

# An Empirical Study of the Use of Visually Enhanced VoIP Audio Conferencing: The Case of IEAC

Xianghua Ding,<sup>1</sup> Thomas Erickson,<sup>2</sup> Wendy A. Kellogg,<sup>2</sup> Stephen Levy,<sup>2</sup> James E. Christensen,<sup>2</sup> Jeremy Sussman<sup>2</sup>, Tracee Vetting Wolf<sup>2</sup>, William E. Bennett

<sup>1</sup>Dept. of Informatics, UC Irvine

<sup>2</sup>Social Computing Group, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Hts, NY  
dingx@ics.uci.edu, {snowfall | wkellogg | levy | ibmjim | jsussman | tlwolf | n7dz}@us.ibm.com

## ABSTRACT

IBM Enhanced Audio Conferencing (IEAC) is a VoIP-based audio conferencing system that, like several other systems, provides a visualization showing who is present and their states (e.g., speaking, muted). This paper presents the first study of the use of such a system. Drawing on log files collected over six weeks of use by over 1300 corporate employees, and interviews with 10 of them, we look at how and why various features of the system are used and what sorts of practices are supported. Our findings shed light on the factors that drive the use of visual enhancements to audio conferencing, and suggest further research topics.

## Author Keywords

Audio conferencing, conference call, social translucence, social proxy, social visualization, telephony, VoIP, voice

## ACM Classification Keywords

H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

## INTRODUCTION

Audio conferencing is becoming an increasingly important mode of collaboration. While telephony has long served as a major means of collaboration across distances (e.g., [13]), the advent of Voice over IP (VoIP) makes audio conferencing cheaper and more available. As a consequence, there has been a recent spate of work on ways of enhancing audio conferencing, including SmartPhone [11], MeetingCentral [23], and our own system, IBM Enhanced Audio Conferencing (IEAC).

But although attention has been focused on development, little has been done to understand whether, how and why

enhancements to audio conferencing work. Our aim in this paper is to examine how participants in enhanced audio conferences actually use the new functionality. To this end, this paper reports findings from a study of the reception of IEAC, a visually enhanced VoIP audio conferencing system that is tightly integrated with enterprise applications such as calendaring, directory, and instant messaging.

After discussing relevant literature and a describing IEAC, we analyze six weeks of log data to get an idea of the overall use of the system. Next we use interviews to obtain more insight on how and why the features are being used, and the practices they support. We conclude by discussing the findings, summarizing the factors that drive the use of the visual features of enhanced audio conferencing, and suggesting issues for future research.

## BACKGROUND

### Audio-only Conferencing

There has been quite a lot of innovative work on the audio-only front. Examples include Voice Loops [20], Thunderwire [7] and Mad Hatter [1]. Empirical work indicates that audio-only systems provide an effective medium for collaboration, supporting complex communication and coordination (e.g., [20]) and rich, sociable interactions (e.g., [7, 1]). However, audio-only systems are limited by a lack of visual cues to determine others' presence and availability, etc. [7, 23]. Problems include difficulties hearing speakers, identifying who is speaking, knowing who is listening, getting a turn to speak, etc. (see [23] for a good summary).

### Visually Enhanced Audio Conferencing Systems

An obvious approach to addressing the problems with audio-only systems is to add a visual channel. One approach is to add video. Work in videoconferencing and media spaces has shown that video is valuable for coordinating activity, identifying speakers, and enhancing and personalizing social interaction by offering mutual awareness, media presence and people images [e.g., 5, 21, 16, 2]. However, video has barriers to adoption such as high cost (e.g., bandwidth, setup), lack of mobility, privacy concerns, and scaling issues [11, 17].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2007, April 28–May 3, 2007, San Jose, California USA.  
Copyright 2007 ACM 1-59593-178-3/06/0004...\$5.00.

An alternative approach is to use visual abstractions, such as icons, to enhance audio conferencing. The advantage of this more abstract approach is that it requires less bandwidth, and the visual information can be presented much more compactly, thus offering the potential for visually enhanced calls to be supported on mobile phones and other small screen devices. On the other hand, this approach necessarily lacks the richness of video, and it is not clear whether the advantages afforded by video will carry over to this minimalist approach.

We divide visually enhanced audio conferencing systems into two categories: screen-oriented and voice-oriented. The former emphasize data sharing and co-editing using techniques such as screen sharing, shared whiteboard use, synchronized presentation, and so forth. Examples of this type of system include Sametime [14], Teamspace [18] and LiveMeeting [10]. Voice-oriented systems are focused on improving calling experiences, and while they use visual channels to enhance audio, they are designed so that use of a computer and its screen is generally optional. Examples include the Enhanced Telephony prototype [3], Impromptu [15], and audio conferencing systems such as SmartPhone [11], Meeting Central [23] and IEAC.

For the purpose of this paper, we focus on voice-oriented enhanced audio conferencing. One of the earliest examples is Moors' SmartPhone [11]. Moors suggested that audio-only conferencing was hindered by problems with control functions such as turn taking, and designed SmartPhone to include a visual display to allow callers to take turns, raise hands, and so on. However, aside from reports of a few observations during an initial trial on an unknown number of users for an unknown period of time, no evaluation of the system's use is provided.

Colburn et al. [3] contrasted conversations using one of two graphical depictions of participants (either icon-based or avatar-based) to voice-only conversations. Significant differences in subjects' judgments of the ease of determining who was speaking, when to talk, and who was listening were found, with the graphical interfaces making the judgments easier. In addition, users indicated a preference for the graphically-enhanced interfaces.

The project closest in approach to IEAC is Sun's Meeting Central [23]. Both use a shared visualization to provide presence and status information about the call's participants. Drawing on the concept of social translucence [6], both systems explore ways of providing information to allow participants to conduct a call coherently and create and enforce conventions. The Meeting Central work is notable for its extensive formative empirical work (surveys, interviews, and analysis of data such as trouble tickets) that provides valuable documentation of problems with audio conferencing, and for the comprehensiveness of its shared visualization which in addition to showing presence and status information, permits users to engage in textual or voice backchannels, raise hands, vote, etc.

However, although Meeting Central has been implemented and deployed, we know little about how the Meeting Central visualization is used. The first report on Meeting Central [23], reports preliminary user reactions from a limited deployment to three groups of 4 to 15 people – these consist of a few positive comments and feature requests. Subsequent empirical work by this group has focused on particular aspects of Meeting Central's functionality: the reception of private voice chats and their use relative to private textual chat has been explored in a short paper [22], and efforts to enhance audio conferencing sound quality are evaluated in [24]. Although development has continued to expand the features supported by Meeting Central's (now Meeting Suite's) visualization [8], there are no further reports on whether or how the visualization features are used in everyday audio conferencing situations.

### ABOUT THE IEAC SYSTEM

IEAC is a VoIP-based audio conferencing system with two user interface components: the *assistant* and the *meeting window*. The gist of its functionality is this. First, the *assistant* allows the user, in effect, to press three buttons to connect to his or her current conference call, without having to recall the call's number or passcode. Second, for users with access to a computer, the IEAC *meeting window* shows who is on the call, who is speaking, and provides access to call-related functionality such as muting.

In somewhat more detail, IEAC works as follows. A registered IEAC user dials a single number – using either a conventional or VoIP phone – and enters a single password (both, in practice, usually programmed into the user's phone and accessed via a button press). This connects the user to the IEAC *assistant* which uses information extracted from the employee's corporate calendar (used throughout IBM as the standard way of scheduling meetings) to offer the user a choice of conference calls, beginning with the current meeting. Upon selecting a meeting by pressing "1", the user is transferred into the call (regardless of whether it occurs on IEAC's VoIP bridge, or is a traditional conference call service offered by an outside vendor). If the IEAC user has also installed the *meeting window*, a visual component that runs on top of the corporate instant messaging infrastructure, it will pop up on the user's computer screen 10 minutes before the call's start.

Figure 1 shows the meeting window for a small conference call. At the top is a menu bar that provides access to various preferences, help information, and functionality (most of the person-oriented functionality is more easily accessed by right-clicking on the names below). Next is meeting information taken from the user's calendar. Jumping to the bottom of the window, we see a user information display area and a status line that shows the most recent event.

The central pane of the meeting window contains a *social proxy* for the conference call – a minimalist visualization that shows the states of the meeting's participants [6]. It shows those 'dialed into' the meeting (icons and names of

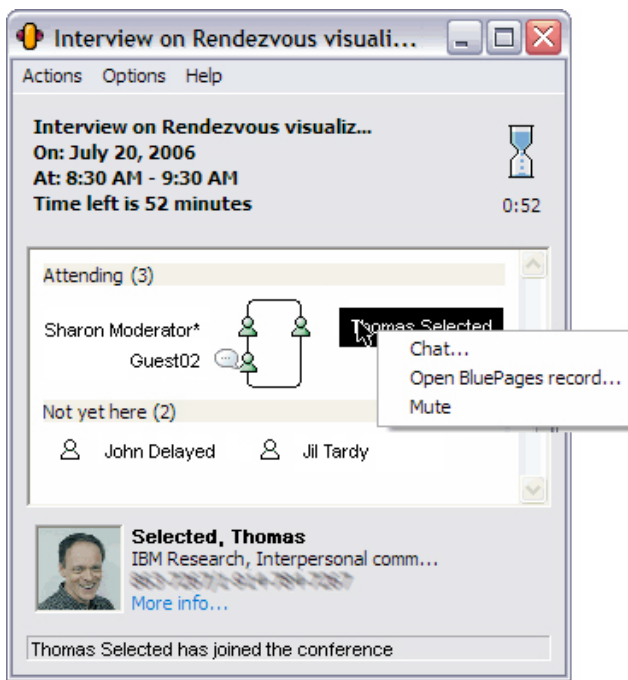


Figure 1. The IEAC meeting window. (“BluePages” is IBM’s corporate directory).

participants are arrayed around a ‘table’), as well as those who were invited to the meeting (according to the calendar entry) but have not yet arrived. Attendees who are not registered IEAC users (and therefore use a number and meeting-specific passcode to dial in) show up as guests, as in the case of “Guest02.” In addition to showing presence information, the conference call social proxy shows who is ‘speaking’— or, more accurately, which line(s) a signal is coming over — and displays a ‘speech bubble’ next to the appropriate icon(s); the icon also changes to indicate if the user is on mute or has disconnected. The conference call proxy also allows users to carry out actions. Clicking on another user displays a locally cached picture, job title, etc.; right clicking on another user’s icon provides a menu (as above) that enables a user to chat with, open the directory record for, or mute the other person.

### STUDYING THE IEAC DEPLOYMENT

IEAC is being gradually deployed within IBM. Currently there are over 1,300 registered users of IEAC who can set up calls and use its functionality. Those not registered can still dial in to IEAC calls, so the number of people dialing in is about four times that of the registered user base.

IEAC is being deployed via IBM’s Technology Adoption Program (TAP), which allows any interested employee to try out systems under development. This has two implications for our study. First, it means that IEAC is being used primarily by non-researchers who constitute the vast majority of IBM employees (a review of the registered users list shows that 7% are in Research). Second, it means that IEAC users are self-selected: they learned of IEAC and requested to be included in the trial. Thus, while IEAC is

being used by ‘real employees,’ the users are decidedly not ‘average users,’ but tend to be technically adventurous and/or heavy users of audio conferencing. In consequence, the results of this study cannot be taken as an indication of how IEAC will fare when more generally deployed. However, as Von Hippel and his colleagues have argued [19], studying “lead users” — those who are ahead of the curve in experiencing a new technology — can provide a valuable source of information on how the technology is being adopted and how to further develop it.

### Methods: Log Analysis and Semi-Structured Interviews

We used two approaches to study the deployment: we analyzed system logs, and interviewed IEAC users.

There were two sources of log data. The first was the IEAC *system log*: it logged all calls, their durations and participants, and various low level events. The second source was the *meeting window server request log*. This log was generated especially for this study, and spans about six weeks from early June to mid July. It captured requests sent from the meeting window to the server, resulting in a log of many (but not all) of users’ interactions with the meeting window. If an interaction required server access to proceed (e.g., IM, directory or calendar access, muting, etc.) it was logged; if it could be done locally by the meeting window on the client (e.g., showing a thumbnail picture of each participant in the call) it was not logged. Occasional server crashes resulted in (different) gaps in both logs, but with the exception of the week of July 4th, gaps are short and few.

The interview protocol was iteratively developed through pilot interviews with 6 members of a cross-organizational team involved in overseeing the IEAC deployment. The resulting semi-structured interview protocol covered three areas. We began by asking about the informant’s job, location and conference calling experience. The majority of the interview was devoted to inquiring about the use of each of the functional elements of the meeting window, probing for specific examples for each function. We ended by asking for feedback on how the system might be improved.

Because the focus of this study is on the ways in which users use the meeting window functionality, we restricted our sample to people who had used the meeting window recently. We generated a list of 116 people who had used the meeting window in the last two days, and randomly selected 10 names from it. Those who did not reply positively to the interview invitation were replaced with other random selections until we had arranged 10 interviews (the acceptance rate of our invitations was 33%). Because our informants were not local, we conducted all interviews using IEAC. Interviews typically lasted an hour, and were conducted by two or three people, one conducting the interview, and the other(s) taking notes; all were digitally recorded.




Figure 2. Calls per day (weekends omitted); box indicates meeting window logging period.

### RESULTS: THE USAGE LOGS

Figure 2 shows the calls per day handled by the IEAC system for the first 7 months of 2006 (these are all calls, regardless of whether the meeting window was used). The steadily increasing usage corresponds to the gradual addition of users. Although weekends are omitted, a weekly pattern with more calls midweek than on Monday or Friday is still visible. The box towards the right end of the graph indicates the period during which IEAC meeting window server requests were logged.

Figure 3 shows the number of meetings in which at least one user did something that required the meeting window to make a special request to the server. This provides a lower bound on the daily number of meetings in which the meeting window was used. That is, the estimate is conservative because the meeting window could provide value without interacting with the server: the window pops up automatically 10 minutes before each meeting, shows participants arriving, speaking and departing, and allows its user to select participants' names to reveal their pictures and job descriptions, all without server requests.

If we compare the highlighted portion of Figure 2 with Figure 3, we see two things. First, the meeting window is being used to access the server in only 15% - 25% of the meetings. Second, Figure 3 shows a gradual decline from an average of around 400 uses of the meeting window per day in the first week, to around 200 per day in the last week, even though the total number of IEAC calls IEAC is about the same. We will return to these observations in the discussion; for now we will note that even 250 meeting window uses per day is a significant body of data on which to draw.

We analyzed the *meeting window server request log* to see which server functions were used most frequently (Figure 4), and for each of the five most frequently used functions, selected 20




Figure 4. Relative use of meeting window server functions.




Figure 3. Calls per day (weekend omitted) in which the meeting window was used to access the server.

instances to examine their distribution relative to the start of the meeting.<sup>1</sup> The most frequently used were:

- **Get person information.** This means that a user used the meeting window to open another person's entry in the corporate directory. About half the instances of directory use occurred just after the meeting started.
- **Start/end calendar.** This means that a user opened or closed a "mini-calendar" which showed his or her upcoming meetings. About a third of the time a calendar "open" event was followed immediately by an "open the meeting window" event, suggesting that the calendar was being used to manually launch the meeting window.
- **Chat.** This event is logged when users use the meeting window to open a chat with another meeting invitee (chats initiated from a user's buddy list were not logged by IEAC). Chat showed no pattern relative to the meeting's start.
- **Set identity.** This happens when an identified user assigns a "guest" (a caller not registered with IEAC) an identity by right clicking their name. About half the instances of "Set Identity" occurred immediately after the meeting's start.
- **Mute.** A user can mute herself, another user or everyone else on the call. Over half the instances of mute occurred in the first half of the meeting.

While we had expected that chat would be a common use of the meeting window, the even greater frequency of the directory and calendar lookups, and the relatively high frequency of the set identity function surprised us.

In summary, log data showed that the IEAC system is supporting around 1500 calls a day, that the meeting window is being used to access the server in 15% to 25% of these calls, and which functions invoked through the meeting window are most frequently used. These data leave many questions unanswered. Why are corporate directory lookups, viewing one's calendar, and setting others' identity so often used? What is happening with respect to muting others, an action which seems potentially rude? How is chat being used in meetings? And, more generally, what value, if any, do users of the meeting window obtain?

<sup>1</sup> The log did not indicate when the call *actually* started. These analyses assume that calls start on the hour or half hour nearest the times each call's meeting window appears.

This brings us to the next section of our study, in which we turn to interviews with active users of the meeting window to understand how and why they are using it.

## RESULTS: INTERVIEWS

Our 10 informants, randomly selected from those who had recently used the meeting window, were generally heavy conference call users. This is not surprising, since users of IEAC have voluntarily adopted it, and since, as a global and highly distributed company, IBM contains a large number of people who work remotely, including a substantial percentage who work exclusively out of their homes (the IBM CIO's office estimates that on any given day, approximately 45% of IBM employees are operating from someplace other than a traditional corporate office). Of our 10 informants, 7 worked primarily out of their homes rather than offices (2 of whom were also frequent travelers), and 7 (not the same 7) played technical rather than managerial roles. It was not uncommon for our informants to report spending 4 to 6 hours a day on conference calls, with conference calls typically ranging from 5 to 10 people, though several reported occasional large scale calls ranging from 30 to 100 participants. To give a sense of the informant we're referring to, we will designate informants by #-{T(echnical) | M(anagerial)} {H(ome-based) | O(ffice-based)}; thus 3-TH designates informant 3, an employee who plays a technical role, working primarily from home. We occasionally quote our pilot interviewees (designated by P-#), all of whom are involved in managing IEAC's deployment.

We will discuss the findings from our interviews by dividing the conference call experience into three temporal segments: getting to the call; managing the meeting process; and facilitating the business of the call.

### Getting to the Call

Almost all informants noted that IEAC made it easier to get to their calls. Although joining a conference call may not seem difficult, frequent callers face a number of issues, including keeping track of multiple numbers and passcodes; remembering when to call (multiple time zones can cause confusion); and knowing when to dial in. These difficulties are exacerbated when a user has back-to-back calls, when a prior call is running over, if a call is during non-business hours, or if the caller is traveling – all circumstances that are not uncommon in a globally distributed business.

As described above, the IEAC assistant automatically connects callers to their calls. As this does not involve the meeting window, we will pass quickly over this, noting only that most of our informants stated that this was an extremely valuable function, some noting that the 'single number' functionality was sufficient reason for them to use the system. Perhaps the most significant indication of the utility of one number is that 9-MH, one of our most mobile users reported that he used IEAC for one to one calls: *"I use audio conferencing very extensively, even for two party*

*calls, simply because as highly mobile as I am, [and] as others are, you know, it's a place we know we can go to meet and count on [a] number that is working."*

The IEAC meeting window helped users get to their calls in two ways: it served as a reminder, and it also helped users decide precisely when they should join the call.

### Remembering when to call

The automatic pop up of the meeting window ten minutes before a call was appreciated by most informants. 7-TH's comment is representative: *"I \*love\* that it pops up ten minutes before my meeting, because I get so deep into my [work] that I'll lose track of time."* Another informant used this feature along with his calendar to create staged reminders: *"IEAC reminds me ten minutes before [a meeting] and I am like 'OK, I've got like five minutes.' Then [my calendar] reminds me at two minutes. So I realize, 'OK, seriously, I really need to get ready to take the next call' kind of thing."* [10-MO].

### Knowing precisely when to call

Once the meeting window pops up on users' screens, they will often monitor it to see when others join. As 10-MO said, *"If I'm joining a call and I look at IEAC, and nobody is in the room yet, I know everybody is running late, so maybe I don't have to rush too much... if I am running over [on] my own call, I don't have to rush to get into the next one when there is nobody else there either."* Similarly, 5-MO remarked, *"Sometimes if I'm late in joining a conference call ... I do bring up the visualization ... for instance, I may have been on another call or I am doing something I need to finish, I can join the call ... 10 minutes late, I ... bring up [the] visualization to see who else has joined, ... or to find a critical time to join the call."*

We experienced this phenomenon during the interviews, which we conducted using IEAC. We typically dialed into the meeting about five minutes early because we didn't want our informants to have to wait for us. However, many of them noticed our presence and immediately dialed in, apologizing for being late. Others double-clicked on our names and IM'ed us to say they'd be a bit longer.

### Managing the meeting process

Once users have dialed into the call, the next stage of activity involves getting the meeting off to a good start. This role is often taken up by the moderator, the person who sent the meeting invitation and thus set up the call. Activities include deciding when to start, monitoring who is present, and managing sound quality.

### When to start: announcing arrivals, summoning latecomers

The decision of when to start the meeting is often a function of how many people have arrived, or whether particular people have arrived. The conference call social proxy allows participants to see who is present, who just arrived and who is late, rather than the moderator having to repeatedly ask "Who just joined?" As one of our pilot

interview subjects explained, “*One of the most beneficial features of IEAC is that bit of information... as a moderator of a call it is nice for me not to say ‘who just joined?’ or having to go back to my ... calendar to see who are all I invited or to see who wasn’t there yet. From my visualization, I know who is not here.*” [P-4]

Surprisingly, the social proxy provides indirect value even for those who do not see it, because many moderators make a practice of monitoring the proxy and announcing who has arrived. 7-TH said, “*Several times there has been a delay between when I see [someone’s name] popping up and when they actually announce themselves, so I could announce to everyone, ‘Here’s so and so.’ As you know, on some of the calls, people might be late, especially with the Directors, [who] are coming from a prior call. Someone might say ‘Where is so and so?’ In that case, I might know ahead of time that they are joining the call, and be able to say ‘Here they come.’*”

The conference call proxy also facilitates coordination via chat. A typical use is for the moderator to open a chat with someone who has not yet arrived. 3-TH reported “*Actually that is one thing that the visualization does help with: You can very easily figure out who is not on the call yet, you know, you have that whole list of people who are not here yet... if they were critical to the call, I would probably ping them and say, you know, ‘Are you joining the call?’*” Similarly, 6-TH said, “*Let’s say somebody hasn’t joined the call yet and I’m expecting them to be here because they have accepted [the invitation], then I select the name and double-click on it, and IM them directly from there to say ‘Hey, are you joining the meeting?’ or whatever. Instead of having to go to look up [in the corporate directory] or whatever, so that is helpful.*” By providing ‘communication handles’ for all participants in the meeting, the social proxy makes coordination via chat a much lighter weight process; this is particularly useful because moderators are often dealing with multiple issues when starting a meeting.

#### ***Monitoring the audience: ‘No sneaking in on a call’***

In addition to making sure that invited people are present, another reported use of the conference call proxy is to ensure that others are not present. Sometimes, especially in large meetings, people are supposed to attend for one part of the call, and then leave. 7-TH said: “*One thing I find interesting is that you can’t hide... [If] you’re on, you are on. No sneaking in on a call. ... There are times ... when some of the participants... the topics may have finished, but participants still want to stay on the call to discuss the particular item or issue, and there are times when you can not be one hundred percent sure that everyone has left the call (when they) should have left the call, but with IEAC, when everyone is using IEAC, you know who is disconnected.*” While many informants weren’t very concerned with confidentiality, those who were – for example consultants who assemble bids and vet contracts – were very concerned.

Even when the content of the call is not confidential, people appreciate knowing who is there. 6-TH commented “*You have different kinds of conversations with different groups of people, and if people can lurk without the visualization, it’s always a concern: do you really know who everyone is there, right? I may add some information that I’m willing to give to some people but not others, or I may even behave differently with some people than others.*” Of course, as participants knew, the conference call proxy does not provide complete surety as it does not reflect the situation in which collocated groups dial in using a speaker phone.

The concern with who is present also explains why setting identity is among the most frequently used functions. People who attend a call, but are not registered with IEAC, show up in the conference call proxy with the label “Guest.” “Set identity” allows any registered IEAC user to right click on the “guest” name and assign it an identity from the set of not-yet-here invitees (if an employee), or a label (e.g., “CIO, Company X”) if from outside IBM. Although this is not foolproof, as it relies on a person to assign the identity, it appears to be sufficient for normal purposes. A number of our informants mentioned that they customarily assigned identities to guests as they showed up at the start of the meeting, which is consistent with the log analysis showing that about half the uses of “set identity” occur immediately after the start of the call.

#### ***Managing sound quality***

In addition to ensuring that meeting participants are present and accounted for, good sound quality is important for a successful call. While sometimes this is only amenable to a technical solution, often it can be socially managed.

Problems with audio quality may be due to a participant calling in from a noisy place, a poor cellular connection, or a participant who breathes heavily while hurrying through an airport (not to mention the occasional person who falls asleep and begins to snore on a cross-time zone call). The meeting window’s social proxy supports dealing with audio quality problems in two ways: first, it makes it easier to detect the source of the problem; second, it provides a means of control. Several informants reported diagnosing the source of annoying background noise by watching the conference call proxy and correlating the appearance of the ‘speech bubble’ with the noise. Once the noise was identified, the person could be informed, or if necessary, another person could mute him or her. Muting another person is useful when the person is using a telephone without a mute function, when they are asleep, when the caller has temporarily left the phone (and background noise occurs), or when a person has put the call on hold (not realizing that her telephone service plays ‘on hold’ music). About half our informants reported putting others on mute, or doing so and then IM’ing the person muted to tell them.



### Facilitating the business of the call

Once conditions necessary for the meeting to proceed are in place, the business of the call begins. While IEAC's meeting window does not directly facilitate the discussion of content, it plays three roles in supporting the meeting: aiding speaker identification, providing a means of gathering ancillary information about participants, and serving as a ready to hand within-meeting backchannel.

#### Speaker and participant identification

As expected, given the difficulty callers have in determining who is speaking [23], our informants found the association of the speech bubble with a person's name very helpful. 6-TH said "[The] most frequent time I look... is when I hear someone talking and either I don't recognize their voice, or I want to identify their voice. So at the beginning of the meeting, I use it a lot – I identify people as they're joining the meeting. Subsequent to that, I look at it kind of idly to see who's talking." This is reinforced by the fact that several informants grumbled about the lag between speech and the appearance of the speech bubble (the lag is a calculated choice to reduce system load). Informants also reported that they used the meeting window to keep track of who was on the call. 7-TH: "Once a month I present the executive report to my Director as well as a Business Director)... and it's very helpful to me to know who's joining, and [know] that they've joined the call as opposed to waiting for [them] to announce themselves."

The conference call proxy also helped callers remember who was in the meeting later. Both 4-TO and 6-TH reported that they captured screenshots of the meeting window attendance list to include in their meeting notes. Other informants reported keeping instances of the meeting window open after the meeting ended to remind themselves of the call, and of who had been present. P-6 reported that he leaves the windows up "to make sure that I get the names right when I'm doing minutes of meetings."

#### Gathering ancillary information about participants

While we expected informants to find the speech bubble useful, as it helps to determine who is speaking, we were surprised by the degree to which informants used the social proxy to obtain ancillary information about other callers.

About half our informants said that the pictures that appear when a participant's name is selected in the meeting window were important. In terms of practical value, some commented that pictures served as memory aids. 1-TH said "Even [if] I don't necessarily ... ever see that person face-to-face, I can usually associate someone's face better than with someone's voice." In addition, some informants reported that the pictures allowed them to address another person by name, either on the call or when they finally encountered them face to face. Some informants were quite passionate about the value of pictures, noting that they were especially important for those who worked from home. 1-TH declared: "As a work from home employee I'm a big fan of [corporate directory] pictures because I almost never get

to see everyone face to face..." 3-TH agreed, saying "I work from home and don't get to interact with people in person very often. It's nice to see what people look like; there are people I've worked with for years that I have no idea what they look like until [the corporate directory] started putting pictures up." 10-MO: "I think it personalizes it. In a lot of cases you end up having phone calls with people who you never met in person...you know, this call has been a perfect example, I've never met either of you, the odds of meeting you are pretty slim, and yet I feel like I have a better sense of connection to you and the work you are doing just by looking at the picture of each of you. It makes it more personal; it is more human."

According to the meeting window server request log, the most heavily used function of the meeting window was to bring up the corporate directory page. Although simply selecting a name provides a picture and brief job description (unfortunately the frequency of this was not logged), users often took the next step of bringing up BluePages, the corporate directory. BluePages provides far more detail, including job descriptions, lists of projects, and the person's co-workers, management chain, and organizational location. Most informants reported that they used it to learn "what people's roles are," "who they report to," or "where people 'fit' or 'belong' in the organization."

Why should people be so concerned about each other's roles? Our data suggest that such ancillary information helps to interpret others' input. Several informants reported checking out a person's directory information when his or her comments were interesting (or weird). 8-TH said "This person's talking and I get curious and look up who they are and what they do. Sometimes, I am curious: OK, these people are weird, what motivations there are, who do they report to – that way, I get some perspective of where they are calling from." In addition, people also determine what and how to speak based on their knowledge of their audience. As 7-TH put it, "I wanted to know what their role was, because I don't like to discuss metrics unless I am certain that the people participating are the correct representatives." Some informants reported looking up people even before the meeting started, as a way to prepare.

#### A ready-to-hand within-meeting backchannel

As Figure 4 shows, chat – and specifically chat initiated by clicking on a name in the meeting window – is one of the most frequent uses of the social proxy (see [9] for a more complete discussion of its use as a backchannel). We have already noted its role in managing the meeting process by permitting the summoning of latecomers, and as a means of informing others that they need to mute themselves.

Informants described a variety of backchannel chat behaviors, often noting that they preferred to privately consult a colleague than disrupt the meeting. Sometimes the ability to speak privately was valued for saving face – for example, to avoid appearing as though one hadn't been listening or hadn't fully understood the conversation.

Informants also described tag-team behavior, in which chat allowed one person to act on another's behalf. 3-TH described playing a role as a backup to someone doing a presentation: "...so if someone asks a question that the presenter doesn't know the answer to I'll use [IM] to find out the answer while they're still on the call." In a less cooperative example, 1-TH describes how he used chat to 'encourage' another person to provide him with information that he could tell he would soon be asked for: "So I'll ... ping the person that was supposed to respond to me [prior to the call] and hasn't yet. Hey did you read my email? Give me an answer now! That would be good because probably in five or ten minutes I'm going to be asked a question and I'm going to have to ask you [publicly] anyway."

More generally, informants spoke of using chat during a call to make sure that they were "on the same page" as their colleagues. 4-TO: "I feel free with my management to express my opinion when it's just them and I, but when everybody's on the call it is best to reserve those opinions... a lot of time those will be the [chat] conversations that will be going on ... conversation with my management to express that I don't agree with them." Similarly, 5-MO said "It can be useful to be able to have a private discussion about the merits of what someone is saying, and to confer with team members... generally my habit in conference calls, especially if it is with multiple organizations ... there are things we talk about on the call, and then we use IM in trying to arrive at an understanding, or sharing info that can't be openly shared on the call." It is useful to remember that meetings aren't just among people but among organizational units. 7-TH: "We've had a few cases where we had [division X] folks on a call, and we were meeting with people in [division Y] who had to sign off on [X's] requests. We might have a few folks sending IM messages back and forth to answer questions, to be prepared to answer questions that [Y] was asking, etc. We'd IM back and forth to get info across, or to remind folks to ask specific questions."

## DISCUSSION

### *Practices facilitated by IEAC*

Our interviews reveal a rich set of practices that are facilitated by IEAC. The process of getting to the meeting is facilitated by the assistant, with its single number for all calls, and by the meeting window's appearance ten minutes before the call's start. In addition, the meeting window's social proxy provides a new capability in allowing a not-yet-dialed-in caller to 'watch' as people arrive, and to thus gauge exactly when to join the call. The conference call proxy is also important in managing the call and getting it off to a smooth start. Call moderators find it useful in monitoring the arrival of participants, and its chat functionality provides a ready-to-hand means of summoning the tardy. It also provides new ways to ensure that the conditions are right for the call to get under way. On the social side, the conference call proxy makes it easier

to detect lurkers who shouldn't be present, or who were supposed to leave; on the technical side, it permits the diagnosis of the source of annoying background noises, and provides social and technical means (private chat with the noise source, or simply muting the noisy line) for remedying the problem. Finally, the meeting window facilitates the business of the meeting, showing who is speaking, enabling users to gather information about others, and allowing users to coordinate with one another (either by taking cues from who is speaking, or via private chat).

### *Many uses by a sub-critical mass of users*

There is a seeming tension between the results of our log analyses and our interviews. The people we interviewed are enthusiastic about IEAC, and report a multitude of ways in which the meeting window makes their daily lives easier. On the other hand, the meeting window is only used to access the server in 15% to 25% of the calls. Furthermore, over the six week study period, the number of uses of the meeting window per day has declined from an average of about 400 to 200.

These results are not as divergent as they seem. One explanation for the limited use of the meeting window is, quite simply, that it isn't always useful. As several informants noted, the conference call proxy is less useful in small recurring calls in which you know the other callers – once people have arrived, speaker ID, coordination and backchanneling are less important than in larger calls with strangers. A second point is that, according to the system log, a maximum of 25-30% of those who currently dial in to IEAC calls are registered (and thus identified); non-registered users can participate in the audio portion of the call, but can not see or use the meeting window. In spite of the fact that some moderators use "Set Identity," on the average, only a minority of participants will be identified, and the social benefits of the proxy – seeing who is present, talking, etc. – are decreased. It will be interesting to see if usage of the meeting window increases as the deployment continues and the percentage of registered users rises. A third point is that many ways of using the meeting window – watching the arrival and departure of users, identifying speakers, and clicking on names to see pictures, etc. – are not counted as 'uses' (because they don't access the server and thus aren't logged).

There are a number of explanations for the decline in usage over the study period. First, it may be that callers abandoning the use of the meeting window due to the lack of a critical mass of users, or for other reasons. Several of our informants commented that the meeting window was less useful when there were few registered users. However, there are at least three temporary factors that may have contributed to the decline. One is that a significant change in the operation of the meeting window occurred during the study period. The earlier version of the meeting window was designed so that it popped up on the screen of any user whether or not they were registered (this was intended to



promote ‘viral spread’). However, as this also caused confusion, the pop-up behavior was disabled for non-registered users in a new version that was gradually downloaded by the user population in July. This shift eliminated uses of the meeting window by non-registered users. A second temporary factor was occurrence of a bug that caused intermittent failure of the meeting window to pop up automatically, meaning that users would have to manually launch it. A third temporary factor is that the end of the second quarter occurred midway through the study period. Quarter endings are characterized by large cross functional meetings, which (as some of our informants remarked) are those in which the meeting window is most useful for retrieving information about unknown participants. That the decline was largely due to temporary factors is buttressed by a recent spot check of the meeting window usage rate that shows about 175 touches per day, for the eighth week after the end of the study period. This represents a slight decline, but had the drop seen in the study period continued linearly, meeting window use should have ceased by this point.

The enthusiasm of our informants for the meeting window should not be a surprise. Since our informants were selected from the set of people who had recently used the meeting window, our sample was biased towards frequent and active users, and this in turn means that they were likely to have meetings that included other registered users (you can’t chat or retrieve information about others if they are not identified). This selection method was deliberately chosen to provide a look at how the meeting window might be used when there is a larger percentage of registered users.

#### *Drivers of how and why the system is used*

In reflecting on our findings we see four factors that drive the use of IEAC’s assistant and meeting window.

- **Reduction in cognitive load.** IEAC makes many currently possible things easier. It eases the burden of remembering numbers and passcodes, makes meeting invitee’s IM handles easily accessible, and facilitates seeing who is present. It is easy to dismiss these as minor improvements, but we suggest that when considered in context – multi-tasking users, under time pressure, in a cognitively demanding situation in which errors are public and potentially embarrassing – they constitute a major shift in the caller’s experience.
- **Support for control and coordination.** The conference call proxy makes control and coordination easier by making the state of the meeting visible, and by providing ready to hand channels through which communication and control are exercised. This supports central control by the moderator (e.g., summoning a tardy invitee via IM), coordination among colleagues (e.g., by using chat as backchannel), and collective coordination (e.g., when people dial in because they see that others have dialed in or decide not to lurk because they’re visible to others).
- **Gathering ancillary information.** Another important use is signaled by the popularity of accessing detailed

information about others’ roles and where they fit in the organization, rather than the just using the summary information provided in the meeting window. This highlights the fact that people benefit from understanding others’ contexts, and that meetings are interactions between organizations, as well as individuals.

- **Affiliation.** Finally, we were struck by the strong attachments our users, particularly remote workers, felt to pictures. This may become increasingly important as remote work increases in quantity and frequency.

#### **CONCLUSION**

The results of our study leave us cautiously optimistic that there is a future for visually enhanced audio conferencing systems like SmartPhone [11], Meeting Central [23] and IEAC. Our study indicates that the value of visual enhancements lies both in making activities that are already possible easier, and in enabling activities not previously possible in audio conferencing. The ability of an individual to dial in at just the right moment by keeping an eye on the meeting window, or to open the corporate directory to find out more about the speaker making the ‘weird’ remarks, or the ability of representatives of division X to use a private backchannel to coordinate their response to division Y, are all means of making the meeting process more effective. These are examples of what Nardi and colleagues have termed ‘outeraction:’ “communicative processes outside of information exchange, in which people reach out to others in patently social ways to enable information exchange” [12]. But while they have positioned outeraction *vis a vis* informal communications, we note that it also plays a role in the more formal meetings we report on here.

While conference calls seem formal, with their scheduled times and invitation lists, participating is not a simple matter of just dialing in and talking. As this study has shown, a lot of work is necessary: from the basic logistics of restoring, managing, and retrieving call information, to more complex social coordination to ensure critical presence; from the creation of a desirable and confidential call environment, to the collective construction of face when different organizations interact. All of this becomes more difficult as the number of conference calls increases, as people become more mobile and distributed, and as more organizations and strangers are involved in calls. One of the most interesting lessons we take away from this research is how much work is going on behind the scenes – by moderators, individuals, and organizationally defined sub-groups – to create a coherent and productive meeting.

This work raises a variety of other questions that are ripe for future work. One area that seems ripe for work is to take a closer look at “formal” meetings. While the field has by no means ignored more formal meetings (e.g., see [13]), especially if they involve design or programming teams, it seems to us that there has been a very strong bias towards studying informal, unscheduled interactions. We believe that the domain of ‘formal’ meetings offers a rich terrain to explore. A second area of interest involves the home

teleworker. We've been struck by the difference in experience reported by those who work full time out of their homes. While there is a large literature on telework, much of it takes an organizational or managerial perspective (e.g., asking what individual or organizational characteristics are likely to lead to successful telework), rather than asking how to support the individual. In this vein, one might inquire about the role and importance of social cues like pictures, and ready access to profile information, on remote workers' ability to collaborate with their colleagues. Finally, our interest in understanding and designing enhanced audio conferencing systems continues. Both of these areas – formal meetings and home workers – offer interesting possibilities for driving the further design of enhanced audio conferencing.

#### ACKNOWLEDGMENTS

Thanks to Jonathan Brezin for timely assistance.

#### REFERENCES

1. Aoki, P. M., Romaine, M., Szymanski, M. H., Thornton, J. D., Wilson, D., and Woodruff, A. The mad hatter's cocktail party: a social mobile audio space supporting multiple simultaneous conversations. In *Proc. CHI 2003*, ACM Press (2003), 425-432.
2. Bly, S., Harrison, S., Irwin, S., Media spaces: bringing people together in a video, audio, and computing environment. In *Communications of the ACM*, 36, 1, (1993), 28-46.
3. Cadiz, J.J., Narin, A., Jancke, G., Gupta, A., and Boyle, M. Exploring PC-telephone convergence with the enhanced telephony prototype. In *Proc. CHI 2004*, ACM Press (2004), 215-222.
4. Colburn, R. A., Cohen, M. F., Drucker, S. M., Tiernan, S. L. and Gupta, A. Graphical Enhancements for Voice Only Conference Calls. *Microsoft Research Technical Report*, MSR-TR-2001-95, October 1, 2001. <http://research.microsoft.com/research/coet/Graphics/TRs/01-95.pdf>. Accessed September 24, 2006.
5. Dourish, P., Bly, S. Portholes: supporting awareness in a distributed work group. In *Proc. CHI 1992*, ACM Press (1992), 541-547.
6. Erickson, T. and Kellogg, W.A. Social translucence: An approach to designing systems that mesh with social processes. *ACM ToCHI*, 7, 1, (2000), 59-83.
7. Hindus, D., Ackerman, M. S., Mainwaring, S., and Starr, B. 1996. Thunderwire: a field study of an audio-only media space. In *Proc. CSCW 1996*, ACM Press (1996), 238-247.
8. Kaplan, J. and Yankelovich, N. Sun Labs Meeting Suite, Executive Edition. Available at <http://research.sun.com/projects/mc/hcic-boaster.pdf>. Accessed September 15, 2006.
9. Kellogg, W. A., Erickson, T., Vetting Wolf, T., Levy, S., Christensen, J., Sussman, J., and Bennett, W.E.. Leveraging digital backchannels to enhance user experience in electronically mediated communication. In *Proc. CSCW 2006*, ACM Press (2006), 451-454.
10. LiveMeeting. <http://office.microsoft.com/livemeeting/>. Accessed September 26, 2006.
11. Moors, T., The Smartphone: Interactive group audio with complementary symbolic control. In *Revised Papers From the 4th international Workshop on Distributed Communities on the Web, Lecture Notes In Computer Science, 2468*, Springer-Verlag (2002), 100-112.
12. Nardi, B. A., Whittaker, S. and Bradner, E. Interaction and outeraction: instant messaging in action. In *Proc. CSCW 2000*, ACM Press (2000), 79-88.
13. Olson, G.M. and Olson, J.S. Distance matters. *Human Computer Interaction*, 15, 2&3 (2000), 139-178.
14. SameTime. <http://www.ibm.com/lotus/sametime>. Accessed September 26, 2006.
15. Schmandt, C., Kim, J., Lee, K., Vallejo, G, Ackerman, M. Mediated voice communication via mobile IP. In *Proc. UIST 2002*, ACM Press (2002), 141-150.
16. Sellen, A. Remote conversations: The effects of mediating talk with technology. In *Human Computer Interaction*, 10, 4 (1995), 401-444.
17. Tang, J. , Isaacs, E.. Why do users like video? Studies of multimedia-supported collaboration. *Technical Report TR-92-5*, Sun Microsystems (1992).
18. TeamSpace. <http://www.teamspace.com/>. Accessed September 26, 2006.
19. Von Hippel, E. *Democratizing innovation*. The MIT Press, Cambridge, MA, USA, 2005.
20. Watts, J. C., Woods, D. D., Corban, J. M., Patterson, E. S., Kerr, R. L., and Hicks, L. C. 1996. Voice loops as cooperative aids in space shuttle mission control. In *Proc. CSCW 1996*, ACM Press (1996), 48-56.
21. Wellman, B, Experience in the use of Media Space. In *Proc. CHI 1991*, ACM Press (1991) 203—209.
22. Yankelovich, N., McGinn, J., Wessler, M., Kaplan, J., Provino, J., and Fox, H. Private communications in public meetings. *Ext. Abstracts, CHI 2005*, ACM Press (2005), 1873-1876.
23. Yankelovich, Walker, W., Roberts, P., Wessler, M., Kaplan, J., and Provino, J. Meeting Central: Making distributed meetings more effective. In *Proc. CSCW 2004*, ACM, (2004), 419-428.
24. Yankelovich, N., Kaplan, J., Provino, J., Wessler, M. and DiMicco, J. M. Improving Audio Conferencing: Are Two Ears Better than One? *Proc. CSCW 2006*, ACM, (2006), 333-342.