Back to School: Design Principles for Improving Webcast Interactivity From Face-to-Face Classroom Observation

Jeremy Birnholtz
Knowledge Media Design Institute
University of Toronto
40 St. George St. – Rm 7204, Toronto, Ontario M5S 2E4
jeremy@kmdi.utoronto.ca

ABSTRACT
This paper presents an observational study of face-to-face university classrooms and provides preliminary design principles for improving interactivity in "webcast" presentations. Despite the fact that participation and interaction patterns appear to depend heavily on presentation style and class size, useful patterns were observed and analyzed. Design principles presented include the need to support rapid changes in floor control, multiple types of presentation technologies, and the subtleties of awareness between the audience and presenter.

Author Keywords
Remote audience, presentation, classroom, observational study, interactivity, webcasting

ACM Classification Keywords
H.5.3 Group and Organization Interfaces: Synchronous interaction

INTRODUCTION
Despite growing interest in "webcasting" [25] and "telepresentation" [8] technologies, there remain significant differences in user experience between face-to-face and webcast presentations. Webcasting tools, for example, do not support frequent and fluid interaction between remote audience members and the presenter. While some have suggested that such experience gaps may never be completely closed [1, 15], there are steps that may improve remote user experience.

Recent research, for example, has focused on improving interactivity in presentations to distributed audiences via improved presenter/audience awareness [5, 8], automated camera operation/selection [17], and the capability for two-way voice interactions [2].

DIS 2006, June 26-28, 2006, University Park, Pennsylvania, USA.

There is little agreement, however, about how to improve interactivity and participation in distributed lecture environments. It will be argued here that this stems, in part, from an incomplete understanding of human behavior in face-to-face lecture environments. Surprisingly, little prior development work has been based on a systematic understanding of expected behavior in face-to-face presentation environments.

If we are to effectively improve the interactivity of webcasting tools, however, we must begin with a principled set of design guidelines that accurately reflect likely user expectations and behaviors. Specifically, we must improve our understanding of participation patterns and the subtleties of communication in presentation environments.

The paper that follows represents a modest step in that direction. Observational data presented here suggest that participation and interaction patterns vary substantially between settings, but relate strongly to class size and presentation style. Many subtle aspects of communication in the classroom are observed and used to derive preliminary design principles.

BACKGROUND AND RELATED WORK
Few studies of presentation environments have focused on deriving design principles for interactive systems. Nonetheless, there are lessons that may be learned from relevant work in three areas.

Research in Education
Despite recent focus in the education literature on novel education techniques, such as distance education [24] and inquiry-based learning [19], it is nonetheless the case that lecturing is one of the most common instructional techniques in university education [14].

Research that has been conducted on lectures [13] has identified a range of discrete lecture styles [18] and suggests the importance of interactivity [21] in engaging students. These authors, however, focus largely on pedagogy and provide limited insights for how this might be achieved in webcast environments.

Presentation system developers

While there is a growing body of research literature from developers of remote presentation tools, little of this work is grounded in systematic observation and analysis of data from face-to-face classroom environments. Instead, these studies have been carried out in environments where technology interventions of one sort or another have already taken place, and have likely caused participants to modify their behaviors in some ways. Thus, it is difficult to generalize these findings to other environments.

In some cases authors have relied on observations or personal experience with existing presentation tools. Chen [5], for example, observed a small number of existing distance education classrooms in order to identify potential enhancements. Similarly, guidelines presented in Baecker's work [3] are derived from the author's observations and experience with a prototype presentation tool.

Another set of studies focus primarily on user experience with particular presentation technologies, but they do not carefully examine the face-to-face experience they seek to replicate, Isaacs et al. [7], for example, present a pioneering interface for remote presentation and demonstrate the value of such systems. Zijdemans, et al. [26] describe their experience with the use of their own tool in conducting an academic workshop. Others have described experiences with classroom technologies geared toward specific functions, such as lecture capture [4] and awareness [8].

Note further that this research area stands in contrast to other interactive systems research in lacking systematic study of face-to-face behavior on which designs can be based. Designers in other areas, for example, frequently cite studies of face-to-face meetings [20], collocated collaborative tasks [22], and how individuals organize their physical desktops [10].

Meetings/CSCW

A third stream of relevant work that bears mentioning here stems largely from the CSCW community and is focused on meetings. A wide range of tools have been developed and tested in this area, several of which are based on observations of face-to-face meetings with and without technological interventions [16, 20]. Additional work has focused on novel communication and control technologies in meeting room environments[9, 11].

As Baecker [3] points out, however, meetings and lectures have very different requirements and expectations for behavior. Thus, conferencing technologies supporting continuous, many-to-many interaction are typically used for distributed meetings, while webcasting technologies that support primarily one-to-many interaction are used primarily for distributed presentations and lectures. Moreover, webcasting boasts the advantages of low barriers to entry and easy scaling for large audiences [25].

Webcasting: The state of the art

Observations and design principles presented here are based on an understanding of technologies currently available, commercially and otherwise, for the webcasting of lectures. A more detailed description of these technologies is available in [25], but a brief description is provided here. The majority of these tools support one-way streaming of audio and video to a large number of participants, with two-way interaction limited to primarily to text chat and polling. Audience members wishing to ask questions or comment generally indicate their desire to participate by clicking on a button which signals the presenter, and he or she can "call on" the audience member. Static PowerPoint slides can typically also be streamed, and screen or browser sharing is permitted by some systems.

While these tools were not generally designed to support extensive interactivity during webcasts, evidence from the education literature (see above) suggests value in doing so. Such improvements must be made, however, without sacrificing the advantages of webcasting over conferencing or other approaches. .

RESEARCH CONTEXT AND METHOD

Observations were conducted in four courses at a large university in Canada. The sample was partly a convenience sample of colleagues willing to allow an observer in their classrooms. Courses were selected deliberately, though, to obtain a range of academic disciplines, class sizes, and student standings. Courses were in the areas of Sociology (50 students), Computer Science (100 students), Political Science (500 students) and an interdisciplinary seminar in media design (20 students). Students included were first-year undergraduates, advanced undergraduates, and master's students. Instructors ranged from advanced graduate students to senior faculty. While this wide variety cannot allow reliable attribution of observed differences to the factors of differentiation, it was felt appropriate for an exploratory study.

Observations

Observations were conducted during the fall, 2005. The first classes of the term were deliberately avoided because these are frequently atypical in format/style [14]. Four sessions of each class were observed by the author, for a total of 16 classes, or 26 hours of observations. Extensive field notes were taken during the classes, and these were typed and expanded after each session. Particular attention was paid to the use of visual aids of varying sorts, changes in speaker/floor control, and potential disruptions that were present. Other events were noted as deemed appropriate.

Analysis

Once data collection was complete, the typed field notes were read several times to check for accuracy and begin to find themes for later coding. Analysis was then carried out in two phases. First, participation and the use of visuals were quantified and counted. Results from these analyses are presented below. Next, specific episodes were coded

into categories, based on themes observed in reading through the notes. These categories were iteratively refined throughout the process, and notes were re-coded as needed.

RESULTS

Results will be presented in three sections. First, participation statistics will be summarized. This will be followed by discussions of what works well in the face-to-face lecture environment, and of what problems could potentially be improved on by presentation technologies.

Look Who's Talking: A Summary of Participation

One important question in this work is improving understanding of participation patterns such as how many people typically speak in a lecture, and the extent to which this varies based on parameters such as class size and presenter style. Participation statistics were derived primarily by closely examining field notes for changes in floor control, which were carefully tracked during observation. A change in floor control is considered to have taken place any time there was a speaker change.

Student Contributions

Because the observed courses were primarily lecture-oriented, it was accepted as a given that the presenters would, on the whole, speak more than any audience member. More interesting from the standpoint of interactivity was how floor control broke down between the presenters and students. About half (46%) of the 971 changes in floor control gave control to the instructor and about half gave control to students. This does not mean that instructors and students spoke for equal amounts of time (which was not the case), but rather suggests that, when a student was speaking, floor control almost always came back to the instructor before another student started speaking. Moreover, because floor control changes were roughly equally split between presenters and students, counting the number of contributions by either of these groups gives a reasonable indication of how much interaction took place in any lecture. Thus, student contributions are used to measure lecture interactivitity in the remainder of this section.

To normalize for variations in class length, the overall number of student contributions in each lecture was divided by the number of minutes in that lecture. Overall, there were a mean of .29 student contributions per minute (SD=.20). As can be seen in Table 2, there was substantial variation along this dimension between courses. An ANOVA analysis comparing the mean student contributions per minute for all lectures (N=16) showed a statistically significant main effect for field, $F(3,12) = 3.5$, $p < .05$. When contrasts were tested, however, only the difference between sociology and the combined mean of the other fields was significant, $t(12) = 3.20$, $p < .01$.

Generally, these differences seem to be less a function of the nature of the course content in each discipline, however, and more closely related to individual instructor presentation style and the overall size of the class. In the sociology class, for example, the instructor had a tendency to engage in frequent, rapid, instructor-initiated (i.e. "cold call") dialogs with students. These accounted for a large fraction of the observed student contributions in those lectures. This instructor would periodically speak for several minutes, and then engage in a dialog that might involve five or six student comments in a single minute.

Other presenters tended to carry on with their presentation until interrupted by a student. One effect of this teaching style in several of the observed courses was relatively long, uninterrupted periods where the instructor spoke to the class. This was particularly true in the political science and computer science classes, where the mean student contributions per minute are lowest of all. It also seems likely that class size plays a role in that students are arguably less likely to interrupt a speaker in a room full of 500 students than in a seminar with 15 or 20 students.

Also of interest here is the fraction of the total number of participants who may wish to speak. Knowing this can provide a rough estimate of webcast participation expectations and an estimate of the need to provide remote users with interactive functionality. Table 1 shows that this appears to depend at least somewhat on class size, as the two smallest classes (sociology and the interdisciplinary seminar) have the highest fraction of participants as speakers (68% and 51%, respectively, compared with 3% and 6% in the two largest classes). This is not surprising because there is a finite amount of class time that can be devoted to class participation (i.e. it simply would not be possible for 50% of 500 students to speak) and in that instructors in smaller classes may invite more participation.

Repeat Participation

Another aspect of participation important to system designers is understanding how often students speak more than once in a class period, and what is the nature of their

Table 1. Summary of lecture participation

	Student Contributions (per minute)		Unique Student Speakers		Fraction of students who spoke	
Field	Mean	SD	Mean	SD	Mean	SD
Poli. Sci.	.19	.20	17	19.35	.03	.04
Interdisc.	.25	.13	14	5.20	.68	.26
Computer Sci.	.19	.16	6	2.83	.06	.03
Sociology	.52	.12	26	5.92	.51	.12
Overall	.29	.20	15.44	11.88	.32	.32

repeat participation. Is there, for example, a small group of students that dominates the "air time?" To examine this, student contributions were coded based on whether the student had spoken previously during that class period and whether they had spoken within the previous five floor control changes.

These results are presented in Figure 1, as fractions of the total set of student contributions. In the first place, the fraction of repeat speakers is far lower in the largest class (political science) than in the other three. When tested in an ANOVA comparing mean fractions of repeat speakers across all lectures (N=16) with field as the independent factor, this contrast is statistically significant, $t(12) = 2.66$, $p < .05$. This means that any given speaker was less likely to have spoken before in the political science class than in the other three classes.

It is also noteworthy that the highest frequency of repeat speakers occurred in the two smallest classes, sociology and the interdisciplinary seminar. The contrast between these two and the larger two classes was also statistically significant, $t(12) = -2.53$, $p < .05$. This means that, as we would expect, there were more repeat speakers in the smaller classes than in the larger ones. Moreover, the higher fraction in sociology than in the interdisplinary seminar is likely due to that instructor's presentation style.

Another common behavior in face-to-face lecture environments is "follow up" or repeat participation. Student speakers periodically ask multiple questions in a row in order to clarify or engage in dialogue with the presenter. This can be difficult in some webcast environments, however, because delays from streaming make real-time dialogue nearly impossible. If designers had a better understanding of how frequently and in what environments repeat participation occurred, however, this behavior could be better supported.

Thus, Figure 1 also illustrates differences in the fractions of student speakers who had spoken within the previous five floor control changes, or "recent repeat speakers." There

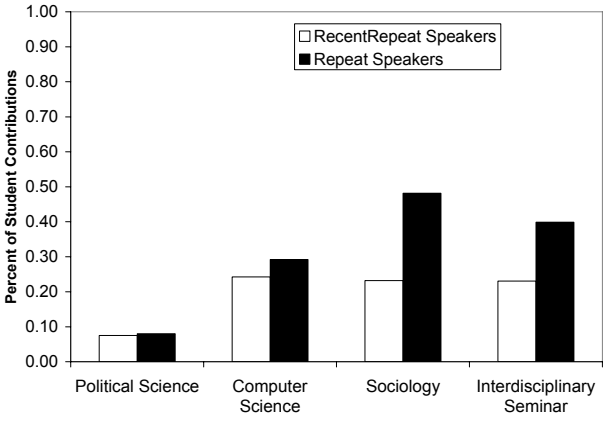

Figure 1. Mean Fractions of Student Contributions Attributable to Repeat Speakers and Recent Repeat Speakers

was, on the whole, surprisingly little difference between the fields on this dimension. There were fewer recent repeat speakers in the political science class, but this difference was not statistically significant. Fractions of recent repeat speakers in the other classes were nearly equal.

Usage of Visuals

There are three aspects of the presenters' usage of visual aids that are of potential interest to designers of interactive systems for lecture environments: 1) how frequently visuals are used, 2) what types of visuals are used and, 3) what are the ways in which they are used.

How frequently are visuals used

All presenters observed made use of PowerPoint slides projected on at least one large screen. There were differences, however, in the extent to which PowerPoint was combined with the usage of other visual aids, as is described below.

To better understand the use of visuals generally, the frequency of "visual changes" were carefully tracked during observation. For PowerPoint slides, a "visual change" was defined as a change to a new slide. For other forms of visuals, it was defined as a nearly complete replacement of screen contents.

Interestingly, when the number visual changes were tracked in this manner and normalized by dividing into the number of minutes in the lecture, there was little difference between the four courses observed. The overall mean number of minutes per visual change was 3.04 (SD = 1.41), with individual lectures ranging from 1.33 to 6.00 and mean values for each field ranging from 1.85 to 4.47. These differences were tested for statistical significance using a one-way ANOVA comparing mean values for all lectures (N=16) with field as the independent factor, and testing contrasts. The only significant difference was between the computer science course and the sociology course, $M_{CS} = 1.85 < M_S = 4.47$, $t(12) = -3.39$, $p < .01$).

All of this suggests that, on average in all classes, the contents of the screen were changed by the presenter about every three minutes. Presenter and presentation style appears to account for the observed differences between fields, as will be described below.

What are the types of visuals used?

While traditional PowerPoint slides were the most common visual aid in use, 27 instances were observed where other types of visual aids were used (of 571 total visual changes). While these instances represent only 5% of the total, 20 of them (74%) occurred in the sociology class.

This is interesting in two respects. In the first place, it validates overall the assumption by many webcast software designers that PowerPoint slides are the primary visual aid in use and therefore the most important to support. At the same time, however, these results suggest that it may be far more difficult for some instructors to adapt their

presentation style to online delivery of lectures without improved support for additional media.

The additional types of visuals used in these courses were also tracked. The most common alternative visuals were web pages, PDF files, and Microsoft Word files. A few presenters also made use of physical actions and artifacts.

Web pages accessed were used to illustrate assignment requirements, the location of readings, or to discuss particular web-based technologies. Recommended "blogs" and web sites from research projects were also common. PDF files were used primarily to display specific documents, such as letters or assignment specifications. Microsoft Word files were used primarily by the sociology instructor, who had two uses for these. He periodically displayed research papers he had written that seemed to serve as a guide for his own comments about the work. Second, he typically kept a Word file open on the podium PC in which he typed new terms as he introduced them.

Finally, a small number of presenters made relatively infrequent use of their bodies or of physical objects as visual aids. For example, one presenter jumped animatedly into the air, landing hard to illustrate that you can't "stomp your feet and get angry" in an online dialogue.

It also bears mentioning that many webcasting tools support only static PowerPoint slides, though presenters here were observed to make frequent usage of animations and videos. One lecture in the computer science course used a series of approximately 30 short videos to illustrate a novel user interface. In another lecture in the computer science course, precisely choreographed PowerPoint slides were used to replicate cognitive psychology experiments (e.g. to illustrate the Stroop effect).

How were visuals used?

In reviewing field notes from the observed lectures, it became apparent that there were two very distinct styles of presentation in use, with some occasional overlap.

Most of the observed presenters spent most of their time with the podium computer in the "slide show" mode of PowerPoint. In this mode, slides are the only thing displayed on screen, and the visuals are presented in a very polished manner (i.e. with little or no time spent locating the next slide, sometimes with formal transitions and animations, etc.).

On the other hand, some presentations that integrated multiple types of visual resources were less "polished" in that the audience had a full view of the "behind the scenes" work of finding, opening and switching between files. In one class, the presenter was able to display multiple resources without revealing his "behind the scenes" activity by switching periodically between two computers connected to the classroom projector. One had his PowerPoint slides and remained in slide show mode. The other was used to display related web sites, which were pre-loaded by a teaching assistant while the other computer was being used.

What works well in the co-located environment?

One frequent complaint from users of remote presentation tools is that the experience of being remote is "not the same as being there." Improving our understanding of this difference requires an improved understanding of what makes physical copresence so desirable in a setting where the bulk of the time is spent listening to a single presenter. This question was considered in both observation and analysis, and episodes deemed relevant by the author were noted and later coded for analysis. Coding divided relevant episodes into these categories:

Subtle communication and awareness

Despite the fact that information flows primarily in one direction (from the presenter to the audience) in the lecture environment, there is a great deal of subtle communication that is taking place constantly.

Hand raising

One particularly apparent example of this in the classes observed here was hand raising by the students. Outwardly, hand raising does not appear to be a subtle form of communication. To the casual observer, a student's hand is either up or down.

Observations in this study, however, revealed that there are many subtle dimensions to this practice for both the presenter and the student. In observing students, it first becomes clear that hand state (raised or not) is not a binary condition. While it appeared to be acceptable in most of the observed classes to raise one's hand while the instructor was speaking, students raising hands while other students were speaking was much less common. Instead, when another student was speaking, many students who wished to speak appeared to put a great deal of concentration into timing the raising of their hand so it coincided precisely with the conclusion of the student speaker's point. This was evident both in the stringent attention paid by these students to the current speaker, and in the number of "false starts" to hand raising that resulted in the student "scratching" his or her neck or hair.

One interpretation of this scenario is that it represents a problem easily solvable by using a queue to ensure that students wishing to speak are handled in order. This interpretation, however, is likely an oversimplification.

Rather than simply pointing to a desire to be "next," the variable states of hand-raising observed seem to point to varying states of desire to speak. When a student is particularly engaged by another student speaking, she may start to put her hand up. If the speaker then switches topic or makes the point the student wished to make, however, she will not likely continue to try to raise her hand.

For the presenter, hand raising gives the instructor a sense of who wishes to speak. The instructor then has the freedom to call on a student when he or she pleases.

Second, the varying states of "hand raisedness" provide the presenter with some information about who is thinking about speaking. This is potentially useful both in moving discussion forward and, possibly, in computing participation grades.

Third, even students with their hands in a fully lowered state provide the instructor with some useful information. Two of the observed instructors engaged in frequent "cold calling" on students. One common use of this technique is to engage students who do not seem to be paying attention or those who are otherwise unlikely to participate. Instances of both of these categories were observed. In one case, the instructor seemed to deliberately call on a student she knew to be asleep, and then asked a nearby student to "poke him," much to the class's delight. When the student woke up, he was unable to answer the instructor's question.

In this way, a lowered hand can signal a lack of engagement or lack of desire to speak. Subtle visual information (e.g. gaze, likelihood of being awake, etc.) then becomes important for the instructor in deciding whom to call on.

Chiming in

Despite reasonably strong social norms that govern floor control in presentation spaces, there were several observed episodes where these norms were violated, arguably to the benefit of the entire group. These episodes fall into three categories as follows:

First, there were several cases where one person "talked over" another. Most often this was the instructor's way of indicating to a student that she had spoken long enough, or that a point was off-topic and it was time to move on. In one class, for example, there was a student who spoke up very frequently, but his comments were often off-topic. In these cases, the instructor would simply resume talking where he had left off. This student understood the implicit message and stopped talking, but there were also a small number of episodes in other courses where both student and instructor talked over each other until one or the other eventually stopped talking.

Second, there were cases where it was clear that the instructor and a student speaker misunderstood each other (Clark [6] would refer to this as a breakdown in common ground). In these cases, particularly in the smaller two classes, it was typical for another student to "chime in" and help to clarify what was being said. In one case, for example, there was a misunderstanding about an assignment specification. One student walked up to the instructor to point out the discrepancy on a printout and the issue was resolved.

Third, there were cases where several students "chimed in" simultaneously to convey a message of some sort to the instructor. In one example of this, a guest making a brief announcement in the political science course (of 500 students) elected not to use a microphone. Students could not hear him and they shouted out to let him know this.

Lecture as theatre/entertainment

In many respects great lectures can be considered a form of theatrical entertainment. Like performers in a theatre, great lecturers engage, derive feedback from, inform and entertain their audiences. Several aspects of this were observed in this study:

Humor was frequently used by presenters, who seemed to anticipate and enjoy the audience response. In some cases, humorous stories or examples were used that illustrated relevant points, but also served to entertain. One instructor, for example, showed several 15-20-year-old research videos of technologies that seemed quaint by today's standards, but nonetheless illustrated salient points from the lecture. In other cases, humor was used simply to entertain (and perhaps to wake up and/or engage) the students, as when another instructor did an impression of a former prime minister's accent that had most of the room laughing.

This may seem unimportant or trivial, but humor was also observed to have an important function in one episode. In this case, the presenter was under the mistaken impression that the class could see his PowerPoint slides on the large screen behind him. The slides were not on the screen, however, and the presenter did not realize this until he switched to a slide containing a cartoon, which the students could not see and therefore did not laugh at.

Building rapport with the students

Another aspect of the theatrical environment is the extent to which instructors were observed to draw on the shared physical environment in building rapport with students. One common example of this was when cell phones rang loudly during class. This was sometimes an occasion for visible display of annoyance, but it was also an occasion to build rapport at times. One instructor, whose own cell phone rang during a lecture, mockingly asked, "Does this mean I'm in trouble with Professor [Jones]?" (a pseudonym for his co-instructor, who had explicitly forbidden phones in class). The students laughed and seemed to appreciate what was arguably an attempt to reduce the formality of the large lecture and break down barriers between the instructor and students.

Beyond Being There: What doesn't work?

Thus far the physical presentation space has been described as extremely effective. This is not to say that it is an absolute "gold standard" to be replicated online in every respect, however. There were also several problems observed that could be addressed in both face-to-face and remote presentation tools and environments.

Technology troubles

A range of difficulties with presentation and room control technologies were observed in 37 specific episodes. On three occasions, the difficulties caused a long enough disruption that the instructor halted the lecture and took a scheduled break at an unanticipated time. These problems included difficulties with projection equipment, light controls, computer and interface difficulties, and sound equipment. The main point here, is that, like remote presentation technologies that are often criticized for requiring a facilitator to ensure smooth operation, face-to-face presentation environments are also rife with technology troubles. As is suggested below, it may be possible to address some of these in the design of tools to better support both local and remote presentations.

Hearing troubles

In the three largest classes, there were very frequent difficulties with hearing. Most typically, this was dealt with by the instructor by repeating or summarizing each student comment in a loud voice (or using a microphone) so the rest of the class could hear. In at least a few cases in each class, though, not even the instructor could hear the student comment and had to ask for repetition.

Distractions

Just as there are many distractions for remote attendees who may be multi-tasking while attending a lecture or meeting, a great many distractions for local attendees were observed as well. This was true even for those students without access to a computer on which to check their email or chat with friends via a wireless network. These distractions were not tracked systematically, but were noted during observation on an opportunistic basis to capture a sense of their scope.

At least some students in all of the observed classes were periodically engaged in activities not directly related to the presentation at hand. In one class, students seated at the back of the room tended in several sessions to be reading the newspaper, sleeping, or working on what appeared to be homework from another class. In another case, two students seated near the observer passed a cellular phone back and forth to view photographs while another student sent text messages throughout the lecture. In yet another case, two students had their arms around each other and were constantly playing with each other's hair. Laptops, particularly in the small seminar, were common and students periodically seemed (from a distance) to be checking email or using Google.

DISCUSSION: TOWARD DESIGN PRINCIPLES FOR IMPROVING INTERACTIVITY

This section uses the data described above to derive preliminary design principles for improving interactivity in webcasting technologies. Design tradeoffs are then discussed for each principle, along with suggestions for future work to refine and expand on this exploratory study.

Participation ratios and rates

Designers of interactive presentation systems have historically had little empirical data on which to base expectations of audience participation. While these data do not necessarily generalize to the larger set of all lectures, they do suggest that class size and presenter style are significant factors in understanding audience behavior. While conventional wisdom might suggest that larger classes will tend to have fewer student participants, though, this was not always the case. As shown above, fewer students in larger classes (as a percentage of the class size) spoke than in smaller classes on the whole, but there tended to be more unique speakers in larger classes. Thus:

Principle #1: All remote participants must be provided with the capacity to interact with the presenter, even in very large events.

Certain classes observed here also suggest that some instructional styles may be quite difficult to support with current technologies. This was particularly apparent in one political science lecture where there were 44 unique student speakers called on by the instructor, and, more generally, in the sociology class where the instructor's tendency to engage in dialogues with students appeared to have a significant impact on participation trends. From design standpoint, these behaviors present several challenges.

First, most presentation systems are designed to let instructors know when students want to speak, but the reverse is not generally true. Even where a presenter can click to designate a speaker (or, in the case of Chen's virtual auditorium, use gaze to provide some semblance of eye contact), this designation would likely be quite jarring to a remote student ("Who, me? Can you hear me?"), and potentially embarrassing to everybody if the remote attendee is multi-tasking in another application and does not even know she has been called on. Thus, there is a need to improve not just presenter awareness of the audience, but also the audience's awareness of the presenter's focus of attention. In cases where multi-tasking is expected (such as large lectures), this could be as simple as displaying a flashing "notifier" icon (such as when new email or instant messages arrive) when the student is called on, and even when the presenter is "looking at" or seems likely to click on that student's name/image.

Principle #2: Awareness goes both ways. Audience members and presenters both need awareness of each other's state and focus of attention.

Second, presentation systems that rely on webcasting, because of compression and transmission delays inherent in that technology, do not effectively support the "rapid-fire" exchanges observed here. While many do support text chat, this is awkward for presenters and is frequently primarily accessed by a moderator for that very reason. Thus, another important dimension in improving interactivity is allowing for smoother, faster and more natural transitions between speakers. Faster streaming technologies or hybrid solutions,

though, are only part of the answer to this problem. There is also a need to improve the speaker's ability to engage students directly, via an interface that is as effortless as pointing at a student in a face-to-face classroom.

Principle #3: Seamless interactivity requires support for rapid and spontaneous changes in floor control, without the intervention of a moderator or requiring the presenter to navigate a detailed interface.

Violations of social norms

Several instances were described above where the deliberate violation of social norms regarding floor control was seen as a positive event. Most presentation software, however, enforces these social norms via rules embedded in the code. In other words, the software is written in ways that do not allow for the periodic violation of norms, even in cases where such violation might be beneficial. While it is certainly true that most presentation technologies have, a public chat stream that allows remote participants to "talk" any time they want, it is frequently the case that the presenter is not paying attention to this chat activity. Thus, the ability of remote participants to interrupt the lecture is substantially constrained. Nonetheless, there are times when interruption might be of value.

Principle #4: It should be possible for remote audience members to interrupt the presentation.

Enabling this capability, though, would involve several interesting tradeoffs. On the one hand, simply providing the additional functionality is not difficult. On the other hand, more study would be needed to better understand two attributes that would be substantively different for remote users. First is the extent to which interrupting in unison (as when several students simultaneously shouted to the presenter that they could not hear in the episodes described above) is important in face-to-face settings and could be replicated for remote participants. Second is the possibility that audience members' propensity to interrupt will be altered by being physically isolated from the larger audience. In other words, one wonders if people are less socially "guarded" when physically alone but in virtual company, than when sitting in a space full of people.

Even without such additional research, though, this is also a problem that could be addressed with a continuum of interruption capabilities (as described below) instead of the binary hand up/hand down system currently in use. One could easily imagine, for example, giving remote participants the capacity to interrupt the presenter's PC display or even the webcast audio stream to draw attention. This would be particularly useful in smaller classes, or in "clarification episodes" where there is a breakdown in shared understanding between the presenter and audience.

There are, of course, two significant risks inherent in this approach. One is that the capability will be abused by students, but one would hope that, particularly in university courses, this could be addressed via systems of accountability just as it is in face-to-face classrooms. The second is that remote participants will mischaracterize observed technical problems (e.g. bad sound, lack of video, etc.) as global in impact, when the problems are actually occurring only locally. In other words, an interruption to fix a problem might be mistakenly thought valuable to the rest of the class when it will really be disruptive because nobody else is experiencing the problem. This can be addressed, though, by establishing social norms whereby global interruption is a strategy of last resort for technical issues, after using, for example, chat functionality to speak with other students, a moderator or technical support staff.

Subtle communication

The data presented here clearly demonstrate the centrality of subtle communication, particularly with regard to hand-raising. This issue has two primary components.

First, there is the signaling issue. Raising one's hand is a signal of wishing to speak, but, in some environments, simply sitting there with a hand up can be a sign of rudeness—perhaps an implicit indication that one already knows what the speaker is going to say, or that one has something more important or interesting to say. Thus, several states of "hand raisedness" were observed that illustrate that the desire to speak is not a binary condition, but rather one that varies in terms of time and urgency.

Second, there is the issue of timing and interruption. As those who have studied interruptions in other environments [12] are well aware, timing is crucial. Situations are common where a sudden shift in discussion can make a recently crucial and relevant comment seem suddenly inappropriate. It was noted earlier that floor control almost inevitably shifted back to the instructor before going to another student, even during discussions. This suggests that instructors have a capacity for tracking those who want to participate and calling on speakers at an appropriate time.

Many presentation tools and case studies cited above, however, frequently take this control away from the presenter and give it to a moderator, who may or may not have the same sense of what audience members might say, or even when a good time is for interrupting the presenter.

From a design standpoint, this points once again to the importance of improving the presenter interface. A combination of the audience display wall used in Chen's virtual auditorium [5] and large display interaction techniques described by Vogel and Balakrishnan [23] could be quite useful in this regard. For example, motion tracking or a glove could be used to allow the presenter to point at and call on audience members displayed on a large wall at the back of the live presentation room.

It also suggests the value of extending systems to provide for a continuum of "hand raisedness" and to make it easier to realize that one's hand remains raised. One nice affordance of real hands in this regard is that they are attached to bodies that are constantly aware of their state.

The same is not true of interface sliders, which must be consciously updated and are frequently forgotten. As computer vision and motion tracking technology become cheaper, there may be significant value in using these to allow remote participants to actually raise and lower their hand when they wish to speak. All of this leads to:

Principle #5: Signaling one's desire to participate should be supported in ways that allow for variations in the intensity of expressed desire but,where possible, do not require constant, explicit updating of this parameter.

As with any design solution, there are tradeoffs here. While this may be a method that effectively captures the subtlety of hand raising observed in this study, it also alters the remote user experience by, for example, making multitasking more difficult. An upward reach to grab a book off a shelf could easily be misinterpreted as a desire to speak. Thus, different interaction techniques may be appropriate for different users and settings. In large lectures, subtlety may be less important, and multi-tasking may be more desirable than in smaller classes where participation may be more likely.

Presentation technology usage

Another issue of importance to designers is the usage of visual aids. While these results do indicate that static PowerPoint slides (the primary medium supported by presentation tools) comprise the vast majority of visual aids used by instructors observed here, the small number of usages of other media or advanced PowerPoint functionality (e.g. animation, timing, etc.) are well worth considering for two primary reasons.

Additional media are critical

First, these visuals did not appear to be ancillary to the lecture, but rather were an integral part of the instructor's ability to present materials. Videos were used to illustrate novel technologies and interfaces. Microsoft Word files were used to explain and define new terms as they were introduced. As described above, the timing and color features of PowerPoint were used to replicate cognitive psychology experiments and illustrate valuable points. And websites were frequently referred to and highlighted for the class to see. The obvious design implication here is to improve the support for additional media and file types, as well as their seamless integration into presentations.

Principle #6: A range of presentation technologies and media beyond static PowerPoint slides must be supported.

This presents an interesting design possibility that borrows from the notion of "virtual desktop managers." These tools allow users to leverage some of the benefits of having multiple screens by toggling placing windows on multiple "virtual desktops" that can be toggled on a single display. Presentation tools could have a similar functionality, whereby the output of certain "desktops" on the presenter's computer could go directly to the presentation software. Such an approach would allow for many media, allow for presenters to have multiple resources "cued up" and ready to go, and also secure the privacy of the presenter's primary desktop screen. Using "desktop" as the unit of display has the additional advantage that desktops being selected need not necessarily be from the same computer. Thus, it is possible to support student presentations, or have an assistant prepare visuals on a second machine.

Presentation style

In these results presentation style appeared to play a significant role in the way visual aids were used and the nature of these visuals. One instructor had a very "behind the scenes" approach to using the podium computer that appeared quite effective, but differed sharply from the "on stage" or "slide show" model that is typically used by presenters and assumed by software designers. More study is needed to see how widely this approach is used, but it does call for some support in webcasting tools for presenters to share this "behind the scenes" material with the audience. This could also be addressed through the "virtual desktop" scenario described above.

Improving on FTF

It is obvious that certain distractions and problems observed here in face-to-face presentations can be eliminated or improved upon in remote presentations. Microphones and headsets eliminate hearing troubles, and the isolation of audience members eliminates many additional sources of distraction. In addition, a chat stream can allow for dialogue and resource sharing during a presentation. Another advantage involves room control.

Even in face-to-face environments, presentation and room control technologies are rife with difficulties. Adding software to enable presentation to remote audiences typically adds another layer of complexity to an already difficult situation. Theoretically, however, there is no reason for remote audiences to be a late addition or afterthought. There would be significant value in considering the possibility of remote audiences in the design of presentation and room control technologies. If done well, this would vastly simplify the task of webcasting lectures to remote audiences, improve the experience of presenters trying to make technologies work, and also improve the experience of both local and remote audiences, who would spend less time waiting.

Principle #7: Webcasting should not be treated as an "add-on" to existing presentation and room control technologies.

CONCLUSION

Observations presented here revealed a range of classroom behaviors that extend understanding of the classroom setting from the unique perspective of system designers. By focusing on subtle aspects of communication and awareness, and improving on support for presentation technologies, it has been argued that the interactivity and, by extension, the potential effectiveness of webcast

presentations can be improved. Moreover, such improvements can be made without sacrificing the advantages of webcasting over conferencing.

ACKNOWLEDGMENTS

This work was supported by the NSERC-funded NECTAR research network. I wish to thank Ron Baecker for his ideas and feedback that were central to this study, as well as Dan Horn and Mike Massimi for comments on earlier drafts. I also thank the instructors I observed for allowing me into to their classrooms.

REFERENCES

1. Ackerman, M. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human Computer Interaction 15,* 2/3 (2000), 181-203.
2. Baecker, R., Baran, M., Birnholtz, J., Chan, C., Laszlo, J., Rankin, K., Schick, R. and Wolf, P. Using VoIP to Enhance Interactivity in Webcasts. To appear in *Proc. CHI 2006,* ACM Press.
3. Baecker, R. A principled design for scalable internet visual communications with rich media, interactivity, and structured archives. In *Proc. CASCON*, (2003), 16-29.
4. Brotherton, J.A. and Abowd, G.D. Lessons learned from eClass: Assessing automated capture and access in the classroom. *ACM Transactions on Computer-Human Interaction 11,* 2 (2004), 121-155.
5. Chen, M. Design of a virtual auditorium. In *Proc. Multimedia* 2001, ACM Press (2001), 19-28.
6. H.H. Clark *Using language.* Cambridge University Press, New York, 1996.
7. Isaacs, E., Morris, T. and Rodriguez, T.K. A forum for suppporting interactive presentations to distributed audiences. In *Proc. CSCW 1994*, ACM Press (1994), 405-416.
8. Jancke, G., Grudin, J. and Gupta, A. Presenting to Local and Remote Audiences: Design and Use of the TELEP System. In *Proc. CHI 2000*, ACM, (2000), 384-391.
9. Johanson, B., Fox, A. and Winograd, T. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *Pervasive Comp. 1,* 2 (2002), 71-78.
10. Malone, T.W. How do people organize their desks? Implications for the design of office information systems. *ACM Transactions on Office Information Systems 1,* 1 (1983), 99-112.
11. Mark, G., Grudin, J. and Poltrock, S. Meeting at the Desktop: An Empirical Study of Virtually Collocated Teams. In *Proc ECSCW 1999*, Kluwer Academic Publishers (1999), 159-178.
12. McFarlane, D.C. and Latorella, K.A. The scope and importance of human interruption in human-computer interaction design. *Hum. Comp. Int. 17,* 1 (2002), 1-61.
13. McKeachie, W.J. Research on college teaching: The historical background. *Journal of Educational Psychology 82,* 2 (1990), 189-200.
14. W.J. McKeachie *McKeachie's Teaching Tips.* Houghton Mifflin, New York, 2002.
15. Olson, G.M. and Olson, J.S. Distance Matters.*Human.-Computer.Interaction. 15,* (2001), 139-179.
16. Olson, J.S., Olson, G.M., Storrosten, M. and Carter, M. Groupwork close up: A comparison of the group design process with and without a simple group editor. *ACM Trans. on Information Systems 11,* 4 (1993), 321-348.
17. Rui, Y., Gupta, A. and Grudin, J. Videography for telepresentations. In *Proc. CHI 2003*, ACM Press (2003), 457-464.
18. Saroyan, A. and Snell, L.S. Variations in lecturing styles. *Higher Education 33,* (1997), 85-104.
19. Scardamalia, M. CSILE/Knowledge Forum. In *Education and technology: An encyclopedia,* ABC-CLIO, Santa Barbara, CA. 2004, 183-192.
20. Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S. and Suchman, L. Beyond the chalkboard: Computer support for collaboration and problem-solving in meetings. *Comm.s of the ACM 30,* 1 (1987), 32-47.
21. Steinert, Y. and Snell, L.S. Interactive lecturing: strategies for increasing participation in large group presentations. *Medical Teacher 21,* 1 (1999), 37-42.
22. Tang, J.C. Findings from Observational Studies of Collaborative Work. In *Computer-supported Cooperative Work and Groupware,* S. Greenberg Ed. Academic Press, London. 1991, 11-28.
23. Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large high resolution displays. In *Proc. UIST 2005*, ACM, (2005), 33-42.
24. Wallace, R.M. Online learning in higher education: A review of research on interactions among teachers and students. *Education, Communication and Information 3,* 2 (2003), 241-280.
25. Weinstein, I.M. Effective Enterprise Webcasting. Wainhouse Research, October, 2005.
26. Zijdemans, A., Moore, G., Baecker, R. and Keating, D.P. ePresence Interactive Media and Webforum 2001: An accidental case study on the use of webcasting as a VLE for early child development. In *International Handbook of Virtual Learning Environments,* J. Weiss, J. Nolan, J. Hunsinger and P. Trifonas Eds. Springer, New York. 2005, 1365-139.