

Exploring Non-Speech Auditory Feedback at an Interactive Multi-User Tabletop

^{1,2}Mark S. Hancock

¹Chia Shen

¹Clifton Forlines

¹Kathy Ryall

¹Mitsubishi Electric Research Labs
201 Broadway
Cambridge, MA, 02139, USA
{shen,forlines,ryall}@merl.com

²Department of Computer Science
University of Calgary
2500 University Dr. NW
Calgary, AB, Canada T2N 1N4
msh@cs.ucalgary.ca

Abstract

We present two experiments on the use of non-speech audio at an interactive multi-touch, multi-user tabletop display. We first investigate the use of two categories of reactive auditory feedback: affirmative sounds that confirm user actions and negative sounds that indicate errors. Our results show that affirmative auditory feedback may improve one’s awareness of group activity at the expense of one’s awareness of his or her own activity. Negative auditory feedback may also improve group awareness, but simultaneously increase the perception of errors for both the group and the individual. In our second experiment, we compare two methods of associating sounds to individuals in a co-located environment. Specifically, we compare localized sound, where each user has his or her own speaker, to coded sound, where users share one speaker, but the waveform of the sounds are varied so that a different sound is played for each user. Results of this experiment reinforce the presence of tension between group awareness and individual focus found in the first experiment. User feedback suggests that users are more easily able to identify who caused a sound when either localized or coded sound is used, but that they are also more able to focus on their individual work. Our experiments show that, in general, auditory feedback can be used in co-located collaborative applications to support either individual work or group awareness, but not both simultaneously, depending on how it is presented.

Key words: Computer-Supported Cooperative Work (CSCW), awareness, tabletop displays, non-speech audio.

1 Introduction

Interactive tabletops are well suited for co-located group collaborations. Technological advances, such as DiamondTouch [5] and DiamondSpin [22], now support system development and exploration for these interactive

tabletop environments; tabletop research is now an active area in the HCI community. Although people have begun to explore issues of size, orientation, territory and interaction for interactive tabletops [18, 19, 20], little work has been done on the importance and impact of multi-user feedback. The designer of a co-located environment must consider that multiple users perceive auditory feedback and that this simultaneous perception might affect performance, understanding, and user experience.

Interactive tabletops offer a different set of affordances and provide a different user experience from both single-user systems and groupware systems with personalized input devices and a shared wall display [15, 24]. In a multi-user interactive tabletop setting, the table serves as both a shared display and a shared input device. Because the display is visual, designers often focus on presenting information to users through the visual channel. Although visual feedback is the primary modality through which to communicate information in this environment, we believe auditory feedback may also serve an important role. However, it is as yet unclear what that role is and how it might enhance users’ experiences.

Multi-user interactive tabletop displays offer a unique environment in which to present information. Simultaneity in actions carried out by multiple people can be both efficient and interfering. In an initial user study of the UbiTable [7], one of the key findings was that users were confused when auditory feedback was provided to indicate operational errors. An identical system beep for both users was the source of a common reaction by most users: “Who was that sound for?”

In collaborative settings, users often work in parallel on individual tasks and may be unaware of the actions of their peers. For example, users at the table may want to be informed when other users are trying to access or manipulate objects on the screen that may be outside of their visual attention. Auditory feedback may be useful in these circumstances. While the use of redundant auditory feed-

back may, in some cases, increase group awareness, it may also hinder the performance of individual users. Alternatively, sounds that are useful for the individual may contribute to an overload in the auditory channels of the entire group. Thus, it is essential to consider this tradeoff when designing co-located, collaborative applications.

In this paper, we present two experiments on the use of non-speech audio at an interactive multi-touch, multi-user tabletop display. The goal of our studies is to understand the effects of auditory feedback on performance and awareness in a co-located group setting. Our study focuses on the dimensions of sound association, simultaneity, codification and localizations of sound played at the table, not the specifics of sound design.

2 Related Work

We separate the relevant literature into three categories: group awareness in collaborative applications, the use of auditory feedback in user interfaces, and previous use of auditory feedback in collaborative applications.

2.1 Awareness

Typically, group awareness – “the up-to-the-moment understanding of others’ activities in a shared space” [10] – has been discussed in the context of a distributed computing environment. Clearly, the remote physical distance of the individuals in the group will more drastically affect awareness than if they are co-located. However, the use of technology in a co-located environment can remove some of the available awareness information. For example, markers on a whiteboard produce distinctive sounds that inform those not looking at the board that it is in use. Electronic whiteboards only produce this sound artificially, if at all, and in doing so may lose some implicit information, such as force and speed of the strokes.

Gutwin and Greenberg [9] discuss the tradeoff between supporting the work of the individual and supporting the work of the group. The point of tension that most relates to co-located awareness is in their discussion of artifact manipulation. They claim that artifact manipulation is best for the individual when the system is designed to optimize speed of interaction, but that this reduces group awareness, due to the lack of feedback. For example, an animation to show that a window has been minimized slows down the individual, but may improve awareness for the group. They also describe the presence of consequential communication, “the characteristic movements of an action communicate its character and content to others”, and feedthrough, the feedback provided to the other individuals in the group when artifacts are manipulated, in co-located environments.

Our work expands distributed groupware literature to include co-located groups. A tabletop display setting

is unique in groupware, since users necessarily provide some amount of consequential communication. However, the use of a large-screen, horizontal display removes some of this communication, primarily in the auditory channel. For example, the movement of a digital document no longer produces the rustling sound normally caused by the movement of paper documents across the surface of a table. When working in groups, members often do individual work and may not notice the reduced feedthrough. Our study shows that providing auditory information can improve group awareness in this environment, at the cost of decreased individual performance.

2.2 Auditory Feedback

Auditory feedback has been used extensively to convey information in computer applications. Buxton [4] classifies non-speech audio messages into one of three categories: alarms and warning systems; status and monitoring indicators; and encoded messages and data. These messages can be conveyed using either abstract synthetic sounds called earcons or naturally occurring sounds that may be related to an action or event called auditory icons.

Brewster et al. [3] compared users’ ability to recall information about icons and menu items with the aid of earcons. Lemmens et al. [11] compared the use of earcons and auditory icons in a visual categorization task.

We perform a study that explores the use of auditory icons in a co-located environment. In particular, we are interested in how auditory icons affect simultaneous actions and group awareness in a multi-user environment, and potential solutions to problems of interference.

2.3 Auditory Feedback for Collaboration

Auditory information has been used to enhance awareness in some distributed collaborative applications. ShrEdt [6], a collaborative writing tool, handles “cursor collisions” (when one user attempts to edit the document at the same place as another) by playing a sound and displaying a pop-up dialog to the user making the change. The ARKola simulation [8] uses auditory icons to support awareness in a distributed collaborative bottling plant management application. An experiment involving eight pairs showed that users notice warning sounds, but not the absence of background sounds, while confirmation sounds provide an awareness of the partner’s actions.

Ambient auditory feedback has also been used to provide awareness. AROMA [17] presents an abstraction of captured data to display auditory cues about the remote presence of individuals in a distributed group. Audio Aura [16] allows for awareness of remote presence by providing “serendipitous information, via background audio cues, that is tied to people’s physical actions in the workplace”.

We wish to support the kinds of auditory feedback suggested by these systems. Our work presents an initial step toward integrating this type of auditory feedback in a co-located environment where the cognitive overload may be increased due to the number of users generating and perceiving the sounds heard.

Auditory feedback has also been used in other co-located collaborative applications. The Pond system [23] provides auditory feedback for actions such as queries and file retrievals in the form of an aquatic soundscape at a tabletop touch-screen display. Although they report that users were able to use the sounds for functional feedback, the sounds were not intended to improve performance or awareness, but instead were used simply to enhance the experience of using the system. The DynaWall [14], a large-screen tiled display wall that supports direct input, uses auditory information to provide awareness of others who may be working in the users' periphery due to their proximity to the screen. Sounds are presented when users perform gestures or move objects. Sound is presented through three speakers, one behind each of the screens.

Jam-O-Drum [2] is an interface that allows users to collaboratively create music around a circular tabletop. Each user is able to create their own music and mix it with the sounds made by the group. This system was tested with three different sound setups: a global mix, where all users' sounds were mixed to two stereo speakers, heard by all; distributed sound sources, where each user wore headphones and heard a mix of the sounds made by other users', spatialized to match their physical location; and individual speakers mounted in front of each user. They found that the use of a global mix made it difficult for users to identify their own sounds. Distributed sound sources showed some improvement of ability to identify sounds, but only for musicians and at the expense of inhibiting communication between users. The individual speaker setup proved to provide both identifiable sounds while allowing communication between users.

Sotto Voce [1] provides auditory feedback in a museum exploration application using a single headphone in one ear. This setup addresses the issue of lack of communication found in Jam-O-Drum by using two channels of input, one for sound input from the computer and the other for communication with a group. Ringel et al. [12] compare a similar setup to the use of public speakers at a tabletop display. They found that the use of earbuds in one ear allowed for more equal participation, greater individual involvement, and more group discussion.

Our interest is both in supporting awareness and reducing interference caused by simultaneous actions resulting in sound. We present a concrete evaluation of the effect of different types of feedback, as well as a comparison of

two different methods of associating a sound with a particular user. Our work differs from these previous studies in that it is not focused on an individual application, but rather is intended to inform the design of any co-located collaborative application that uses auditory feedback.

3 Parameters of Multi-user, Co-located Feedback

Sellen et al. [21] present a parameterization of feedback. We have adapted that list to focus specifically on providing feedback to multiple users in a co-located, collaborative environment:

Modality – the combination of auditory, visual, or other sensory mechanism through which the feedback is presented.

Number of users – the number of users that will perceive the provided feedback.

- *Source(s)* – which user (or users) initiated an action.
- *Target(s)* – who will be able to perceive the provided feedback.
- *Association* – some other subgroup (e.g. who is affected).

Type of feedback – the types of feedback available to the designer fall into three categories:

- *Proactive* – feedback provided to the user before the user performs an action.
- *Reactive* – feedback provided to the user after (or as) the user performs an action. This type of feedback can be further classified:
 - *Affirmative* – affirmative reactions confirm the actions of the user (e.g. highlight when selected)
 - *Negative* – negative reactions indicate errors (e.g. beep when moving past end of document)
- *Ambient* – information provided in the periphery of a user's attention.

4 Exploring Auditory Feedback

In our first experiment, we investigate the use of two categories of reactive auditory feedback: affirmative sounds that confirm user actions and negative sounds that indicate errors. In our second experiment, we compare two methods of associating sounds to individuals in a co-located environment. We compare localized sound, where each user has his or her own speaker, to coded sound, where users share one speaker, but the waveform of the sounds are varied so that a different sound is played for each user.

4.1 Individual Focus vs. Group Awareness

In single-user applications, auditory feedback can be useful for identifying errors and to verify success of a particular action. In this setting, it is useful to provide sounds that help to improve performance and to eliminate redundant sounds that may annoy users or overload their auditory channel. In collaborative settings, users often work in parallel and may be unaware of the actions of their peers. Sounds that improve individual performance in this setting may hinder group awareness, while sounds that improve awareness may hinder individual performance.

4.2 Task Selection

On tabletop displays it is possible to display a vast number of objects simultaneously. These objects can originate from many different users. Many different access control policies can be used to mediate this sharing of material, many of which involve restricting specific actions from certain users [13]. Thus, it is important to provide some feedback mechanism to the user to indicate which objects he or she is allowed to access and from whom an object originated.

In this paper, we use auditory feedback for the purpose of classifying on-screen objects according to their pre-assigned access permissions. The task (explained in more detail below for each experiment) was designed to incorporate the basic subtasks of searching, manipulating and organizing documents – subtasks that are likely to be performed in a co-located, collaborative environment. In addition, the classification of on-screen objects lends itself easily to a divide-and-conquer strategy that allows for the simultaneous exploration of individual performance and group awareness.

5 Experiment: Affirmative vs. Negative Feedback

In this experiment, we tested the use of reactive auditory feedback in a multi-user environment. We compared the use of negative feedback which indicates an error for the individual to affirmative feedback which indicates the success of an operation.

5.1 Participants

Eight groups of three students from local colleges and universities participated in the first study. There were four male-only, one female-only and three mixed groups for a total of 17 males and 7 females with an average age of 22.8 years ($SD = 3.6$).

5.2 Apparatus

A 76 cm x 60 cm DiamondTouch [5] table, a multi-user-simultaneous-touch surface, was used for input. The displayed image was projected from above at a resolution of 1024 by 768 pixels. Sound was presented through com-




Figure 1: Three participants sat with two at the narrow ends and one at the wide end of the DiamondTouch table. The speaker was placed opposite the middle user.

puter speakers placed at one of the wide ends of the table opposite one participant and adjacent to the other two (see Figure 1). The system was powered using a 3.0 GHz Pentium IV processor.

5.3 Method

In each trial, three participants sat around the table as shown in Figure 1. The objects on the table were assigned an access policy so that some subset of the three users (perhaps all three) was “allowed” to move each object (note that the assigned permissions were only superficial and did not prevent users from actually moving objects). Participants were instructed to move each object into one of the color-coded bins in front of them. The participants were given a one minute time limit in which they were asked to classify as many objects as they could, making as few errors as possible. An error was explained to be either moving a picture that you are not allowed to move, or placing a picture in the incorrect bin.

In all trials, the access permissions were displayed at all times by the presence or absence of colored tabs on the sides of the object facing the allowed users (see Figure 2).




Figure 2: (left) Color-coded tabs indicate the access permission of each picture. (right) Participants placed each picture into the bin corresponding to its assigned access permission, indicated by color.

The presence or absence of two classes of auditory feedback determined the four conditions in the experiment. Affirmative auditory feedback was provided under the following two circumstances:

- When an object was moved by a user who had permission, the sound of “rustling paper” was played.
- When an object was placed in the correct bin, a “pop” sound was played.

Negative auditory feedback was provided under the following two circumstances:

- When an object was moved by a user who did not have permission, a “scraping” sound was played.
- When an object was placed in the incorrect bin, a “buzz” sound was played.

These two classes were crossed to determine the four conditions of auditory feedback: none, affirmative-only, negative-only, and both. The order of the four trials was counter-balanced using a Latin Square design. Before the trials began, each group performed four practice trials with auditory feedback presented in the same counter-balanced order as in the experiment.

After each trial, participants were asked to write down how many pictures they had classified and how many pictures were classified by all three users. They were asked to report on both correct classifications and incorrect classifications for a total of four numbers after each trial.

After all trials were completed, participants were asked to complete a background and user-feedback questionnaire. The questionnaire was followed by an informal interview.

5.4 Results

A 3 (user location) x 2 (affirmative audio) x 2 (negative audio) within-groups factorial analysis of variance (ANOVA) was used to analyze the dependent measures described in Table 1.

	Speed	Error
Actual	Total pictures correctly classified	Total pictures incorrectly classified
Perceived	Reported number of correctly classified pictures	Reported number of incorrectly classified pictures
Awareness	perceived speed – actual speed	perceived error – actual error

Table 1: The dependent measures used for analysis in the first experiment. For the perceived and awareness measures, scores for both the individual (self ratings) and for the entire group (group ratings) were separately analyzed.

One participant did not report a score for self error in any of the four trials and another user did not report a score for group speed for two of the four trials. These data points were removed from the analysis.

Due to the number of dependent variables analyzed, all non-significant results are excluded to improve readability. However, all significant results we found are reported and discussed.

Actual Speed

A significant interaction between user location and affirmative auditory feedback was found for speed ($F(2, 12) = 8.48, p = .005$). Post-hoc analysis showed that the middle user was significantly slower than both the left user ($p = .040$) and the right user ($p = .038$) when affirmative auditory feedback was present. When affirmative auditory feedback was absent, the right user was significantly slower than the left user ($p = .045$).

Perception

Significant main effects of negative auditory feedback were found for perceived error, both for the individual ($F(1, 6) = 24.94, p = .002$) and for the group ($F(1, 6) = 22.85, p = .003$). Participants reported higher error scores for *themselves* with negative auditory feedback ($M = 1.40, SD = 0.27$) than without ($M = 0.62, SD = 0.15$). They also reported higher *group* error scores with negative auditory feedback ($M = 4.95, SD = 0.42$) than without ($M = 2.24, SD = 0.80$).

Awareness

A significant main effect of affirmative auditory feedback was found for awareness of speed for the group ($F(1, 6) = 14.82, p = .008$). Participants more accurately reported the group’s speed when affirmative auditory feedback was present ($M = 17.93, SD = 4.79$) than when it was absent ($M = 24.74, SD = 4.54$).

A significant main effect of affirmative auditory feedback was found for awareness of self error ($F(1, 6) = 9.75, p = .021$). Participants more accurately reported self error scores when affirmative auditory feedback was absent ($M = 1.19, SD = 0.17$) than when it was present ($M = 2.12, SD = 0.28$).

A marginal main effect of negative auditory feedback was found for awareness of errors made by the group ($F(1, 6) = 5.61, p = .056$). Participants more accurately reported group error when negative auditory feedback was present ($M = 2.57, SD = 0.55$) than when it was absent ($M = 4.71, SD = 0.72$).

User Feedback

Participants were asked which of the four sounds in the experiment they would choose to have in an application. Almost all participants reported that they would choose

the “pop” sound (91.7%), most participants reported that they would choose the “buzz” sound (70.8%), and few participants reported that they would choose to have the “scraping” (29.2%) or “paper” (20.8%) sounds.

5.5 Discussion

Our results verify that a tradeoff between support of individual work and group awareness exists, even in a co-located environment. Affirmative sounds were shown to improve users’ ability to accurately determine how many correct classifications the group made, while inhibiting their ability to determine how many incorrect classifications they themselves made. The redundant information provided in the auditory channel allowed users to process information about what the group was doing, at the cost of being unable to process information about their own actions. The speed of the middle user was also reduced when affirmative feedback was present. Because the middle user is most affected by the actions of the group, with partners to both the left and the right, an awareness of their actions may be most distracting to this user.

Negative sounds were not shown to inhibit group awareness, nor were they shown to improve speed or accuracy of the individual. However, negative sounds were shown to change the perception of incorrectly classified objects. Users perceived more errors when negative auditory feedback was provided. Users also reported that both affirmative sounds were more enjoyable than the negative sounds, which were reported to be more annoying. In contrast, negative sounds did improve upon group awareness of incorrectly classified objects. This improvement is likely because negative sound is not only informative, but can be redundant in this context, and therefore useful to the group.

A performance difference was also found in detriment to the right user when affirmative auditory feedback was absent. This effect is likely due to the increased activity of the middle user with the absence of affirmative audio. Since most of the middle users were right-handed, increased activity interferes with the performance of the right user.

A strong user preference for the “pop” and “buzz” sounds over the “paper” and “scraping” sounds is perhaps due to interference caused by the sound of the projector, causing the latter two sounds to be difficult to distinguish. This preference may also be due to the fact that participants were not asked to report on the number of correct and incorrect access attempts, only on completed classifications.

A common complaint heard in the informal interview following the experiment was that users mistakenly thought that they had made an error when it was in fact their partner. They reported difficulty in identifying who

caused a sound to occur for all sounds, and even more difficulty identifying who caused both the “paper” and “scraping” sounds. In contrast, they were generally able to determine what the cause of the “pop” and “buzz” sounds were, but remained neutral about the cause of the “paper” and “scraping” sounds. This observation led to the following study, which explores different means of improving a user’s ability to identify the initiator of a particular sound.

6 Experiment: Localized vs. Coded Sound

In a co-located environment, providing auditory feedback introduces the issue of sound association (i.e. what sound belongs to which user). In this experiment, we explored two methods of supporting sound association with sounds available to all users simultaneously.

In the first experiment, we discovered that auditory feedback may help to improve group awareness, but may hinder individual work or perception. Furthermore, users reported difficulty identifying to whom the sounds belonged. Associating particular sounds with particular users may both improve awareness and support individual work. In designing the setup for sound output in this environment, we wish to support the association of a sound with a particular user, without limiting the perception of that sound by the other users at the table.

6.1 Association of Sound with a Particular User

There are (at least) two methods of associating particular sounds with individual users in a co-located environment:

Localized the location of the sound (i.e. speaker position) can be associated with a particular user or subset of users.

Coded sounds can be codified so that the waveform of a sound is used to identify with whom the sound is associated (e.g. varied pitch or tempo).

For localized presentation of sound, the sound can be played either publicly (e.g. through unshielded speakers) or privately (e.g. through headphones or shielded speakers). Although private presentation of sound may be useful for some applications, it may not allow for improved group awareness, so we do not consider that condition.

Note that these two methods are not mutually exclusive. That is, coded sound can be localized to individual users. Note also that they are not independent. Some codifications of sound may affect the ability of the user to localize the sound. For instance, high frequency sounds are easier to localize than low frequency sounds. Loudness may also affect the localizability of sound, since loudness cues are sometimes used for this purpose.




Figure 3: Participants were seated as in the first experiment. Localized sound was presented through speakers mounted in front of each user. Sound that was not localized was presented through a single speaker placed opposite the middle user.

An alternative decomposition of audio feedback might be to utilize the parameterization of sound itself (pitch, rhythm, tempo, dynamics, timbre, location [4]) where location is one of these parameters. We have chosen this separate decomposition to reflect the physical setup that may be required to produce these effects.

6.2 Participants

Eight different groups of three participated in the second study. There were three male-only, no female-only and five mixed groups for a total of 17 males and 7 females with an average age of 21.5 years ($SD = 3.6$). One participant reported colorblindness for shades of brown and red and a second reported green/red colorblindness. One participant had a hearing impairment of 20% in his right ear. These three participants were in different groups. 15 participants reported having musical training of some kind.

6.3 Apparatus

The DiamondTouch input and top-projected display was identical to the setup from the first experiment. Similarly, non-localized sound was presented in the same way as in the first experiment. Localized sound was presented through three additional speakers placed one in front of each participant (see Figure 3). For coded sound, we used timbre to distinguish correct and incorrect placement and pitch to distinguish between different users.

6.4 Method

The task in this experiment was identical to the task in the first, with a few minor exceptions. The “buzz” sound and the “pop” sound from the first experiment remained and the other two sounds were removed. This choice was made because of the user preference results from the first study, which suggested that the paper and scraping sounds were difficult to distinguish.

Table 2 describes the factors used in this experiment.

	Present	Absent
Localized	Three small speakers were placed in front of each user.	A single larger speaker was placed centrally.
Coded	Sound was coded using a different pitch for each user (right=low, middle=medium, left=high).	The “buzz” and “pop” sounds were played at the same pitch for all users.

Table 2: A summary of the levels of each factor presented in the second experiment. The experimental conditions were the four combinations of these two factors.

The four conditions were counter-balanced using a Latin Square and four practice trials were performed, as before.

After each trial, participants were asked to write down how many pictures each user had classified, both correctly and incorrectly, for a total of six numbers. This differs from the first experiment in that participants were asked to report numbers for all three users and not just to differentiate between themselves and the group.

The completion of the four trials was again followed by a questionnaire and an informal interview, but the questions asked were modified to correspond to the sound setup and not the individual sounds used.

6.5 Results

A 3 (user location) \times 2 (localized on/off) \times 2 (coded on/off) ANOVA was used to analyze the data. The same dependent measures were used as in the first experiment, except that instead of splitting the perception and awareness data into individual scores and group scores, the data were separated by user location (left, middle, and right). We similarly remove all non-significant results for the purpose of clarity.

Perception

A significant main effect of user location was found for the perceived number of errors made by the right user ($F(2, 14) = 5.25, p = .020$). Post-hoc analysis revealed that the right user ($M = 0.81, SD = 0.15$) perceived less right-user (self) errors ($p = .003$) than the left user ($M = 1.94, SD = 0.27$).

A marginal main effect of localized sound was found for the perceived number of errors made by the middle user ($F(1, 7) = 4.83, p = .064$). All users perceived more middle-user errors when localized sound was on ($M = 1.96, SD = 0.36$) than when it was off ($M = 1.38, SD = 0.36$).

A significant interaction was found between user location and coded sound for the perceived number of er-

rors made by the middle user ($F(2, 14) = 7.30, p = .007$). Post-hoc analysis revealed that, when coded sound was off, the left user perceived more middle-user errors than either the right user ($p = .007$) or the middle user ($p = .020$).

Awareness

A significant main effect of user location was found for awareness of errors made by the right user ($F(1, 7) = 7.59, p = .065$). Post-hoc analysis revealed that the right user ($M = 0.44, SD = 0.12$) more accurately reported right-user (self) errors ($p = .003$) than the left user ($M = 1.38, SD = 0.18$).

A marginal main effect of localized sound was found for awareness of errors made by the middle user ($F(2, 14) = 4.83, p = .064$). All users more accurately reported middle-user errors when localized sound was on ($M = 0.62, SD = 0.17$) than when it was off ($M = 1.19, SD = 0.24$).

User Feedback

Participants rated nine 7-point Likert scale questions for each of the four conditions in the experiment. A Kendall's W test was used to determine consistent differences between responses for the different conditions. Participants consistently rated the sounds in the coded condition as more annoying than the other three conditions ($\chi^2(3, N = 24) = 8.6, p = .035$). Participants also consistently stated that they were less able to tell which partner caused a sound in the non-coded, non-localized condition than in the other three conditions ($\chi^2(3, N = 24) = 8.7, p = .033$). Table 3 shows a summary of a subset of the questions and their average rating over all four conditions.

6.6 Discussion

All three users were able to perceive that the middle user had committed fewer errors when localized sound was on. This heightened perception also led to better aware-

I found the sounds helpful in completing my individual work	5.2
The sounds distracted me from my assigned task	2.8
The sounds helped me to understand what my partners were doing	4.4
I was able to determine when I caused the sounds to be played	5.3
When one of my partners caused a sound to happen, I was able to tell which partner it was	3.6

Table 3: Overall ratings for a subset of the questions asked in the follow-up questionnaire in the second experiment (1 = strongly disagree, 7 = strongly agree).

ness scores for errors caused by this user. This effect may be due to the physical setup of the localized sound system. Both the left and right users were at a disadvantage because their speaker was collinear with the speaker of the user sitting directly across from them. Only the middle user sat in a position in which all three speakers were in different directions (see Figure 3). This unique location is likely the cause of heightened awareness in errors made by the middle user. This result suggests that the use of localized sound requires careful planning about the placement of the speakers. With the addition of a greater number of users, the utility of such a setup may become reduced.

The main effects of user location for both perception and awareness suggest that the right user correctly identified that they caused fewer errors than either the left or middle user believed they had caused. Because the auditory information was available about which user caused the error in three of the four conditions, the right user may have been able to process this information. The other two users, however, were likely too focused on their own tasks to be able to process the same information. Although this finding is consistent with the previous finding that the support of individual work may hinder the awareness of the group, it does not demonstrate a difference between the setups used in this experiment. Similar patterns were observed in the awareness scores of the other two users, but were not statistically significant.

No effects of actual, perceived or awareness of correct classifications were found in this experiment, and the significant results that were found had small effect sizes (usually a difference of no more than 2 errors). We believe this may be partially due to the complexity added in the second experiment with the requirement that participants report scores for all three users instead of just themselves and the group as a whole. In practice, the simultaneous evaluation of awareness and performance presents a significant challenge in user testing. We have demonstrated here that this small change in task requirement can largely impact the potential for significant results.

The study proved to be qualitatively rich. Participants in the second study reported being able to more easily identify sounds caused by their partners with both coded and localized sound. They also claimed to be able to identify sounds made by themselves more easily than sounds made by others. This suggests that the use of either coded or localized sound may provide more support for individual work and simultaneously reduce awareness of others. Because the user can easily identify their own sounds, they are not required to process the awareness information that is provided, and therefore can more easily ignore the available sound cues. This explanation also accounts




Figure 4: A graph of speed vs. awareness for the eight groups in the second experiment shows that some groups sacrificed speed for group awareness. Groups outlined in a circle had low (i.e. more accurate) awareness scores at the cost of low speed. The group outlined in a square likely opted for greater speed at the cost of awareness.

for the higher annoyance rating for coded sound.

Speed vs. Accuracy vs. Awareness

During the experiment, one of the groups attempted an interesting strategy. The participants synchronized their classifications so they were able to more accurately determine the number of correct classifications made by all users (because they were all the same number). This strategy demonstrates at a low level the tradeoff between group awareness and individual performance. This group was willing to completely sacrifice individual performance in order to maintain the maximum amount of awareness. Upon further investigation, the only group that had a similarly good awareness score for correct classifications had also sacrificed for speed (see Figure 4). An opposite, less easily identifiable strategy was to minimize awareness and ignore the actions of the group, and therefore maximize either speed or accuracy.

7 Conclusion & Future Work

Table 4 describes the implications of the combined set of studies to the use of auditory feedback in a co-located collaborative environment.

Our study shows that the tradeoff between group awareness and individual work is not unique to distributed groupware, and must be considered in the design of co-located, collaborative applications. In particular, we have shown that although the presence of auditory feedback may improve group awareness, it can also hinder the work of the individual. Our study also provides evidence that different group strategies may favor either performance or awareness at the cost of the other. Users can decide to maximize awareness of the group and, in so doing, simultaneously sacrifice either speed or accuracy.

The use of auditory feedback in a co-located collaborative environment introduces issues of interference that

	Benefit	Tradeoff
Affirmative Sound	group awareness heightened	individual awareness reduced
Negative Sound	group awareness heightened	individual and group perception of error increased
Localized or Coded Sound	awareness of individual errors heightened + improved focus on own work	awareness of group errors reduced + loss of attention to group

Table 4: A summary of the combined findings.

do not exist in either single-user or distributed groupware applications. We have presented two possible methods of associating sounds with individual users in a co-located environment. User feedback suggests that these methods enable users to identify which sounds they themselves produced and therefore allows them to focus on their own task. However, the cost of using an association method may be a reduced awareness of group activity.

We do not directly present a use for a “universal” speaker in the localized sound setup. The designer may wish to provide auditory feedback that is not associated with an individual at the table. A universal speaker may be useful for providing information relevant to all users, but is not necessarily associated with any one user (e.g. system out of memory error). Future work will include the design and evaluation of this style of feedback.

We wish to further explore the tradeoff between speed, accuracy and awareness in co-located collaborative environments. The relationship between these three dependent measures may be key to a better understanding of the evaluation of collaborative technologies. In the more immediate future, we intend to enhance existing tabletop applications with auditory feedback and to use the findings of our study to inform the integration process.

There is a definite need for additional research on appropriate feedback for co-located, shared-display environments such as interactive multi-user tabletops; our findings lay a solid foundation for future work.

Acknowledgements

We would like to thank Merrie Ringel for feedback on experimental design and Paris Smaragdis and Bill Buxton for help on sound design.

References

- [1] Aoki, P. M., Grinter, R. E., Hurst, A., Szymanski, M. H., Thornton, J. D., and Woodruff, A. Sotto voce: exploring the interplay of conversation and mobile audio spaces. In *Proc. CHI 2002*. ACM Press, 431–438.
- [2] Blaine, T. and Perkis, T. The Jam-O-Drum interactive music system: a study in interaction design. In *Proc. Designing Interactive Systems 2000*. ACM Press, 165–173.
- [3] Brewster, S. A., Wright, P. C., and Edwards, A. D. N. An evaluation of earcons for use in auditory human-computer interfaces. In *Proc. CHI 1993*. ACM Press, 222–227.
- [4] Buxton, W. Speech, language and audition. In R. Baecker, J. Grudin, W. Buxton, and S. Greenberg, eds., *Human Computer Interaction: Toward the Year 2000*. Morgan Kaufmann Publishers, 1995. 525–537.
- [5] Deitz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. In *Proc. UIST 2001*. ACM Press, 219–226.
- [6] Dourish, P. and Bellotti, V. Awareness and coordination in shared workspaces. In *Proc. CSCW 1992*. ACM Press, 107–114.
- [7] Everitt, K. M., Forlines, C., Ryall, K., and Shen, C. Observations of a shared tabletop user study. In *Proc. CSCW 2005: Interactive Poster*. ACM Press.
- [8] Gaver, W. W., Smith, R. B., and O’Shea, T. Effective sounds in complex systems: the ARKola simulation. In *Proc. CHI 1991*. ACM Press, 85–90.
- [9] Gutwin, C. and Greenberg, S. Design for individuals, design for groups: Tradeoffs between power and workspace awareness. In *Proc. CSCW 1998*. ACM Press, 207–216.
- [10] Hill, J. and Gutwin, C. Awareness support in a groupware widget toolkit. In *Proc. Supporting group work 2003*. ACM Press, 258–267.
- [11] Lemmens, P. M., Bussemakers, M. P., and de Haan, A. Effects of auditory icons and earcons on visual categorization: the bigger picture. In *Proc. ICAD 2001*. ICAD, 117–125.
- [12] Morris, M. R., Morris, D., and Winograd, T. Individual audio channels with single display groupware: Effects on communication and task strategy. In *Proc. CSCW 2004*. ACM Press, 242–251.
- [13] Morris, M. R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. Beyond “social protocols”: Multi-user coordination policies for co-located groupware. In *Proc. CSCW 2004*. ACM Press, 262–265.
- [14] Müller-Tomfelde, C. and Steiner, S. Audio-enhanced collaboration at an interactive electronic whiteboard. In *Proc. ICAD 2001*. ICAD, 267–271.
- [15] Myers, B. A., Stiel, H., and Gargiulo, R. Collaboration using multiple pdas connected to a pc. In *Proc. CSCW 1998*. ACM Press, 285–294.
- [16] Mynatt, E. D., Back, M., Want, R., Baer, M., and Ellis, J. B. Designing audio aura. In *Proc. CHI 1998*. ACM Press/Addison-Wesley Publishing Co., 566–573.
- [17] Pedersen, E. R. and Sokoler, T. AROMA: Abstract representation of presence supporting mutual awareness. In *Proc. CHI 1997*. ACM Press, 51–58.
- [18] Rogers, Y., Hazlewood, W., Blevis, E., and Lim, Y.-K. Finger talk: collaborative decision-making using talk and fingertip interaction around a tabletop display. In *Proc. CHI 2004: Extended abstracts*. ACM Press, 1271–1274.
- [19] Ryall, K., Forlines, C., Shen, C., and Morris, M. R. Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. In *Proc. CSCW 2004*. ACM Press, 284–293.
- [20] Scott, S. D., Carpendale, M. S. T., and Inkpen, K. M. Territoriality in collaborative tabletop workspaces. In *Proc. CSCW 2004*. ACM Press, 294–303.
- [21] Sellen, A., Kurtenbach, G., and Buxton, W. The prevention of mode errors through sensory feedback. *Human Computer Interaction* 7 (1992), 2 141–164.
- [22] Shen, C., Vernier, F. D., Forlines, C., and Ringel, M. DiamondSpin: an extensible toolkit for around-the-table interaction. In *Proc. CHI 2004*. ACM Press, 167–174.
- [23] Ståhl, O., Wallberg, A., Söderberg, J., Humble, J., Fahlén, L. E., Bullock, A., and Lundberg, J. Information exploration using the pond. In *Proc. Collaborative virtual environments 2002*. ACM Press, 72–79.
- [24] Streitz, N. A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. i-LAND: an interactive landscape for creativity and innovation. In *Proc. CHI 1999*. ACM Press, 120–127.