Hanging on the 'Wire: A Field Study of an Audio-Only Media Space

Mark S. Ackerman[1]
Debby Hindus[2]
Scott D. Mainwaring[2]
Brian Starr[1]

[1]Information and Computer Science
University of California, Irvine
Irvine, CA 92717
{ackerman, bstarr}@ics.uci.edu

[2]Interval Research Corporation
1801 Page Mill Road
Palo Alto, CA 94304
{hindus, mainwaring}@interval.com

Contact person: Mark Ackerman

Hanging on the 'Wire: A Field Study of an Audio-Only Media Space

Mark S. Ackerman
University of California, Irvine

Debby Hindus
Interval Research Corporation

Scott D. Mainwaring
Interval Research Corporation

Brian Starr
University of California, Irvine

The primary focus of this paper is an analysis of an audio-only media space from a Computer-Supported Cooperative Work (CSCW) perspective. To explore whether audio by itself is suitable for shared media systems, we studied a workgroup using an audio-only media space. This media space, called Thunderwire, combined high-quality audio with open connections to create a shared space for its users.

The two-month field study provided a richly nuanced understanding of this audio space's social use. The system afforded rich sociable interactions. As well, users were able to create a useful, usable social space; however, through an analysis of the social norms that the participants formulated, we show that they had to take into account being in an audio-only environment. Within the field study, then, audio by itself was sufficient for a usable media space and a useful social space, but users were forced to adapt to many audio-only and system conditions. The paper also considers audio's implications for privacy.

Categories and Subject Descriptors: H.1.2 [Models and Principles]: User / Machine Systems--human factors, human information processing; H.4.3 [Information Systems Applications]: Communication Applications; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems--audio input / output,

evaluation / methodology; H.5.2 [Information Interfaces and Presentation]: User Interfaces--evaluation / methodology, interaction styles; H.5.3 [Information Interfaces and Presentation]: Group and Organizational Interfaces; J.4 [Social and Behavioral Sciences]

Additional Key Words and Phrases: Audio, audio spaces, media spaces, electronic social spaces, social presence, speech interactions, mediated communication, computer-mediated communication, CMC, telepresence, social interactions, rich interactions, norms, privacy, computer-supported cooperative work, CSCW.

A preliminary report of this study was presented in the CSCW'96 Conference Proceedings.

Authors' addresses: M. Ackerman and B. Starr, Information and Computer Science, University of California, Irvine, Irvine, CA, 92697; email: ackerman@ics.uci.edu; D. Hindus and S. Mainwaring, Interval Research Corporation, 1801 Page Mill Road, Palo Alto, CA, 94304; email: hindus@interval.com

---

1. INTRODUCTION

Media spaces—systems that use various media technologies to support collaboration—offer an intriguing possibility [Bly, Harrison, and Irwin 1993]. These media systems have an extraordinary potential for creating new types of social spaces, as Meyrowitz asserts:

> ...The introduction and widespread use of a new medium of communication may restructure a broad range of situations and require new sets of social performances. ([Meyrowitz 1985], p. 39)

Each type of media and by extension, each type of media space system, may allow people to interact in ways quite different from those occurring in face-to-face situations, or even in other types of media spaces.

However, considerable research is still required to allow these systems to fulfill their promise. Relatively little is known about what the right mix of media might be for developing useful and usable social spaces [Olson, Olson, and Meader 1995]. Almost all media spaces have used video and audio together; however, audio by itself, if usable, would be attractive because of its lower complexity and cost. As well, numerous prototype systems have demonstrated that audio is a particularly intriguing medium for collaboration and interaction (e.g., [Hindus and Schmandt 1992, Schmandt 1994, Schmandt, Ackerman, and Hindus 1990, Yankelovich, Levow, and Marx 1995]).

In this study, we explored the potential of audio by observing and analyzing a workgroup's use of an audio-only media space. This media space system, called Thunderwire, combined high-quality audio with open connections to create a shared space for its users.

Our two-month field study provided a richly nuanced understanding of the social uses of this audio space. The space was lively, sociable, and usable, as one user reported:

> Hanging on the 'Wire.... I think that the way I am on Thunderwire is the way that I would be if we were out at lunch or dinner or just hanging out. It's not an easily attainable state in the normal kinda-pass-people-in-the-hall [company]. -- Patty

However, the users were forced to adapt to many audio-only conditions in order to use the system. This paper presents the results of Thunderwire's use during the field study.

This paper begins by examining the literature that supported and informed the design of Thunderwire as an audio-only system. After arguing for the utility of the system and its affordances, we then describe the important characteristics of the working system along with the specifics of the field study. The majority of the paper analyzes Thunderwire's use from a Computer-Supported Cooperative Work (CSCW) perspective. The first part of this analysis examines the interactions among Thunderwire participants, and the second part examines the norms that the participants formulated to govern their use of the system. We follow this analysis with a speculative consideration of audio's implications for privacy, and we close with a discussion of the design and research implications resulting from this study.

2. AUDIO MEDIA SPACES

The central premise of this study is that audio by itself might be suitable for media system use. The potential of an audio-only media space intrigued us for a number of design, theoretical, and empirical reasons. From a design viewpoint, an audio-only system is considerably less complex, and therefore more practical, than one that includes video. Furthermore, there is already a noticeable example of a successful audio-only system, the telephone. Indeed, potential users already have experience with audio-only

communications, due to the ubiquity of the telephone. As well, there are theoretical and empirical reasons, outlined below, why audio alone might be suitable for a shared media system.

Despite the potential, to our knowledge there have been few field studies of audio-only media spaces, and no field studies of systems with good-quality audio. The remainder of this section summarizes the empirical and theoretical findings that led us to believe that such a study would be informative. We present support for an exclusively audio media space, the potential for such a system to result in a mediated social space, and the specific audio characteristics needed for such a system.

2.1 Audio-only media spaces

There is substantial empirical support for the argument that audio alone is sufficient for a viable media system. Sellen [Sellen 1995] provides a comprehensive summary of the numerous studies comparing audio with video and with face-to-face interactions; here we will survey only a representative handful.

In general, findings have either pointed to the primacy of audio in communication, or the results have been mixed. Starting from Chapanis' [1975] comparison of media, audio has been found to have a primary role in communication. (See Rutter [1987] for a summary.) In Chapanis' work, for example, task completion was found to be primarily dependent on having an audio channel; the inclusion of video was significant only in tasks requiring negotiation.

These results have been confirmed in recent studies of computer-mediated communication. In her experiments, Sellen examined such conversational mechanisms as number of turns, turn duration, turn distribution, percentage of simultaneous speech, conversational switches, and types of simultaneous speech under face-to-face, video, and audio-only conditions. She found that:

> In these two experiments, the *nature* of the mediating technology, whether it was a videoconferencing facility or an audio-only system, had no significant effects on these kinds of measures. This suggests that, with respect to this set of dependent variables, if a conversation takes place via a technological system, the way in which

> the visual information is presented, or even whether visual information is present, makes no difference to the conversational process. ([Sellen 1995] p. 433)

Sellen, however, continues by noting that only ten percent of her subjects preferred the audio-only condition.

More mixed results were reported by Olson et al. [1995] who studied face-to-face, video, and high-quality audio conditions for distributed meetings. Their findings suggest that the quality of the task was lower in the audio-only condition than in face-to-face. Again, users preferred the video condition over audio-only. The authors suggested that users were responding to hidden process requirements. The latter suggestion was also made by Isaacs and Tang [1993]. All of these studies, then, report similar results: Users prefer video, but audio in general performs either as effectively or almost as effectively as face-to-face communication or video.

2.2 Mediated social spaces

Two sets of studies suggest that *sustained* use of an audio media system could lead to an interesting social space. The long-term "office share" studies, reported by Dourish et al. [1996] and Adler and Henderson [1994], are one example of an audio-video mediated space. These office shares were successful, and communicative practices arose that were adapted to the media and to the participants' work practices.

There is also evidence, albeit very little, that social spaces can be created with even low-quality audio. Orr's [1993] six-month field study showed how limited-bandwidth audio (worse than telephone quality) could create a social space. Service technicians were given two-way radios for continuous communication amongst group members over the regional service district. The technicians reported feeling less alone, and the radios were used for both work-related and social communication. Strub's [1996] field study, where teenagers utilized portable two-way radios, also showed that social spaces can be created with even low-quality audio technology.

2.3 Specific design considerations for an audio-only media space

Specific characteristics of the Thunderwire design, namely open connections, desktop mikes for ambient audio, fluidity of use, and good-quality audio, were informed by several additional studies. Dourish et al. [1996] found continuous open audio was important to long-term interaction patterns. Gaver [1993], in his exploration of auditory icons, pointed out the importance of ambient audio in the workplace, and how audio could ease the transitions between individual and shared work. Similarly, Whittaker, Frohlich, and Daly-Jones [1994] concluded from their study of informal workplace communication that persistent audio and video links would support frequent, brief interactions at minimal cost. They characterized workplace interactions as one long intermittent conversation.

Evidence from studies that examined the audio-only condition in media spaces with multiple media (e.g., [Gale 1990; O'Conaill, Whittaker, and Wilbur 1993; Olson, Olson, and Meader 1995]) argued for good-quality full-duplex audio without any transmission lag. For example, Gale [1990] found that having high-quality audio resulted in faster group task completion times than having audio and video.

Furthermore, there is evidence that low-quality audio could adversely affect communication. O'Conaill, Whittaker, and Wilbur [1993] looked at remote workgroups using videoconferencing over both ISDN (the low-quality condition) and broadcast (high-quality) networks, and compared them with the face-to-face condition. After examining conversational mechanisms such as backchannels, interruptions, overlaps (simultaneous speech), handovers, and turns, they conclude that low-quality, half-duplex audio with discernible lag causes more formality and conversational awkwardness. These results are similar to those found for telephone use (e.g., [Rutter 1989; Hopper 1992]).

2.3 Summary

The prior work on audio's communicative properties and on audio collaboration in media spaces influenced the overall and detailed design of the Thunderwire system, and led us to believe that such an audio-only media space would be usable. It also led us to choose a field study approach. This study differs from earlier work in that it examines an audio-only media space with good-quality audio, using

modern technology. This study, therefore, differs from the various office share studies, by exploring whether mediated office sharing would succeed in the absence of video. Furthermore, unlike many comparative mediated communication studies (including [Sellen 1995; Olson, Olson, and Meader 1995; O'Conaill, Whittaker, and Wilbur 1993]) and unlike many technology studies (e.g., [Smith and Hudson 1995]), we report here on the extended use of this audio-only media space by an existing workgroup performing a range of activities.

In summary, a number of studies have produced empirical findings that audio has sufficient communicative capability for an interesting and useful shared media system. Such a system is described in the following section.

3. THE THUNDERWIRE SYSTEM

Thunderwire was an audio-only communication system conceptually similar to a telephone party line or conference call. The system was built by Interval to facilitate communication within a small group spread throughout two buildings. Thunderwire permitted any number of group members to be simultaneously connected, and anything said at any time by any member was heard by all.

The following were the important system characteristics of Thunderwire. They are critical for understanding the field study results:

- Thunderwire was a purely audio medium. Except for an “on” light, it had no other visual interface or cues.
- The audio was high quality, such that users could easily distinguish one another’s voices as well as overhear background sounds. The sound quality made it possible to hear everything one might hear sitting in a person’s office, including private vocalizations, phone calls, bodily noises, and background noise.
- All messages were public on Thunderwire.
- System use was fluid. People could connect or disconnect themselves from Thunderwire any time they wished, simply by flipping a switch.

- The act of connecting or disconnecting was indicated only by a barely audible click. In fact, there was no way to know exactly who was listening without asking.

Users used desktop microphones, headphones, and controllers with three settings: Off, Listen-only, and On. There were on-off indicator lights for the microphones, and the sound volume could be adjusted. Ten Thunderwire stations could be linked together.

Finally, the Thunderwire system was robust. It was continuously available during the field study.

4. RESEARCH STUDY AND USE

As mentioned, we chose to study Thunderwire within a field setting, to better understand how potential users would view an audio space (if one were to exist). This field study of Thunderwire lasted slightly over two months, and because of the technical requirements, included only one group within Interval.

Before describing the findings from the field study, however, it is important to provide the necessary background. This section describes the field setting, basic usage patterns, and the users' self-reported evaluations.

4.1 The study group

The group using Thunderwire consisted of nine people, seven of whom were engaged in video editing and analysis. The group members themselves were generally young (often just out of college) and were not permanent staff members. In fact, two group members worked for subcontractors to Interval.

Each person concentrated on separate tape segments, so the work was independent yet closely parallel. Tight coordination was not required, although sharing analytical approaches and problems was helpful, as it might be in, for example, a newspaper office. One consequence of the work was that users were already listening to videotape over headphones, so using headphones for Thunderwire fit well into the existing work practice.

The group also included a manager as well as a software engineer. The manager of the group was older and more authoritative. The software engineer reported to another group, and was not well known to the group before the field study. He was actively supporting and improving an analysis tool for those group members editing video.

This team was largely cohesive before the system was introduced. Most of the group knew each other well; they had spent the summer collecting field data together. As the manager said, “They spent the summer brushing their teeth with one another.” All reported that over the summer, they had formed a cohesive social unit. Nonetheless, there were a number of personality tensions at the beginning of the study. As well, two of the nine group members were considered quite independent, and three others were outsiders in some way. One person had become an outsider during the summer, even self-identifying in that role repeatedly. There was also one other outsider who was a subcontractor and incidental to Thunderwire usage. The last outsider was, as mentioned, the software engineer who was not known to the group before Thunderwire use began.

There were several possible obstacles to adoption and use. Most of the Thunderwire participants sat in cubicles within 100 feet of one another; therefore, the Thunderwire system vied with face-to-face interaction for many members of the group. (The group manager, as well as the software engineer, were in another building.) Moreover, while it was in the group leader's interest to have closer communication with group members, it was not necessarily in the staff members’ interest, a discrepancy that Grudin [1989] noted often leads to adoption failures.

4.2 Methods and data

Data were gathered using multiple methods. Users' experiences with Thunderwire were studied over the two-month period through interviews, transcripts of use, usage logs, and direct observation, as follows:

- Semi-structured interviews were conducted with the Thunderwire participants before they obtained the system and at the end of the study period. The pre-introduction interviews

examined group cohesion, group tensions, and communication patterns. The final interviews with the participants were conducted after the field study. Additional interviews with key members were conducted approximately two weeks after the field study. The final interviews were audio-taped and transcribed; the other interviews were captured in detailed hand-written notes.

- Approximately two weeks of conversations, after the adoption phase, were tape-recorded (with the participants' permission), and approximately 18 hours of audio were selected for literal transcription. Conversations entirely and partially on Thunderwire were transcribed in detail; Appendix A shows a portion of one transcript. Short interactions, those taking less than a minute, were missed because of the sampling technique for the tapes. (However, one of the authors listened to several tapes to get a sense of these brief interactions.)

  Conversations were transcribed literally. Some of the tapes were transcribed in depth, marking length of pauses and overlaps in conversation.
- Additional qualitative data were obtained through direct observation, examination of source materials, and data reviews. The pre-introduction direct observation was extremely limited.
- Quantitative data were obtained through general usage logs and survey data.

The qualitative data analysis included a careful examination of the transcripts and field notes. The transcripts and field notes were coded for common topics and interaction patterns [Drew and Heritage 1992; Miles and Huberman 1994]. The quantitative and qualitative data were used to corroborate each other during analysis.

Only the non-Interval authors have had access to the audio data, as well as much of the interview data, to provide confidentiality to the study participants. All data discussed below have been made anonymous.

4.3 System use

In general, system acceptance by the socially central players was quickly accomplished. This was absolutely critical to adoption and use, as Kraut et al.[1994] and others have found. Central members of the group, where centrality implies social position rather than a work or task attribute, adopted the system earlier and used the system more. The exception was the software engineer, who was at first marginal to the group, then played a central role in the system, and later was accepted as part of the core group.

Two users formed the core of Thunderwire participation, using their stations actively and through much of the day. Three others often used the system (for a considerable portion of the day but not every day), and three others seldom used the system. The station in the group's collaborative space was used extensively but passively by the last group member and in varying amounts by the other participants. This pattern roughly followed the exponential curve found in many computer-mediated communication systems [Hiltz and Turoff 1981].

Continued use of the system was fragile. For example, during one observation period, one and then the other key participant was absent for a day and a half, and as a result, there was almost no use of the system.

4.4 Usage patterns

System use varied considerably. It was possible to be continuously connected to Thunderwire; it was not uncommon to find study participants connecting to Thunderwire first thing in the morning and disconnecting when they left at night.

"Live mike" time fluctuated somewhat from week to week, but averaged just over two hours per station per day, or approximately 25% of working hours. Since stations could be on without being used (e.g., left on in an empty office) and since listen-only mode could not be distinguished from the system being completely off, these are only approximate measures of system use. This average also masks the large amount of individual variance; the standard deviation was slightly larger than the average itself.

Figure 1 about here

Figure 1: Mean times per 8-hour day with at least 2, 3, and 4 live microphones, indicated by the upper,middle and lower lines, respectively. (No data were collected over the two-week holiday period.

A better measure may be the number of live microphones. Figure 1 shows the mean times per eight-hour day that Thunderwire had at least two, three, and four live-miked stations (the upper, middle, and lower lines, respectively). The system often supported two or three active participants, but use by larger groups was relatively rare.

Several uses of the system were innovative and surprising. One participant had Thunderwire jacked in as one of four audio inputs. Not surprisingly, this person was often a passive user of the system. Participants discussed whether they could share music through the system; they discovered that they could, but it would drown out conversation. Users also used the system to ask one another where a person was in transit through the buildings. Participants often inquired after their manager and one another.

4.5 Users' evaluations of Thunderwire

In the final interviews, participants were asked to describe Thunderwire. Most of them provided a functional description. Additionally, all but one added comments about the sociability of the Thunderwire space:

> Let's see, I'd probably describe it as ... like an intercom that you can turn on and off and sort of plug into a party line and where ... you can hear other conversations or participate in a group conversation or try to find a single person on the system to exchange information. -- John
>
> It's kind of an open line so whoever's on, you can hear -- Patrick
>
> Sometimes it's just like a plain old backyard gossip fest. -- Mike

Users not only noted the sociableness of the environment, they also noted the telepresence aspects of the system:

> I think another, in a way a social atmosphere is helped with Thunderwire to create kind of a "hi, how are you doing" ... [a] lightweight-sort of social space that people can join or not join if they so choose. -- John

> ...it's kind of replicating a situation where you're in close proximity with lots of other people. It's as if in certain ways if there were four or five other people with their desks all right around mine without these little partitions between them, except that we don't raise our voices to talk to each other, and it's kind of everyone's at their own discretion for how much they're participating. -- Rob

While these self-reports are interesting, it should be noted that the adoption process could have promulgated and sharpened this view with the users. Because the manager obtained the system to consciously create a virtual space for the work group, the users could have merely reflected the manager's arguments about community and telepresence. All that can be safely concluded is that the system did not counter this initial understanding on the part of the users.

Not all of the evaluations were positive. In fact, most were mixed. While the system afforded useful interactions—and these interactions, as will be seen below, were enjoyable as well as useful—the users had to struggle to define and regulate their space. These struggles will be detailed below.

The paper next turns to the possibilities and problems of Thunderwire. The following analysis shows that this audio space was both quite similar to, and quite different from, everyday social spaces as well as other types of media spaces.

5. INTERACTION CAPABILITIES

One important consideration for any media space is how well people are able to converse and interact over the system. If users are unable to conduct suitable conversational interactions or are unduly restricted in their range of interactions, then the system can hardly be called usable or useful. This section, therefore, analyzes the interaction capabilities of the Thunderwire system, as shown in the field study.

5.1 Rich interactions

We begin with an analysis of conversational interactions on the system. The following analysis draws on the work of Isaacs and Tang [1993], who describe the "rich interactions" possible on their system. In their comparison of video against telephone and face-to-face interactions, they defined rich interactions to be those requiring visual cues and geographical presence. In our study, we have extended their consideration of interaction types to include non-visual considerations such as conversational fluidity, topic flexibility, and conversational conventions.

Providing rich interaction is likely to be critical to workplace acceptance of mediated communication systems. As Isaacs and Tang state:

> ...we suspect that richer interactions are likely to lead in the long run to more and / or higher quality results. ([Isaacs and Tang 1993], p. 199)

Rich interaction may include lengthy, intimate conversations, but within a workplace setting, rich interactions may include very informal and unstructured exchanges. Whittaker, Frohlich, and Daly-Jones [1994] showed that informal workplace interactions are frequent, brief (usually two minutes or less), and do not include specific initiation and closing utterances.

An extended example (see Appendix A) demonstrates the sociability and richness of Thunderwire interactions. The interactions presented in this example are quite common in the data, although Appendix A's example does bring together a number of issues we wish to discuss.

It is very difficult to convey the tone of this conversation in print. The banter flows naturally, although with a staccato rhythm as the conversation lulls and restarts. Within a fragment, the conversation easily bounces back and forth among group members, one turn often following immediately after another, almost without pause. The following excerpt is typical in its informality and spontaneity. The two "Welcome back" greetings overlap considerably. (For a complete explanation of the transcription notation, see Appendix A.)

1	Mike	I'm back.

2	Rob	Mmm,
3	Patty	[with mock heartiness] Wel/come back\
4	Rob	[joining Patty] \welcome back/

There are several fragments to the conversation in the extended example, with little connection among the fragments. In turns 5 through 11, the three participants discuss a cartoon. From turn 12 through turn 30, the participants discuss names. From turn 32 through 36, they discuss keyboards. (Some turns have been omitted to save space, and some fragments are longer than they appear.)

Many conversations on Thunderwire show no formal beginning and ending points, unlike telephone conversations [Hopper 1992]. For example, after a one-minute pause, Rob reopens the conversation in turn 32 with a gambit, but this gambit is informal and conversational.

31		[1:04 min. pause. Typing, short amount of someone whispering to himself.]
32	Rob	Such a difference a real keyboard makes.

The interaction pattern is similar to being face-to-face in a common room.

Of particular interest is the overlapping speech pattern. In addition to the overlap seen in turns 3 and 4, above, overlap occurs prominently in turns 10 and 11, 14 and 15, and 20 and 21. Several other pairs include overlaps between speech and laughter. This pattern is quite similar to that of face-to-face interaction [O'Conaill, Whittaker, and Wilbur 1993] and dissimilar to telephone or low-bandwidth media [Hopper 1992; O'Conaill, Whittaker, and Wilbur 1993; Rutter 1987]

The playfulness in the exchanges is also notable. In turns 12 through 30, the participants unself-consciously make fun of their own names and those of others—perhaps by extension, their identities. In the conversation fragment from turn 39 to 45, excerpted below, Patty is waiting for a phone call to be completed.

39	Patty	[evidently on phone] Yes, is Sarah Altman there? ... It's her friend Patty Chapman. ... [Into Thunderwire] They always ask me [mock politeness, official tone] And what organization are you with?

40	Mike	I see ... and so you say you're a friend
41	Patty	So I'm, I'm trying to cut out that line of questioning.
42	Mike	I see. Or you should say, my dear, her deeply rooted enemy [Patty and Mike laugh.] [1 sec. pause] Just say [in crabbed, old voice] Ven-det-/ta\

Interestingly, she continues on Thunderwire commenting on her call, until the call goes through, at which point she disconnects from Thunderwire. (This disconnection norm for phone calls will be discussed below.)

43	Patty	[using the same crabbed voice] \It's a/ personal call [both laugh]
44	Mike	[3 second pause] tiz pretty funny. [4 sec. pause.] [repeating in old voice, half to self?] Ven-det-ta
45	Patty	[22 sec. pause, typing] Hello. ... Is this the woman who I assume is turning 29 today? [slight male laughter] Happy birthday. So are, you like, being showered with gifts and presents and ... food and stuff? ... [surprised] Twelve! ... Man, you did [Patty disconnects]

In this fragment, Mike and Patty joke back and forth in assumed voices. The use of mediated voices and wordplay was quite common with Thunderwire participants.

Again this exchange is friendly, playful, and close; the data clearly show spontaneous, personal, and highly social interactions. In this conversation and many others, there appeared to be a high level of interaction and informality.

5.2 Range of interactions

Another important consideration for media spaces is the range of interactions possible within the system. Based on Dourish et al.'s [1996] experiences with a continuous audio connection, one important type would be sociable conversation. As mentioned, this was also the case with Thunderwire. These sociable conversations were quite remarkable; they show considerable social interchange, play, and personal warmth.

There were also several other conversation types. A number provided Thunderwire members with a sense of participation and background information. These include mouse clicks, paper rustling,

background conversations, phones ringing, and people moving through the building. Several group members, in their final interviews, spoke at some length of enjoying knowing what was going on in their group, and users spoke of overhearing their manager's conversations with other group members. Of course, these exchanges and their social advantages trade off against individuals' needs for privacy; this was a constant tension for the group. These privacy concerns will be discussed below.

There were also many exchanges about coordinating activity and work-related events, as can be seen in the following example:

Rob	Judy's not in her cubicle is she?
John	Last I saw her she was in the lab space.
Rob	Oh. Okay.
Mike	And, she's not in her cubicle at the moment.
Rob	Okay. Thank you.
Vicki	[In the lab space] I think Judy went to lunch.
Rob	Oh, okay.

One of the authors was surprised one day when Thunderwire participants tracked him down by noting his movement through the building. Participants also used the system to announce the starts of staff meetings, provide important schedule changes for meetings, collect one another for lunch, and note when they completed assignments.

Rather surprisingly, there were only a handful of information exchanges about work tasks. Users seldom interacted over the system about the details of their work or asked "how do I?" questions over the system. The following is one of the few examples of such an exchange in the data. It is, as usual, part of a much longer conversation about a variety of topics. The most recent topic had been the use of mail headers, when Patty signals a problem:

Patty	[Macintosh chimes] Oh, my god, now my Eudora's like, gone crazy. How do you...stop this when it goes crazy? [Macintosh chimes continuing]
John	Hold on, I'm coming. [laughs slightly]

Patty	[in funny voice] A-a-a! ... Jesus. It's, like, opening every message I ever got in my life.
Mike	Oh, Jesus.
Patty	Can you hit, like, Control- [laughs]
Mike	You can... control... command-period might work.
Patty	[Continuing to laugh] Okay. [laughing hard] This happened to me a couple days ago, too.

Originally, it was believed that the system would further these types of exchanges, with their quick exchange of task-related information, and so the paucity of work-related information exchanges was very surprising. In follow-up interviews, Thunderwire participants' suggested reasons were the visual nature of the group's work (video-editing), individuals' work patterns (their editing and composing was largely individual work), and the group's work cycle (they were in a phase where most work was individual tasks). It is also possible that this is a result or partially a result of the system's affordances. This issue should be examined further in subsequent studies.

Nonetheless, the data overall, as demonstrated in the above examples, indicate that the study participants had both an extensive range of interactions and very rich interactions. Particularly notable is the degree to which people were able to fluidly socialize and interact through Thunderwire, unlike with low-quality audio and telephone systems

The ability of participants to conduct suitable interactions is an important aspect of audio spaces. If people could converse freely and interact in everyday ways on a system, they should be more willing to use it over an extended period. The following section further examines interactions within the social space and suggests how audio spaces might differ from other media spaces.

6. NORMS FOR A SHARED SPACE

The data, especially the audio transcripts, show a set of norms related to Thunderwire use. As one might expect, the group negotiated shared understandings [Strauss 1991] of what to do on the system, developing norms in reaction (a topic also explored in [Ackerman and Palen 1996]).

A general definition of norms is “group-supported definitions of expected behaviors in specific situations” [Meier 1990]. More importantly, as Feldman states:

> ...a group does not establish or enforce norms about every conceivable situation. Norms are formed and enforced only with respect to behaviors that have some significance for the group. ([Feldman 1984], p. 47)

The major norms for the Thunderwire users concerned the side effects or results of a shared social space: dealing with noise, knowing when someone was present, knowing when someone was listening, and limiting violations of personal privacy. While the Thunderwire participants were only partially successful (from an outsider’s viewpoint) in dealing with these issues, that they struggled with these issues reflects not only some problematic aspects of an audio-only space but also the social possibilities of being within a common space.

The remainder of this section discusses three sets of norms that best typify Thunderwire’s shared audio space. These norms were repeatedly observed within the data. The first set concerns announcing signing off and on to the system, the second set concerns inattention and withdrawal, and the third set concerns regulating private information.

6.1 Announcing oneself

One of the major problems for Thunderwire users was not being able to easily tell who else was within the space.

John	Is Judy here today?
Patty	Judy's here, ... and she was on ... earlier. She is either listening, or ... is not ... here with us.

Thunderwire shares this problem with other shared media systems that do not make public who is on the system. (Even some systems that do make users publicly known can have this problem, if the system is inconsistent or tardy in updating the user list, or if the system is not completely believed by the users.) However, not knowing who is on the system is particularly acute in an audio-only space.

There were three methods for knowing who was on the system: verbally signing onto Thunderwire, asking who was on the system, and verbally signing off the system. All were inherently imperfect, since a user could evade the norms. However, social sanctions (e.g., derisive comments) were applied by group members to those who were caught.

Signing on. Because Thunderwire participants were not visible to each other, it was common for group members to notify and greet each other when they entered Thunderwire. Additionally, participants often thoughtfully updated new arrivals when they signed on, letting them know who else might be on the system. In the following fragment, Rob signs on, and Mike quickly informs him who else might be on, and who Mike physically has in his office with him:

Rob	Hi
Mike	Hi, there Rob, how's it going?
Rob	Good
Mike	Patty was on a little while ago, I don't know if she's still there
	[dialog omitted]
Mike	And, also I'm talking to Bob Corbin in my office...

This norm is also seen in the Appendix A example in turn 1.

Interestingly, the audible click of a microphone being switched on or off served as a resource for group use. Although one might have perceived the click as annoying and assume it should be removed, participants used the click to know when people were joining or leaving the discussion. It was usual for someone to announce their presence. If a click was heard without an announcement (as the first author did once early on), someone asked who it was.

Signing off. Thunderwire participants usually let one another know when they were leaving the system as well, as can be seen in the following example, as well as in turns 37 and 38 of Appendix A:

Patty	Alright, I'm on the move again.
Mike	See ya <indistinguishable>.

This enabled people to know who was still on Thunderwire.

Telephone calls. One of the most obvious norms concerned telephone calls. Simply stated, if a user received or initiated a personal phone call, he was supposed to leave the Thunderwire space for the duration of the call. The sign-off norm did not extend in this situation; telephone calls were signaled on the system by the ringing sound or the participant's initial interaction with the call, as in the following conversational fragment:

Jake	\Yeah, sure,/
Patty	Yeah, I mean, I don't [switching to mock dramatic voice] I don't regret it. I have no regrets.
Jake	[laughs a little]
	[1:25 pause, typing starts after 7 seconds. Background conversation.]
Patty	[evidently into phone] Hi, Cath, it's Patty. ... How ya doin'? [click of someone going off line]

While this norm superficially appears obvious—listening to other people's phone calls is disruptive and invades privacy—the norm has some peculiarities. In shared, common rooms, one generally does not leave under similar circumstances. It is normally acceptable to attend to a personal call while in the presence of others (especially when there is not a shared conversation underway). Indeed, it might actually be considered rude, somewhat odd, or suspicious if one were to excuse oneself. There is a strong possibility, then, that this norm was formed in response to the particular challenges posed by Thunderwire.

Failures to disconnect when accepting a phone call were common. Near the end of the study, participants did not disconnect for very short calls when there was no conversation, perhaps as a result of both privacy or disruption concerns becoming more relaxed. However, failure to disconnect also formed the basis of Thunderwire's privacy problems, which will be discussed at length below.

A very interesting example of the telephone norm occurs in turns 39 through 45, shown above and in the appendix. Patty disconnects only after the call goes through, spending the intermediate time discussing the phone call over Thunderwire.

6.2 Inattention and withdrawal

In an audio-only environment, many of the visual cues that a person may normally use to avoid unwanted interaction in a socially acceptable way are not available. It is not possible to either establish or avoid eye contact, and it is not possible to adjust one's visibility. Therefore, there is no easy way to signal from a distance one's willingness to interact—one cannot close the door, avoid another's gaze, appear busy, frown, or appear intensely engaged. On the other hand, every utterance by any group member is heard perfectly and equally well by every other group member, and therefore one cannot plausibly pretend he did not hear another participant talk to him. Nonetheless, for the system to work, group members needed to find mechanisms through which they could signal when they were more or less desirous of social contact.

When users were preoccupied, it was important for them to be able to pay minimal attention and to be able to both signal this request and have it be accepted in a socially acceptable manner. Without this ability, participants would have to withdraw from Thunderwire space whenever they needed to concentrate or pay attention elsewhere. In general, Thunderwire participants signaled their inattention by pausing or uttering fill words in a distracted tone. For example, in the following exchange, John mutters "uh-hm" very slowly and without interest to signal his preferred inattention. This is the more remarkable because Patty addresses him directly:

Patty	John, are you there?
John	Um-hm.
Patty	Thanks for this New York thing.
John	Um-hm
	[Typing, mouse clicks, male clears throat]

John's brief responses signal that he is uninterested in further conversation, and the exchange ends. Nonetheless, John is still capable of paying some attention (i.e., maintaining his system presence without having to completely withdraw). Indeed, John later asks Patty about her conversation with the group's manager, a conversation heard on Thunderwire.

Another mechanism for signaling partial attention was to just let the conversation lapse. While this might seem odd compared to many face-to-face interactions, this type of conversation lapse does occur in many workplace conversations [Whittaker, Frohlich, and Daly-Jones 1994]. This occurs in the Appendix A example. A lull in the conversation is followed by turn 7, a new opening where Mike starts another extended conversation over the system. In other situations, the conversation slowly winds down as the participants do not revive it:

Patty	Yeah, I've seen that version. A couple of years ago I watched that.
Mike	It has everyone's heartthrob, Olivier, in it. ...And I think it might have been, is it Olivia de Haviland, could that have been Elizabeth? That's the bouncy, spritely sister who gets Darcy in the end.
Patty	[thinking] U-uh
Mike	I can't remember that piece of casting.
Patty	Hu-uh
	[12 minute pause]

In face-to-face behavior, one of the ways that people avoid blatant rejection is to use subtle, non-verbal cues to signal and to determine the possibilities of social contact [Koneya 1977]. Inattention and withdrawal were very public in the Thunderwire space, and the signals were more explicit than one might prefer. Therefore, group members appeared to have devised other means to gracefully terminate unwanted social interaction. Problems with the headphones were a common experience, and withdrawal blamed on apparent equipment failure was generally unquestioned.

6.3 Noise and unwanted disclosures

This is not to say that Thunderwire participants were able to successfully negotiate all issues in their space by formulating well-followed norms. On the contrary, unwanted disclosures and noise were two constant problems during the field study. Users did try to formulate norms to deal with these problems, but were not completely successful.

In a social environment, one must protect himself against others' unwanted information. In the simplest case, one has to be able to reduce unwanted noise by others. As one might imagine, loud noises can be a problem in an audio space, especially when the user is wearing headphones. While the users found low-level background noise to be acceptable, indeed useful and sociable, they did not like loud, disturbing noises. This problem partially resulted from placing the sound at an apparent location that was correct, or nearly so, for voices; the apparent location could be disconcerting for other sounds. This was especially true for nuisance sounds, such as computer-generated tones, the sounds of falling books, and telephone bells, especially if placed near the microphones. Eating, coughing, and other body noises were also problems.

For noise that was predictable or for which advance warning could be provided, it was the norm to announce that such a noise was about to occur. In the following conversation fragment, Mike provides warning that he is about to make some noise:

Mike	I'm going to make a little noise. Let me know if it drives you nuts... [breaths out]
Patty	[small tsk sound] Okey-dokey.

Unfortunately, one does not always know when he is going to sneeze. Nor does one know when his telephone will ring. These noises are impossible to predict and provide few clues that they will occur (especially in a non-visual environment), so participants were unable to formulate any norm or other social construction to deal with the issue. For example, in the following conversation fragment, Patty seems embarrassed that John overhears her sneeze:

Patty	[sneezes]
John	Gesundheit.
Patty	[funny voice] Scuse me...tsk.

Another type of unwanted disclosure was participants' accidentally leaking personal information into the environment. Although the users explicitly decided among themselves not to use the listen-only mode, it was still all too easy to leak personal information unintentionally within Thunderwire.

Forgetting to turn off the microphone while answering the phone or having a visitor could provide others with personal information:

Patty	Hi, Charley, this is Patty Chapman calling. [She clicks off.]
	[Rob and Mike laugh a little]
Rob	Bye, Charley [chuckles a little]
Mike	She remembered that time.
Rob	Yeah.

Compounding this problem, avoiding others' disclosures was difficult in Thunderwire, because of the nature of the space. Everything said was broadcast to everyone else, and one could not avoid hearing it. Audio lends itself to this type of shared broadcast space; one can integrate multiple audio sources intelligibly. While this provides a sense of shared space, it also makes managing disclosures more problematic.

Furthermore, an audio space provides fewer cues to know whether a disclosure was intentional or accidental. Personal disclosures do promote closeness, and participants often felt confused and guilty while overhearing others' deeply personal conversations, because they also enjoyed them:

> Stuff like that happened all the time.... Janette did it constantly, like ... she was having a conversation with Rachael Collins [professor at Stanford], broadcasting the whole thing, and we were all getting into [it]. It's like should we call her and let her know? No, we were all enjoying the conversation. -- Mike

At the end of the field study period, if something was disclosed into the public space, it was generally assumed that the participant either wished to do so or did not care, even if this was not true. This norm took some evolution. When the study began, the norm was to signal to the participant that he or she was disclosing. However, to do so required a substantial amount of effort by the entire group (particularly if the violation occurred when the user was not wearing headphones), and it was abandoned. The norm then became to ignore or listen to the accidental disclosure.

This was an imperfect solution, however. For most participants, especially those in the group's social core, leakage to fellow group members was hesitantly tolerated. One user reported:

> ...As a group, we.... we know each other well enough to know, okay [that] this person's probably got this going on, and this person's probably got that going on, and a lot of it is taken in stride. -- John

Yet, others wanted to substantially control and regulate their private information. The system did not allow this control, and in varying degrees, led three users to reduce their participation.

6.4 Norms and social behavior

In this field study, participants formulated and maintained the norms discussed above, as well as others. The particulars of these normative mechanisms may have been idiosyncratic to the individuals involved, but it must be stressed that the underlying requirements for some norms were not [Cohen and Bacdayan 1994]. Users, if they were to use the system, had to find adaptations to many audio-only and system conditions within Thunderwire. Some of the norms discussed above may have been more successful than others, but the participants did devise and maintain them.

It must also be emphasized that the construction of these norms would not have occurred unless the participants felt they inhabited a space together. The Thunderwire participants constructed these norms to deal with real problems in the Thunderwire space. And, they would not have been able to construct these norms or maintain them without their active participation within a common social space. This was clearly a *social* space.

This section demonstrated that users were able to create a social space with Thunderwire, and discussed some of the issues in creating and using the Thunderwire social space. The next section continues this discussion by examining how audio and system affordances affected one issue, privacy.

7.0 PRIVACY AND AUDIO

The previous two sections established that participants were able to use Thunderwire successfully for social (and sociable) interactions. It was "good enough" for everyday interactions. However, in Meyrowitz's [1985] terms, it was not a completely normal stage for the participants' conversational performances. In other words, Thunderwire distorted the normal mechanisms of social interaction.

Below we consider how audio and audio-only conditions change the nature of social interaction and give rise to a tradeoff between privacy and sociability in systems such as Thunderwire. Privacy violations (in accidental disclosures) were clearly a concern for participants, as manifested in the discussion above. Yet, participants also knew that privacy traded off with what they liked about the system -- its telepresence and sociability. In their final interviews, users spoke at length about being sociable and finding out about one another. One user reported overhearing another:

> I heard Rob one day on Thunderwire call KCSA.... And, okay, well, I didn't realize another KCSA lunatic was around here. And that was just a cool thing to find out, and suddenly there's this whole other area of stuff we can talk about. -- Mike

However for another user, that same type of accidental disclosure was an egregious error:

> I had a conversation with an old professor in Ohio; this was a while back. And it was something of a private conversation. ...I came back to my desk to check my messages, heard there was an urgent message, and I called him directly and forgot to turn it [Thunderwire] off. So, everybody heard my conversation. It turned out not to be that big a deal, but I would rather have not had them hear the conversation. -- Patrick

Clearly, there was a tradeoff occurring among sociability and privacy for the Thunderwire participants.

This section proceeds in two parts. The first briefly argues privacy is necessarily related to the ability to interact in spaces such as Thunderwire. We then consider how Thunderwire, as an audio-only system with specific features, might affect privacy. This second, more speculative section teases apart what was caused by the medium (i.e., audio) and what was caused by system design (i.e., Thunderwire in particular), allowing us to make design recommendations for future systems.

7.1 Control over private information

Shared media systems such as Thunderwire are about conversations and social interaction. Social interaction brings with it privacy concerns; the need to control private information is deeply related to the process of interaction.

This can be simply stated. It was very easy to engage others through the system -- one merely entered the Thunderwire space, a very simple act in the system. That it was so easy to do (and so likely to result in a pleasurable interaction) made the system enjoyable:

> Two people anchor Thunderwire as far as I can see. ...so there's always this presence. -- Judy

Following Goffman [1961] however, any social interaction requires the control of information, both outgoing and incoming. People wish to control their private information, presenting the proper "face" to others. Presenting information about side projects, outside work, or personal habit is potentially dangerous within an organizational setting, and such information is normally released with caution and in confidence. In interactions, then, people will want to know with whom they are interacting and then provide only the information they wish to disseminate. Accidental or unintended disclosures can be very disconcerting. It was no different in Thunderwire, as the above norms about knowing who was in the space and about accidental disclosure testify.

Normally, this management of one's information is handled as an "everyday" occurrence by individuals [Garfinkel 1967]. Thus, privacy is seldom seen as a problem.[1] In environments such as media spaces where one's management of his private information cannot be an "everyday" performance, privacy concerns may loom larger. People must then adjust to the changed implications for social interaction.

These privacy concerns are, of course, not specific to audio. They are shared by all computer-mediated communication systems. Yet, it is clear that the properties and affordances of audio interact with these privacy concerns. As Schmandt [1994] states, "Voice announcements, messages, or warnings are much more public than visual notification" (p. 103). He goes on to state, "But this characteristic of speech can also be advantageous in some classes of applications...." (p. 103). We examine the specifics of how audio and privacy are related next.

---

[1]We have adopted the definition of privacy used by the HCI literature (e.g., [Hudson and Smith 1996]). This includes only control over one's private information. Control over intrusions and the decision to engage are considered part of the standard dictionary definition of privacy, but are not included within the term "privacy" for this discussion.

7.2 Audio affordances, system affordances, and privacy

Since participants' efforts to socially organize themselves to avoid privacy violations were not overly successful, we would like to determine whether this issue could be corrected through future system features. To do so, we must determine whether this problem was the result of choosing audio as the media type, design choices about how to use audio, specific system features, or some interaction among all three. Our analysis must be necessarily speculative because of the limitations of a single case study.

However, some of the privacy issues here mirror those found in Dourish et al's [1996] and Adler and Henderson's [1994] office-share studies, and this provides the key to our analysis. In those studies, the continuous nature of the media space played a profound role. In a continuously broadcast space, it is very easy to forget that the system is on. As an innocuous example, in Adler and Henderson [1994], one user forgot that the camera was on and adjusted his clothing in public. Users often ignore a warning light or even video feedback over time; this was true for Thunderwire's on-off indicator light as well.

One would expect then, that the problem of accidental disclosure is not unique to audio, but that audio, without visual cues, exacerbates the problem under some conditions. In Thunderwire, if the user was wearing his headphones, he received audio feedback that he was in Thunderwire space. He was either aware or could be warned that an accidental disclosure was in progress. Unfortunately, a user was most likely to remove headphones while answering the phone or having a conversation with someone physically co-present, and indeed, these were the most troublesome situations in Thunderwire. In these situations, there was no audio feedback that he was in two spaces, one telepresent and the other physical. Furthermore, the telepresent audio space provided no visual cues, since this is a basic audio property. In general, since Thunderwire was audio-only, people revealing intimate details or having private conversations could not be warned through any secondary medium, channel, or interface. (In fact, people occasionally used the phone for this, but could not do so when the accidental disclosure occurred from telephone use.)

However, the roles of audio as a medium and audio-only as a system affordance are not completely clear-cut. Privacy violations in Thunderwire were mitigated by several design decisions. While sound can be heard by anyone nearby (another basic audio property), the headphones prevented the Thunderwire signal from escaping into the general environment, avoiding additional disclosures. Video is harder to control in this manner. Additionally, the Thunderwire signal, as in many shared media systems, was not broadcast indiscriminately; it was restricted to the Thunderwire group by the technical setup of the system. Thus, users had to be concerned about disclosures only to one another, a much less problematic issue than the threat of disclosing to varied people walking by or to varied people listening in.

The accidental disclosures that occurred, however, also exacerbated awareness issues in Thunderwire. As mentioned, people feel uncomfortable without knowing who else is in a shared space with them; they want to present the proper "face" to others. Since Thunderwire lacked a secondary affordance to provide awareness, users could know who was present only by listening to the audio. The system lacked additional cues from a video signal, a visual interface mechanism, or even an audio-only awareness mechanism. Social mechanisms served to ameliorate this lack in general. Yet, after accidental disclosures, users wanted to know who else was on the system without asking. In these situations, the norm of asking people was socially clumsy and was usually avoided. (It is psychologically difficult to ask who heard something that was embarrassing or confidential.)

With media spaces, then, a continuous broadcast environment affords the interactional occasion in which accidental disclosures occur. Audio-only, with its lack of secondary cues, makes such disclosures both more likely to occur and to continue once they start. Awareness is similar: Being in a continuous broadcast environment creates the occasion, but the lack of secondary cues worsens any problems. Audio, then, accentuates privacy issues, but this can be ameliorated through system affordances such as awareness indicators or by restricting the broadcast. While video adds some additional cues and feedback, it is not clear whether the benefits outweigh the costs.

It is important to remember that much of Thunderwire's utility came from both the continuous broadcast and the audio nature of the system. To provide another user's perspective on disclosure:

> But there was definitely, like, a rapid kind of going from not knowing Mike to (pause).... He and John were talking about Hong Kong, movies, and you know.... It was like that stuff sort of happened much faster. I never would have realized that he loved old movies, if we hadn't been talking on Thunderwire, I don't think.
> -- Patty

The broadcast nature of the system clearly contributed to this; this would not result from point-to-point communication systems such as the telephone. As well, the continuous nature of the connection allowed people to enter the system when they wished to engage others. Several features of audio in affording this type of engagement may not be as obvious. As Abel [1990], Heath and Luff [1993], and Dourish et al. [1996] note, it is extremely difficult to get all of the visuals correctly in a shared video system (e.g., the expected visual distance between speakers changes with the situation). Audio is simpler; the lack of visual cues aids in providing a usable space for communication. Furthermore, audio is less obtrusive. Users can use audio in a lightweight manner [Dourish et al. 1996]. Not only can users pay peripheral attention [Schmandt 1994], others cannot see that a participant is not paying full attention. While audio may exacerbate privacy issues, we speculate that it also contributes to ease of sociability in shared media spaces.

We would like, then, to keep the social ease of Thunderwire, while ameliorating this privacy issue. This is the task of future work. It must also be repeated that Thunderwire worked well enough to be used by the group; solving or reducing this privacy issue would make it *more* usable.

8. CONCLUSIONS

This study began by questioning whether audio alone was suitable for shared media systems and their resulting social spaces. The system that we examined, Thunderwire, included good-quality, full-duplex audio without lag; persistently available group communication capable of conveying ambient workplace sounds; an audio-only user interface; and a simple model of user control.

Our field study of Thunderwire use suggests (with the standard limitations of case studies):

- *Audio can be sufficient for a usable media space.* Thunderwire users conversed sociably and in what appeared to be a natural manner (Section 5). Many of the conversation characteristics, such as turntaking and overlapping speech, were in notable contrast to low-quality audio use, such as one finds with the telephone.

  On the other hand, some user interface mechanisms should be improved in future audio media spaces (Section 7). Users clearly would prefer some way of knowing who is present in the audio space. This could be addressed, for example, through the low disturbance audio explored by Smith and Hudson [1995], where users can hear who is speaking without hearing the words themselves. Additionally, some automatic mechanisms for turning off microphone input during an incoming call and for filtering loud noises would have been helpful. Allowing users to set up two-way, private conversations might have been useful as well, as has been suggested for video media spaces. None of these interface changes, however, would necessarily require a visual interface; all could be incorporated within an audio-only environment.

- *Audio spaces can lead to social spaces.* Thunderwire afforded a social space for its users. Evidence for this includes user evaluations (Section 4.5) and more importantly, existence of norms regulating social use of the space by group members (Section 6).

- *The nature of these social spaces is affected by audio's affordances.* Thunderwire users created and maintained norms in response to concerns that are either specific to an audio-only environment or are exacerbated by an audio-only environment (Section 6 and Section 7). However, the users created and maintained these norms with effort. Adapting to a shared public space was an especially acute issue for the Thunderwire participants (Sections 6.1 and 6.3). This was particularly true because of the lack of secondary system mechanisms for ameliorating some of audio's affordances that affected privacy (Section 7.2).

This study adds to the growing body of work on mediated communication theory (e.g., [O'Conaill, Whittaker, and Wilbur 1993]). It also demonstrates the utility of looking at social behavior to supplement the conversational structure measures often used by researchers of mediated communication. Thunderwire, as an audio-only environment, appears to have had certain *media* properties that had to be considered and potentially dealt with by group members. These media properties created certain *social* conditions (or rather allowed them to occur) for the study participants. While one could also imagine other responses (e.g., resistance to adoption or inter-personal conflict, [Markus 1983]), the interesting point is that group members had to change their behavior in response. While the specifics of the response may differ from system to system, and from group to group, the need for some response to the media and technology affordances will not. Not only can similar examinations of social use guide the design of media spaces, it can also bring us to a better understanding of audio and other media in general.

8. ACKNOWLEDGMENTS

This study would not have been possible without Andrew Singer, who had the original insight for the system and for this study. Nor would it have been possible without the support and cooperation of the study participants and the Thunderwire design team. This paper benefited greatly from comments by John Hughes, Jonathan Grudin, Ellen Isaacs, Elin Pedersen, Bonnie Johnson, Sean White, Wayne Lutters, Diane Schiano, Malcolm Slaney, and Hank Strub among others. We would also like to thank the two special editors, Chris Schmandt and Nicole Yankelovich, and several anonymous reviewers for their helpful comments.

9. REFERENCES

Abel, M. J. 1990. Experiences in an Exploratory Distributed Organization. In *Intellectual Teamwork*, J. Gallegher, R. E. Kraut and C. Egido, Ed. Lawrence Erlbaum, Hillsdale, NJ.

Ackerman, M. S., and Palen, L. 1996. The Zephyr Help Instance: Promoting Ongoing Activity in a CSCW System. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'96)*, 268-275.

Adler, A., and Henderson, A. 1994. A Room of Our Own: Experiences from a Direct Office Share. In *Proceedings of the ACM Human Factors in Computing Systems (CHI'94)*, 138-144.

Bly, S. A., Harrison, S. R., and Irwin, S. 1993. Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment. *Communications of the ACM, 36*, 1, 28-47.

Chapanis, A. 1975. Interactive Human Communication. *Scientific American, 232*, 36-42.

Cohen, M. D., and Bacdayan, P. 1994. Organizational Routines are Stored as Procedural Memory: Evidence from a Laboratory Study. *Organization Science, 5*, 4, 554-568.

Dourish, P., Adler, A., Bellotti, V., and Henderson, A. 1996. Your Place or Mine? Learning from Long-Term Use of Audio-Video Communication. *Computer Supported Cooperative Work, 5*, 1, 33-62.

Drew, P., and Heritage, J. 1992. *Talk at Work: Interaction in Institutional Settings.* Cambridge University Press, New York.

Feldman, D. C. 1984. The Development and Enforcement of Group Norms. *Academy of Management Review, 9*, 1, 47-53.

Gale, S. 1990. Human Aspects of Interactive Multimedia Communication. *Interacting with Computers, 2*, 2, 175-189.

Garfinkel, H. 1967. *Studies in Ethnomethodology*. Prentice-Hall, Englewood Cliffs, NJ.

Gaver, W. W. 1993. Synthesizing Auditory Icons. In *Proceedings of the ACM INTERCHI'93 Conference on Human Factors in Computing Systems*, 228-235.

Goffman, E. 1961. *The Presentation of Self in Everyday Life.* Anchor-Doubleday, New York.

Grudin, J. 1989. Why groupware applications fail: problems in design and evaluation. *Office: Technology and People, 4*, 3, 245-264.

Heath, C., and Luff, P. 1993. Disembodied Conduct: Communication Through Video in a Multi-Media Office Environment. In *Readings in Groupware and Computer-Supported Cooperative Work*, R. M. Baecker, Ed. Morgan-Kaufmann, San Mateo, CA.

Hiltz, S. R., and Turoff, M. 1981. The Evolution of User Behavior in a Computerized Conferencing System. *Communications of the ACM, 24*, 11, 739-762.

Hindus, D., and Schmandt, C. 1992. Ubiquitous Audio: Capturing Spontaneous Collaboration. In *Proceedings of ACM CSCW'92 Conference on Computer-Supported Cooperative Work*, 210-217.

Hopper, R. 1992. *Telephone Conversation*. Indiana University Press, Bloomington.

Hudson, S. E., and Smith, I. 1996. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'96)*, 248-257.

Isaacs, E. A., and Tang, J. C. 1993. What Video Can and Can't Do for Collaboration: A Case Study. In *Proceedings of the ACM Conference on Multimedia '93*, 199-206.

Koneya, M. 1977. Privacy Regulation in Small and Large Groups. *Group & Organization Studies, 2,* 3, 324-335.

Kraut, R. E., Cool, C., Rice, R. E., and Fish, R. S. 1994. Life and Death of New Technology: Task, Utility, and Social Influences on the Use of a Communication Medium. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94)*, 13-21.

Markus, M. L. 1983. Power, Politics, and MIS Implementation. *Communications of the ACM,* 26 (June), 430-444.

Meier, R. F. 1990. Norms and the Study of Deviance: A Proposed Research Strategy. In *Deviant Behavior: Readings in the Sociology of Norm Violations*, C. D. Bryant, Ed. Taylor & Francis, New York.

Meyrowitz, J. 1985. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. Oxford University Press, New York.

Miles, M. B., and Huberman, A. M. 1994. *Qualitative Data Analysis.* Sage, Thousand Oaks, CA.

O'Conaill, B., Whittaker, S., and Wilbur, S. 1993. Conversations Over Video Conferences: An Evaluation of the Spoken Aspects of Video-Mediated Communication. *Human-Computer Interaction, 8*, 389-428.

Olson, J. S., Olson, G. M., and Meader, D. K. 1995. What Mix of Video and Audio Is Useful for Small Groups Doing Remote Real-time Design Work? In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'95)*, 362-368.

Orr, J. E. 1993. Ethnography and Organizational Learning: In Pursuit of Learning at Work. In *Proceedings of the NATO Advanced Research Workshop on Organizational Learning and Technological Change,*. Springer-Verlag, Berlin, 47-60.

Rutter, D. R. 1987. *Communicating by Telephone.* Pergamon Press, New York.

Rutter, D. R. 1989. The Role of Cuelessness in Social Interaction: An Examination of Teaching by Telephone. In *Conversation: An Interdisciplinary Perspective*, D. Roger and P. Bull, Ed. Multilingual Matters, Philadelpha.

Schmandt, C. 1994. *Voice Communication with Computers: Conversational Systems.* Van Nostrand Reinhold, New York.

Schmandt, C., Ackerman, M. S., and Hindus, D. 1990. Augmenting a Window System with Speech Input. *IEEE Computer, August*, 50-58.

Sellen, A. J. 1995. Remote Conversations: The Effects of Mediating Talk with Technology. *Human-Computer Interaction, 10*, 4, 401-444.

Smith, I., and Hudson, S. E. 1995. Low Disturbance Audio for Awareness and Privacy in Media Space Applications. In *Proceedings of the ACM Conference on Multimedia*, 91-97.

Strauss, A. 1991. *Creating Sociological Awareness: Collective Images and Symbolic Representations.* Transaction, New Brunswick.

Strub, H. B. In Press. ConcertTalk: A Weekend with a Portable Recreational Audio Space. In *Proceedings of the Sixth IFIP Conference on Human-Computer Interaction (INTERACT '97).* Chapman & Hall, London.

Whittaker, S., Frohlich, D., and Daly-Jones, O. 1994. Informal Workplace Communication: What Is It Like and How Might We Support It? In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'93)*, 131-137.

Yankelovich, N., Levow, G.-A., and Marx, M. 1995. Designing SpeechActs: Issues in Speech User Interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'95)*, 369-376.

APPENDIX A: SAMPLE CONVERSATION

In the transcription, the angle brackets <> indicate a section of the tape that was indistinct. The square brackets [] indicate external sounds or an inference on the part of the transcriptionist, data analyst, or researcher. The slashes \ / and / \ indicate overlapping conversational areas on the tape; a = sign indicates that there was no appreciable pause between the two words. Three dots ... indicate a short pause; longer pauses were indicated by a description of the pause, as in [pause 1.5 minutes]. Portions of the transcript were omitted for lack of space. These are marked where they occur.

1	Mike	I'm back.
2	Rob	Mmm,
3	Patty	[with mock heartiness] Wel/come back\
4	Rob	[joining Patty] \welcome back/
		(segment of conversation omitted)
5	Mike	Oh, there was this hysterical cartoon somebody put on the library bulletin board [Patty laughs slightly] It's a Berkeley Breathed cartoon about Tammy Baker ...
6	Patty	[with interest] Yeah?
7	Mike	and it's [next word drawn out] really [laughs a little] wonderfully mean. So I recommend it if you
8	Rob	/<inaudible>\
9	Mike	\haven't/ seen it
10	Patty	/[giggles]\
11	Rob	\[laughing slightly] Alright./ [20 sec. silence. Sounds like can rattling, typing.]
12	Mike	Well, Susan Belman's middle name is Evangeline. ...

13	Patty	Yes=
14	Rob	=That is a /fact.\
15	Patty	\It's in/ all of her e-mails=
16	Mike	=That's pretty remarkable. ...There are so many special people here. [Patty and Rob laugh]
17	Rob	Maybe I should start calling myself Rob [emphasis] Steven [ends emphasis] McLaughlin
18	Patty	What was it?
19	Rob	Steven.
20	Patty	[confirming, considering] [drawn out] St/even\
21	Mike	\Steven/ just doesn't have it quite like, I mean I'm sorry but it [Rob laughs] doesn't have it like [dramatic pause] Evangeline
22	Patty	/[laughs]\
23	Rob	[breathlessly] \Evangeline/
24	Mike	Well, geez, I think I'm gunna change my name to Evangeline. [slight laughter] A-, in fact, that sounds pretty good, Evangeline Angora Thompson
25		/[Patty laughs]\
26	Rob	\[appreciatively, laughing] Wooo-hooo-hoo-hoo-hoo-hoo!/
27	Mike	Ah, jeez, a whole new me.
28	Patty	It definitely ... is ... evocative of so- ... a certain something, I'm not quite sure=
29	Mike	=Not quite sure, don't want to think of too much about what [laughing] exactly it's evocative /of\.
30	Patty	\[laughs]/
31		[1:04 min. pause. Typing, short amount of someone whispering to himself.]
32	Rob	Such a difference a real keyboard makes.
33	Mike	Yeah?
34	Rob	Yeah, I can type a lot faster now that I have
35	Patty	Yeah, I was making so many mistakes when I had ... [rising] just the little PowerBook.
36	Rob	[2 sec. pause] It's just a different feel to the keys.
		(segment of conversation omitted)
37	Rob	[clears throat, 2 sec. pause] I'm gunna go [tearing sound] try and find a couple tapes.

38	Mike	Take care, Rob. [1:55 min. pause. Patty whispers to self? Typing, mouse clicks.]
39	Patty	[evidently on phone] Yes, is Sarah Altman there? ... It's her friend Patty Chapman. ... [Into Thunderwire] They always ask me [mock politeness, official tone] And what organization are you with?
40	Mike	I see ... and so you say you're a friend
41	Patty	So I'm, I'm trying to cut out that line of questioning.
42	Mike	I see. Or you should say, my dear, her deeply rooted enemy [Patty and Mike laugh.] [1 sec. pause] Just say [in crabbed, old voice] Ven-det-/ta\
43	Patty	[using the same crabbed voice] \It's a/ personal call [both laugh]
44	Mike	[3 second pause] tiz pretty funny. [4 sec. pause.] [repeating in old voice, half to self?] Ven-det-ta
45	Patty	[22 sec. pause, typing] Hello. ... Is this the woman who I assume is turning 29 today? [slight male laughter] Happy birthday. So are, you like, being showered with gifts and presents and ... food and stuff? ... [surprised] Twelve! ... Man, you did [Patty disconnects]