the controls for audience voice interaction, which are essentially the same as the audience's. The speaker usually manages the questions simply by calling on the person at the top of the question queue. The Comments button in the middle region brings up a window that shows written comments submitted by the audience. A speaker can show a comment to the rest of the audience by pressing a button that

reproduces the comment as a slide. The lower part of the canvas shows the list of audience members currently running Forum.

Observing Forum usage

When designing an application, it is important to test it with users and revise it based on their feedback. While developing Forum, we conducted a series of tests in which we presented talks to a group of people throughout the company. We began this process with an early version of Forum by giving relatively informal talks to a pool of about 75 people who were familiar with our project. As Forum became more stable and we added more functionality, we began presenting regular weekly one-hour talks to a larger set of people, sometimes sending invitations to over 1000 people. A total of 26 presentations were given over a 10 month period by people at Sun who volunteered to talk about a topic of their choice. Some were technical talks about projects going on around the company and others were about non-work-related topics of interest to a speaker. Audience size ranged from 10 people to 95 people, with an average of 41 people attending. Most of the early speakers and audience members were sophisticated computer

between the Forum talk and any other application on their desktop.

Forum runs on the Sun SPARCstation™ line of computer workstations running the Solaris® 2 UNIX operating system. The speaker's workstation is equipped with a camera, a microphone and a SunVideo card, which digitizes and compresses the video. Audience members need no special equipment on their workstation other than a microphone. The 320x240 pixel video image is compressed using CellB, a proprietary Sun video encoding. The video is transmitted at 4 frames per second, yielding a video data rate of 400 Kb/sec. The audio is transmitted at 8KHz μ law, with a data rate of 64Kb/s, which gives adequate speech quality. The slides are in PostScript™ format, which makes them easy to scale.

Forum runs over one domain within Sun's Wide Area Network. This domain is made up of standard Ethernets connected by a high-bandwidth backbone. We use the Internet multicast protocol for the broadcasts, which enables us to keep the network load low regardless of how many people attended a talk. Forum has had a negligible effect on backbone traffic, and it has used up to 7% of the available Ethernet bandwidth.

User interface

Figure 1 shows the audience's interface. The left side of the main window shows video of the speaker, and the controls below it manage the audio parameters. The control panel on the right provides mechanisms to interact with the speaker. To get in the queue to speak, a user "raises her hand" by clicking on the button at the bottom of the window; her name appears in the question queue above it. When the speaker calls on her, she presses and

holds down the **Speak** button and her hand is automatically lowered (i.e. her name is removed from the list). While her voice is transmitted, a small picture or "icon" of her face appears in the square region in the upper left corner with her name above it.¹

To vote on a poll, the audience clicks on the **Yes** or **No** option in the upper right, and the bar above it increases one unit per vote. To submit a comment, the user clicks on the **Comments** button, types in a comment and sends it to the speaker. To find out who else is watching a presentation, the user clicks on the **Audience** button to bring up a window that shows a list of attendees (Figure 2). In that window, they can click on a name to see that person's icon, their location and phone number. They can send that person a short message by clicking on the **Message** button, which brings up a small window pre-addressed to that recipient. When they send the message, it pops up on the other person's screen in a small window with a **Reply** button.

Audience members see the speaker's current slide in the slides window (Figure 3), and they can also click on a thumbnail to view a different slide. To get back in synch with the speaker, users either click on the **LIVE** button or click on the thumbnail of the speaker's current slide. Users can click on colored pens to annotate their own copy of the slides with text or free-form lines.

1. If an audience member doesn't have an icon, a default icon is displayed. We wrote an application to help people customize their icon, but most audience members did not have a customized icon.

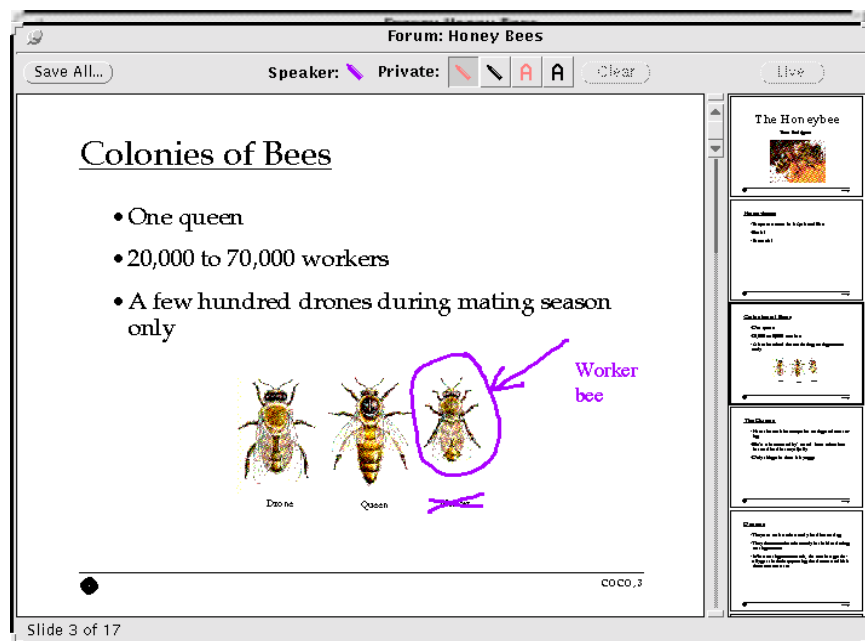


Figure 3. The audience's view of the speaker's slides, with text and line annotations

questions, and sometimes they ask the audience to raise their hands to indicate their level of background knowledge or experience with the topic. Sometimes just seeing the physical characteristics of the audience (such as age, attire, etc.) can help the instructor adjust her material to better interest the audience.

Not only do speakers learn from audiences, but audience members also gain important information from observing and interacting with each other. If an attendee doesn't understand a minor point the speaker made, she can lean over to her neighbor and ask for a clarification rather than stop the class by asking a question. When instructors ask the audience to raise their hands to "vote" on some question, everyone can see where they stand in relation to the others. And of course the audience is likely to be just as interested as the speaker in the other participants' ongoing reactions to the talk (their level of interest, their amusement at a joke, etc.).

In addition, there is evidence to suggest that students in an instructional setting learn more when they interact with the speaker than when they do not [Moore, 1993; Gibbons, et. al., 1977]. Although some researchers have concluded that remote students tend to be at a slight disadvantage relative to local students, there is evidence that the degree of interactivity may be the most relevant factor for quality training. In fact, one study found that remote students outperformed classroom students when they watched a tape of the class and stopped the tape to discuss the content [Gibbons, et. al., 1977]. In light of this finding, it seems worthwhile to look for effective ways to enable interactivity among a presenter and distributed participants.

A tool to support presentations over a network, then, should not only enable speakers to transfer information to audiences, but also enable audiences to interact with

speakers and the rest of the audience in a natural and effective way. However, providing such interactive capability represents a significant challenge, largely because it requires overcoming three limitations inherent to the environment:

- the participants are distributed over potentially large distances,
- the size of the group can be quite large,
- the speaker and audience play asymmetric roles.

This paper describes how we designed an application for broadcasting presentations over a network, which we called Forum. It focuses on how we attempted to enable interaction among participants and evaluates how successfully people were able to use the tool.

FORUM DESIGN

Functionality

Forum is a distributed application that enables speakers to broadcast talks over a network and enables audience members to participate in the talks from their workstations. The audience receives live audio and video of the speaker as well as the speaker's slides and slide annotations. Audience members can interact with the speaker in three ways. They can speak to the presenter, they can "vote" anonymously on an issue raised by the speaker, and they can send in written comments.

Speakers give presentations while sitting in front of a workstation screen that displays the slides, a video preview of themselves and the rest of the Forum interface. Audience members also sit at their desks and watch the talk through their interface. Since they are in a multi-tasking environment, they can switch their attention

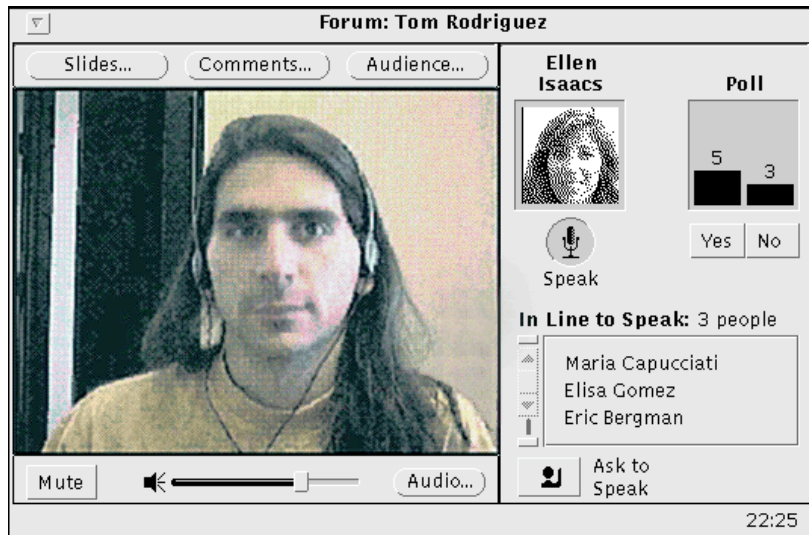


Figure 1. Forum's audience interface. A poll has been taken and an audience member is asking a question while others wait in line to speak.



Figure 2. Audience list with a name selected

A Forum for Supporting Interactive Presentations to Distributed Audiences

Ellen A. Isaacs, Trevor Morris and Thomas K. Rodriguez

SunSoft, Inc.

2550 Garcia Avenue, MS 21-225

Mountain View, CA 94043

Phone: (415) 336-1167; (415) 336-3097; (415) 336-4918

Email: ellen.isaacs@sun.com; trevor.morris@sun.com; tom.rodriguez@sun.com

ABSTRACT

Computer technology is available to build video-based tools for supporting presentations to distributed audiences, but it is unclear how such an environment affects participants' ability to interact and to learn. We built and tested a tool called Forum that broadcasts live audio, video and slides from a speaker, and enables audiences to interact with the speaker and other audience members in a variety of ways. The challenge was to enable effective interactions while overcoming obstacles introduced by the distributed nature of the environment, the large size of the group, and the asymmetric roles of the participants. Forum was most successful in enabling effective presentations in cases when the topic sparked a great deal of audience participation or when the purpose of the talk was mostly informational and did not require a great deal of interaction. We are exploring ways to enhance Forum to expand the effectiveness of this technology.

KEYWORDS: Broadcast video, distributed presentations, distance learning, remote collaboration, user interface design, multimedia.

INTRODUCTION

Over the last few years, the multimedia community has been exploring ways to use distributed, interactive audio and video. Point-to-point desktop video conferencing and multi-way "video spaces" are two of the applications that have received attention [Bly, et. al., 1993; Tang and Isaacs, 1993; Root, 1988]. Another obvious application is to enable people to broadcast live presentations to audiences that are distributed across many locations.

There has been a growing interest in such one-to-many technology. Politicians have recently made popular the

idea of an "electronic town hall" meeting in which voters all around a nation could not only watch but also interact with policy makers. And at a time when many institutions are struggling with budget cuts, both universities and companies have shown interest in providing live training over a computer network. By broadcasting courses, schools can reach a wider set of potential students and companies can train more of their work force while reducing travel costs.

The distance learning industry has been using technologies that provide one-to-many communication for quite some time. The Open University, for example, has been airing courses over British television for many years, and, more recently, such programs as Stanford University's Instructional Television Network and the Chicago TV College have enabled remote students to not only watch but also phone in questions to the speaker [Gibbons, et. al., 1977]. These environments, although quite successful, represent only a step toward supporting fully interactive presentations. What they lack is a rich channel for audiences to interact with speakers and with each other. Although the phone does provide some audience-to-speaker communication, it is a cumbersome mechanism that does not allow for quick, spontaneous questions and discussion. It also does not support audience-to-audience interaction. With the development of technology to support distributed video and audio over existing computer networks, we are now in a better position to support more fully interactive presentations.

THE IMPORTANCE OF INTERACTION

One might question how important it is to provide better interaction capabilities for presentations. After all, the bulk of the information seems to flow from the speaker to the audience. However, there are in fact many ways in which speakers depend on information from the audience to help ensure the success of the presentation. They gauge the audience's level of interest and understanding not only from overt responses, such as applause or laughter, but also from monitoring audience member's facial expressions, body posture, and activity level. Speakers learn about the audience when they accept

Published in the proceedings of Computer-Suported Cooperative Work, 1994, ACM: Chapel Hill, NC, 405-416.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of ACM. To copy otherwise, or to republish, requires a fee and/or specific permission.

CSCW 94 - 10/94 Chapel Hill, NC, USA
© 1994 ACM 0-89791-689-1/94/0010..\$3.50