Investigating Effects of Visual and Tactile Feedback on Spatial Coordination in Collaborative Handheld Systems

Koji Yatani[1,3], Darren Gergle[2], and Khai N. Truong[1]

[1] University of Toronto
Ontario, Canada

[2] Northwestern University
Illinois, USA

[3] Microsoft Research Asia
Beijing, China

{koji, khai}@dgp.toronto.edu, dgergle@northwestern.edu

ABSTRACT

Mobile and handheld devices have become platforms to support remote collaboration. But, their small form-factor may impact the effectiveness of the visual feedback channel often used to help users maintain an awareness of their partner's activities during synchronous collaborative tasks. We investigated how visual and tactile feedback affects collaboration on mobile devices, with emphasis on spatial coordination in a shared workspace. From two user studies, our results highlight different benefits of each feedback channel in collaborative handheld systems. Visual feedback can provide precise spatial information for collaborators, but degrades collaboration when the feedback is occluded, and sometimes can distract the user's attention. Spatial tactile feedback can reduce the overload of information in visual space and gently guides the user's attention to an area of interest. Our results also show that visual and tactile feedback can complement each other, and systems using both feedback channels can support better spatial coordination than systems using only one form of feedback.

Author Keywords

Visual feedback, tactile feedback, collaboration, spatial coordination, mobile/handheld devices, touch screen.

ACM Classification Keywords

H5.3. Group and Organization Interfaces: Computer-supported cooperative work; H5.2. User Interfaces: Haptic I/O.

General Terms

Human Factors

INTRODUCTION

Mobile and handheld devices are becoming commonplace as tools to support remote collaboration. People now use their mobile devices to coordinate and interact with others through email, instant messaging, and video conferencing. They use their mobile devices to communicate and collaborate in similar ways to those with desktop computers.

However, the physical form-factor of mobile devices can impact user collaboration. For instance, mobile devices offer a much smaller visual workspace in comparison to those available on desktop computers or tabletop displays. Furthermore, when the user interacts with a touch-screen interface, her hand can occlude much of the screen. As a result, the visual feedback often used in synchronous collaborative tasks to help users maintain awareness of their partner's activities [8] may impact collaboration differently on mobile devices.

We thus see many design opportunities for improving collaboration through mobile devices. This includes using different communication channels for providing feedback to collaborators such as vibrotactile feedback: Tactile feedback using vibration motors spatially arranged in a device can be used to convey spatial information [19, 22] (we refer to this as *spatial tactile feedback*). This can be used to improve the expressiveness of tactile feedback on a given device. Furthermore, tactile feedback can be used to design more accessible user interfaces, particularly for visually impaired users who may not be able to access the visual feedback channel entirely. Understanding the effects of tactile feedback on collaboration can contribute to better design of collaborative systems for visually impaired users as well as sighted users.

We investigate how different feedback channels can affect collaboration on mobile devices. We are specifically interested in spatial coordination in a visual workspace [8], an important aspect of collaboration. Such collaboration can happen when users are sharing screen views and discussing the content within them. To measure the effectiveness of feedback for supporting spatial coordination, we built a collaborative game played by two users on mobile touch-screen devices, and we observed how users played the game with and without receiving visual or tactile feedback about their remote game partner's action.

Through two user studies, we gain an understanding of the different benefits of each feedback channel and also demonstrate better performance through the combination of visual and spatial tactile feedback for spatial coordination in collaborative handheld systems. Visual feedback can provide precise spatial information about a collaborator's action, but degrades collaboration when feedback is occluded and sometimes can distract the user's attention. Spatial tactile feedback may not be appropriate for

CSCW'12, February 11–15, 2012, Seattle, Washington, USA.

conveying precise spatial information; however, it can reduce the overload of visual information in the workspace and can gently guide a user's attention to an area of interest. Furthermore, visual and spatial tactile feedback can complement each other, and systems using both feedback channels can offer better spatial coordination support than systems using only one form of feedback.

RELATED WORK

Vibrotactile feedback has been widely used to provide tactile sensations on touch-screen devices. Active Click was perhaps the first demonstration of vibrotactile feedback on touch-screen devices [7]. Poupyrev *et al.* later brought this idea to mobile devices, and extended it by using different vibration patterns to increase the expressiveness of tactile feedback [17]. Yatani and Truong employed multiple vibration motors on the backside of mobile touch-screen devices to convey semantic information about the object which the user is touching [22]. These vibrotactile feedback systems are useful for a variety of tasks, such as item selection in a linear list [17], text entry [10], and eyes-free interaction [22].

These systems primarily provide the user with feedback about her own interactions on a mobile device. But, in this work, we focus on the use of tactile feedback for offering information about another user's interaction on a remote device to support collaboration between users.

Collaborative Systems with Tactile/Haptic Feedback

The use of tactile and haptic feedback has been explored in inter-personal communication systems, such as HandJive [6] and InTouch [1]. HandJive consists of two joystick-like devices enhanced by haptic feedback [6]. One person's movement on the device is propagated onto another device as an orthogonal movement (*e.g.*, when a user moves a HandJive device forward, the partner will feel horizontal movement from the device). InTouch [1] is a device with three cylindrical rollers, and the rotational velocity of each roller is synchronized with the paired device. Thus, remote users can feel each other's interaction with the rollers over a distance. Oakley *et al.* implemented several haptic effects for supporting communication in tasks with a collaborative editor using the PHANTOM device [14]. For example, one user can produce haptic feedback (called *haptic gestures*) on the other user's device to guide her to a specific point of the screen. They found that the participants frequently used haptic gestures to communicate with each other about the region of interest within the shared visual workspace or objects they wanted to discuss.

Tactile feedback is accessible to visually impaired users, and thus collaborative systems with tactile feedback often extend to this population. Plimmer *et al.* [18] developed McSig, a system to support visually impaired users in learning how to write letters through collaboration with a sighted user. Through a PHANTOM device held by a visually impaired user, the system haptically reproduces the trajectory of a letter written by a sighted user on a Tablet PC. Their user study showed that visually impaired users could successfully learn how to write some alphabets. Pielot *et al.* [16] built a belt-like device with multiple vibration motors and studied the effects of using it to provide tactile feedback of where fellow teammates are located in a 3D multiplayer game. Their user study revealed that users could sense the other players' locations, but did not show clear evidence on whether collaboration between the players was improved through their tactile system.

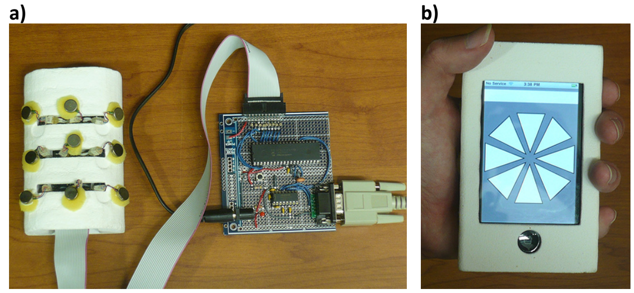

Figure 1. The mobile touch-screen device used in this study: a) the backside of the mobile device and the hardware to control the nine vibration motors embedded in the mobile device; b) the front side of the mobile device.

Communication Strategies Supported by Tactile Feedback

Research has focused on the use of a haptic or tactile channel for communication and collaboration as well. Chang *et al.* [4] investigated the effect of vibrotactile feedback on remote voice communication. In their ComTouch system, they explored the use of vibrotactile feedback as a supplemental communication channel for voice communication. Their user study revealed that the participants used tactile communication for five purposes: emphasizing a voice message, turn-taking, duplicating a tactile message, responding "yes" or "no," and conveying integer numbers. Brown *et al.* [2] explored how couples would develop communication protocols with their mobile audio-tactile messaging system called Shake2Talk. They identified four purposes for audio-tactile messages: coordinating events and calls, maintaining awareness, sharing the fun, and expressing affection. Chan *et al.* [3] examined the effects of haptic icons on turn-taking in collaborative tasks; specifically participants were instructed to solve puzzles while communicating through tactile feedback. They found that vibrotactile feedback can be useful to communicate some messages, particularly urgent requests in a visually-demanding situation.

Collaborative systems with tactile feedback can enable visually impaired users and sighted users to collaborate. McGookin and Brewster [13] explored the use of a collaborative tactile system to allow visually impaired users to explore a bar graph collaboratively. Through a user study with pairs of visually impaired users, they found that haptic feedback was useful for one user to guide the other user to the point of interest effectively. However, collaborative systems for visually impaired users are still in the early stage, and understanding the effects of different types of feedback including tactile feedback could be useful for

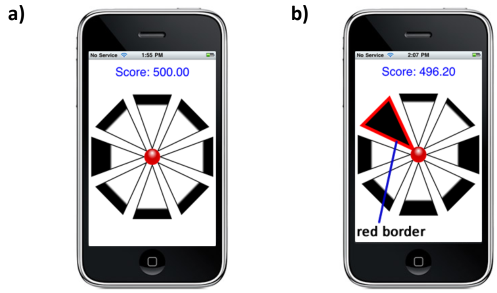

Figure 2. Screenshots of the game: a) The screen contains eight wedges, and they fill with the black color from the outside edge toward the center; b) When a wedge becomes completely filled, the system shows the red highlight and the players start to lose points.

designing collaborative systems for visually impaired users as well as sighted users.

Our work focuses on how tactile feedback can support spatial coordination in collaborative tasks. Spatial coordination is an important aspect of interaction between users to successfully accomplish collaborative tasks [5], but it has not been studied deeply in the context of tactile feedback. This paper contributes to the design of tactile feedback for collaborative systems.

SYSTEM

Hardware

We developed a spatial tactile feedback system for mobile touch-screen devices similar to the SemFeel device [22] (an iPhone in our system; Figure 1). We developed a special sleeve to embed the vibration motors in the backside of a mobile touch-screen device. The vibration motors are aligned in the 3 × 3 grid with 2 cm of separation. We determined this placement to satisfy a psychological limitation that it is hard to distinguish two distinct vibration sources with a gap smaller than 1 cm [15, 20]. It also offers eight different vibration points needed in the game design explained in the next section.

Game Design

The goal of the system is to measure the effectiveness of feedback on spatial coordination between remote users. We decided to build a game involving an abstracted spatial coordination task which can happen on mobile devices (*e.g.*, tapping the same region of the screen). This abstraction allows us to examine the effects of different feedback channels, while minimizing the effects of outside factors which could affect spatial coordination between users. At the same time, it maintains the generalizability of the study results to systems needing tightly-coupled spatial coordination in the workspace (*e.g.*, a collaborative drawing application or a system supporting visually impaired users in developing their motor skills to perform gestures or handwriting [18]).

Figure 2 shows screenshots of the game used in our study. The game screen is shared on the two mobile touch-screen devices when users play the game. Each screen shows eight interactive wedges and the score. The pair initially starts the game with 500 points. When the game starts, the system gradually fills a wedge (with black pixels) from the outside edge towards the center. The game fills each wedge at a different rate. When a wedge becomes completely filled, the system highlights that wedge with a red border, and the pair starts to lose points. One point is deducted for each completely filled wedge per second, and players do not have any way to gain points.

Success in this game depends on effective spatial coordination between the players. The objective of this game from the player's perspective is to keep as high a score as possible, given five minutes of play. To prevent a wedge from becoming completely filled, both must touch the same wedge, and at least one of them must perform a scrubbing gesture on that wedge. The game calculates the amount of black pixels to remove based on the length of the user's scrubbing gesture. The system ignores all scrubbing gestures when both game partners are not touching the same wedge. In this manner, the game requires frequent, tightly-coupled coordination between the players.

The game always provides the user with feedback about whether the touch screen has registered her touch by highlighting the selected wedge with a green border. Additionally, the game uses visual or tactile feedback to inform the user which wedge their game partner touches.

Visual Feedback

Figure 3a and 3b show how our system provides visual feedback. When one player touches a wedge, that wedge is highlighted by a blue border on the game partner's screen (Figure 3a). When both players touch the same wedge, its border color turns orange on both devices (Figure 3b).

Tactile Feedback

Figure 3c and 3d show how our system provides tactile feedback. When one player touches a wedge, the game partner's device activates the vibration motor associated with that wedge to generate a localized discontinuous vibration by turning it on/off every 200 msec (Figure 3c). When both players touch the same wedge, both devices generate continuous vibration with the motor associated with the contacted wedge (Figure 3d). The motor positioned at the center was not used in this study.

System Architecture

The mobile devices connect to a server machine through wireless communication, and they report all contact events to this server. The server computes how fast each wedge gets filled based on the predefined game pattern. When the players are scrubbing the same wedge, it also computes how much of the fill needs to be removed. This information is sent to each mobile device, which renders the game screen. When visual or tactile feedback is enabled, the server also sends each device the information about which wedge the game partner is touching. The mobile device then provides visual or tactile feedback as we described above.

STUDY1: EFFECTS OF EACH FEEDBACK TYPE

We conducted a laboratory study to examine how visual and tactile feedback individually affects spatial coordination on mobile devices.

Conditions

We designed three feedback conditions: *No feedback* (the system did not provide any visual or tactile feedback about the game partner); *Visual feedback* (Figure 3a and 3b); and *Tactile feedback* (Figure 3c and 3d). For each type of feedback, we also controlled the availability of the audio channel which allowed the participants to talk with each other: *Audio* (the audio channel was provided and the participants could talk with each other); and *NoAudio* (the audio channel was disabled). Table 1 shows the five conditions we studied. We excluded the condition of *No feedback* without the audio channel because it does not allow the participants to communicate in any way and they would not be able to collaborate strategically. Our main interest was to understand the effects of feedback on audio-enabled conditions as we were motivated by collaborative scenarios we can see frequently; however, this inclusion of audio-disabled conditions allows us to understand how well visual or tactile feedback could convey spatial information by comparing against the *Audio-only* condition.

The study was a within-subject design where each pair received exposure to all experimental levels, and presentation order of the system feedback was counter-balanced across the participant pairs. The order of the audio channel availability was fixed to *Audio* followed by *NoAudio* within *Visual* and *Tactile feedback* conditions. All of the conditions used a predefined script which specified changes in the rate at which the game would fill each wedge. Although the same script was used, the game rotated the script randomly for each condition. In this manner, we controlled the difficulty of the game to be the same across all conditions.

Procedure

Twenty-four participants were recruited in teams of two persons. Upon arrival to the laboratory, they were given an explanation about the system and the game. After this explanation, we separated the participants from each pair into different rooms. We gave each participant a device to use during the study as well as a microphone and speaker so that they could communicate with each other during the experiment. We asked each participant to wear a head-band with a mounted Web camera (Figure 4). We adjusted this Web camera so that it could record all interactions that each participant performed on the mobile devices. Participants then had a practice session to become comfortable with the system before starting the experiment.

		Feedback		
		No	Visual	Tactile
Audio Channel	Yes	*Audio-only*	*Visual-Audio*	*Tactile-Audio*
	No	(N/A)	*Visual-NoAudio*	*Tactile-NoAudio*

Table 1. The five conditions studied in the first user study.

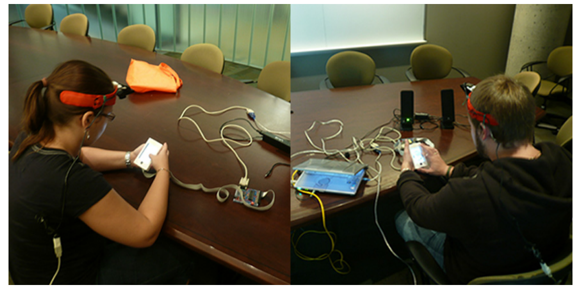

Figure 4. Our laboratory study. Two participants were brought to different rooms. In the audio-enabled conditions, participants were allowed to talk with each other through microphones.

All touch events generated by each user were recorded. The state of all wedges and the point score were logged every 100 msec. The system also audio-recorded all conversations and stored videos recorded by the Web cameras attached to their head-bands for analysis. At the end of the experiment, we conducted a short semi-structured interview to explore the difficulties with collaborating in each condition and their reasons.

Participants

We recruited 12 pairs of participants (PA1–PA12) between the ages of 18 to 39 with a variety of backgrounds (such as students, teachers, engineers, and business persons) for this study. Three of the pairs were both male, one of them was both female, and the rest consisted of one male and one female. All the pairs knew each other before participating in this study. The study lasted approximately 70 minutes. Each participant was compensated with $30 after the study.

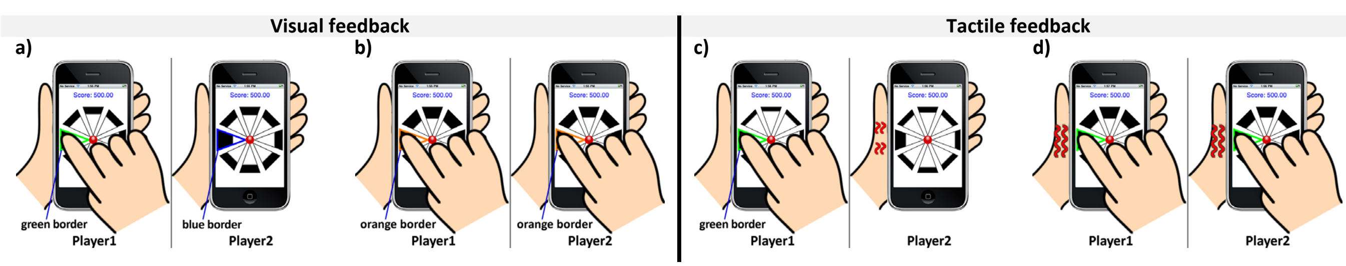

Figure 3. Screenshots of the game with visual and tactile feedback: a) When one of the players touches a wedge, it is highlighted with green, and the system also provides the blue highlight on the contacted wedge (in this example, the left wedge) in the game partner's screen; b) When both of the players touch the same wedge, the highlight turns orange; c) When one of the players touches a wedge, it is highlighted with green, and the system provides discontinuous vibration from the vibration motor associated with the contacted wedge; d) When both of the players touch the same wedge, the vibration becomes continuous.

Theme	Examples
Targeting	“Seven”, “Go to top”
Confirmation	(in response to *Targeting* or *Planning*) “OK”, “Yeah”
Clarification	(in response to *Targeting* or *Planning*) “What?”, “One?”
Strategy switching	(Changing play strategies) “Go counter-clockwise”, “Go this direction”
Awareness	“I'm doing it”, “Where are you?”, “keep going”
Planning	“Seven, nine, and one,”
Prompting (the next move)	(after swiping) “OK, go ahead”

Table 2. Coding scheme.

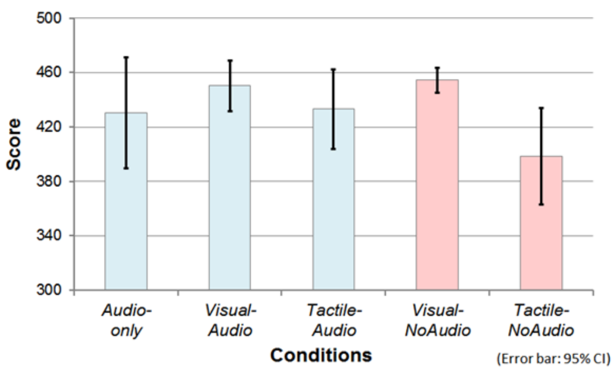

Figure 5. The average scores for the conditions tested in the first user study

Utterance Analysis

All the conversations that the participants had during the experiment were transcribed with timestamps as faithfully as possible. The authors conducted open coding of the quotes to identify seven themes pertaining to coordination, and developed the coding scheme as shown in Table 2. One of the authors and another coder independently categorized the recorded utterances along the scheme; they achieved high inter-rater reliability for every theme (higher than 95% agreement and 0.8 Cohen’s kappa).

STUDY1 RESULTS

Score

All pairs played the game for the full five minutes in each condition. Therefore, we were able to remove playing time from our analysis and focus on the performance scores. Figure 5 shows the average score for each condition. Mauchly’s test did not reveal a violation of sphericity and therefore permits the direct interpretation of the ANOVA F-test results. A one-way repeated-measure ANOVA revealed a significant difference in the scores by condition ($F_{(4,44)}$=.88, p<.05). Tukey’s HSD revealed that scores for the two visual conditions (*Visual-Audio* and *Visual-NoAudio*) were significantly higher than the *Tactile-NoAudio* condition (p<.05). The other conditions were not found to be significantly different.

In addition to the performance scores, it is telling to examine the communication processes that take place in the different feedback conditions, and understand their influence on performance. By doing so, we gain insight into the type of coordinating information that each feedback mechanism provides. We classified the types of spoken content using the coding scheme described in Table 2 and the results are shown in Table 3 which reports the average number of utterances by content type for the three audio-enabled conditions.

Theme	*Audio-only*	*Visual-Audio*	*Tactile-Audio*
Targeting	**78.4** (16.0)	**37.8** (37.1)	**55.4** (36.0)
Confirmation	**12.2** (12.1)	**11.3** (12.5)	**14.0**(15.2)
Clarification	**1.4** (1.6)	**0.7** (0.8)	**0.7** (0.9)
Strategy switching	**2.1** (1.8)	**1.6** (1.8)	**1.3** (2.3)
Awareness	**3.3** (4.9)	**2.6** (2.3)	**2.5** (2.7)
Planning	**3.4** (5.6)	**1.8** (4.6)	**2.2** (6.3)
Prompting	**0.8** (0.9)	**2.4** (5.1)	**0.7** (1.4)
Total	**101.5** (24.0)	**58.3** (54.6)	**77.1** (54.0)

Table 3. The average number and standard deviation of utterances for the three audio-available conditions.

To examine how these contribution types affected the game score in each condition, we used a random effects linear regression model where *Condition*, *Content Theme*[1], and the *Condition* × *Content Theme* interactions were included as independent variables. Because the pairs participated in all conditions, observations were not independent and were therefore modeled as a random effect. The resulting model fit was moderate (R^2=.64, *Adj*-R^2=.40).

This model controls for the types of utterances that were generated in addition to the feedback provided by the condition, and describes their influence on game score. Controlling for language, we see a significant effect of *Condition* ($F_{(2,19.19)}$=3.72, p<.05) where performance increased from *Audio-only* to *Tactile-Audio* to *Visual-Audio*. Of the content types examined, there is a main effect of *Targeting* content ($F_{(1,21)}$=8.37, p<.01), controlling for condition. Higher level interactions revealed significant interaction effects of *Condition* and the number of *Targeting* utterances ($F_{(2,19.96)}$=4.63, p<.05) and *Condition* and the number of *Prompting* utterances ($F_{(2,20.32)}$=4.22, p<.05). The other main effects and interaction effects were not found to be significant at the 95% confidence level.

Further examination of the interaction effects revealed that an increase in the production of *Targeting* utterances helped to improved game score in the *Audio-only* condition, but did not provide a similar improvement on *Visual-Audio* and *Tactile-Audio*. In other words, when the pairs did not have visual or tactile feedback, they had to compensate by increasing their production of spatial information regarding the targets. We also found that the increase of *Prompting* utterances was associated with a lower score in the *Tactile-Audio* condition, but such effects were not found in the other conditions.

[1] To avoid multicollinearity, we first examined the correlations among the possible independent variables, and chose those which were not strongly correlated with each other. As a result, our variables for *Content Themes* were the utterance counts of *Targeting*, *Clarification*, *Planning*, *Prompting*, along with their interaction terms with *Condition*.

Strategies

We analyzed the naming schema that the participants used for specifying the wedges on the screen. Understanding their naming schema is important because they are a part of the participants' strategies for effective coordination. We observed the three following naming schema:

- *Index*: Language using numbers the participants agreed on. For instance, "one" meant the top wedge, "three" meant the right wedge, and "five" meant the bottom wedge.
- *Clock*: Language based on the clock metaphor. For example, "twelve o'clock" meant the top wedge, and "six o'clock" meant the bottom wedge. The use of numbers is similar to *Index*, but the wedges that the numbers correspond with here differ from *Index*.
- *Direction*: Language based on the direction or orientation. "North" or "top" meant the top wedge, and "south" or "bottom" meant the bottom wedge.

PA10-1	Time [sec]	PA10-2
"Nine." 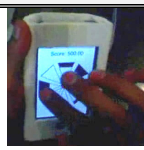 Moving to the left wedge.	46	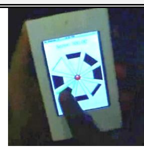 *"Seven."* Moving to the bottom-left wedge.
"What are you doing?" 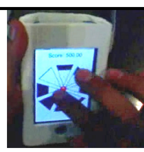 Scrubbing the left wedge, but nothing occurred.	48	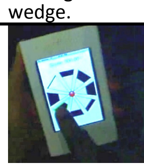 Touching the bottom-left wedge.
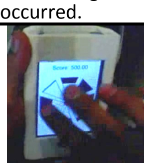 Scrubbing the left wedge.	49	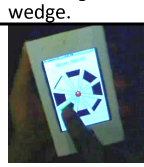 *"Seven, seven."* Holding the bottom-left wedge.
"OK." 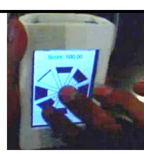 Moving to the bottom-left wedge.	50	 Holding the bottom-left wedge.
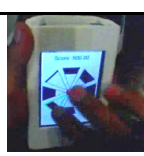 Scrubbing the bottom-left wedge.	51	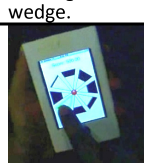 Holding the bottom-left wedge.

Figure 6. Collaboration observed in the *Audio-only* condition. This pair mis-communicated about moving to the next target. PA10-1 then explicitly asked where PA10-2 was holding and re-coordinated the position.

We observed that for most of the pairs, one participant would typically assume the responsibility of deciding the target and the other person would follow her. However, the participants often stop using this strategy when many wedges become almost completely filled. Both participants would then actively communicate about the next target. In terms of scrubbing, one person (who generally decides the target) would touch and continue to press the target while the other player scrubbed that wedge.

Collaboration in the Audio-only Condition

Participants used targeting utterances in the *Audio-only* condition most frequently as shown in Table 3. Because the system did not provide users with any feedback about their game partner's action, participants often failed to figure out their game partner's location. There were different reasons for this. For example, both players may specify different wedges as the next target, and then fail to recognize the need to negotiate or clarify which wedge both should touch (*i.e.,* once a participant expresses the next wedge she will target, she assumes the partner will follow her verbal instruction even though both in the pair spoke at the same time). Figure 6 illustrates a case in which both players

PA7-1	Time [sec]	PA7-2
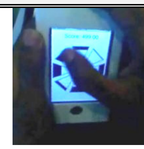 Moving to the top-left wedge. The bottom-right wedge is occluded.	119	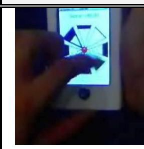 Moving to the bottom-right wedge.
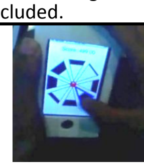 Noticing visual feedback and moving to the bottom-right wedge.	121	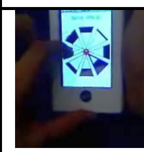 Moving to the top-right based on visual feedback.
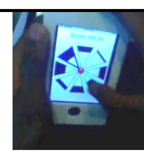 Noticing that PA7-2 had moved to the top-left wedge.	123	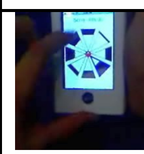 Scrubbing the top-left wedge, but producing no effect.
"Five" 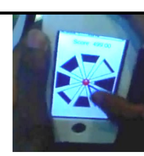 Explicitly mentioning the target.	124	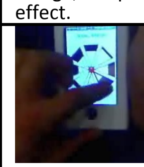 Following PA7-1 after his instruction.

Figure 7. Collaboration observed in the *Visual-Audio* condition. While PA7-1 moved to the top-left wedge, PA7-2 moved to the bottom-left. But PA07-1 did not notice PA7-2's move because the visual feedback was occluded by his thumb. This caused another mis-coordination, and PA7-1 had to tell PA7-2 to come to the bottom-right wedge.

specified two different wedges as the next target, but did not make any clarification because both were focused on their own targets. This resulted in the pair falling out of synch and required explicit re-coordination.

Participants sometimes slipped and specified a wrong target, which led them to touch different wedges. In the post-experimental interview, one participant explained that the *Audio-only* condition was difficult because of possible slips in how to refer to the wedges.

> *"Probably the most difficult condition was audio-only… I keep getting confused by 'top-right' or 'top-left'."* [PA3-1]

Although participants performed well in the *Audio-only* condition, some participants commented that explicit audio communication increased the workload involved with playing the game compared to the other conditions. As shown in Table 3, this condition forced the participants to use more utterances. Although most of the utterances we observed were short, participants felt that frequent conversations often prevented them from collaborating efficiently in this tightly-coupled coordination task.

Collaboration in the Visual Conditions

We observed that the participants generally used fewer utterances in the *Visual-Audio* condition than the *Audio-only* condition. The participants mostly made utterances for targeting and confirmation in the *Audio-only* condition and the game provided visual feedback of where the game partner was touching on the screen; thus, they did not need to perform explicit coordination using voice in some cases.

In both *Visual* conditions, the participant's hand sometimes occluded a large portion of the screen, and the participants often failed to spot the visual feedback indicating which wedge their partner was touching. Figure 7 shows one of the instances in which the occlusion impeded collaboration. While PA7-1 moved to the top-left wedge, PA7-2 moved to the bottom-right wedge. However, PA7-1 did not notice the game partner's move because the bottom-right wedge was occluded by the thumb. PA7-1 then noticed the visual feedback, and tried to follow the partner. But PA7-2 also noticed that PA7-1 was on the top-left wedge, and tried to follow him. As a result, another mis-coordination happened, and P7-1 had to specify the target verbally.

Collaboration in the Tactile Conditions

Similar to the *Visual* conditions, participants generally used fewer utterances compared to the *Audio-only* condition. The tactile feedback helped the participants identify which wedge their game partner was touching. Some participants, particularly those who had small hands, commented that they often had difficulty correctly identifying the location of the vibration; however, participants liked having a separate channel for knowing their game partner's location.

We found that the participants generally were able to use spatial tactile feedback to communicate a location with their partner and often did not need to confirm it explicitly. Figure 8 shows a common interaction that we observed in the *Tactile-NoAudio* condition. After this pair finished scrubbing the bottom wedge, PA3-1 moved to the bottom-right wedge whereas PA3-2 moved to the top-right wedge. PA3-2 immediately noticed the discontinuous vibration coming from the bottom-right. This caused PA3-2 to defer his location; PA3-2 then moved to PA3-1's location, and scrubbed the bottom-right wedge. Because PA3-1 was able to perceive that PA3-2's previous location through the tactile communication channel, after they finished the bottom-right wedge, they moved to the top-right wedge, which was PA3-2's previous intended target.

PA3-1	Time [sec]	PA3-2
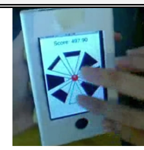	158	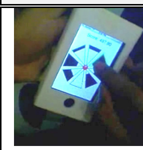
Moving to the bottom-right wedge.		Moving to the top-right wedge.
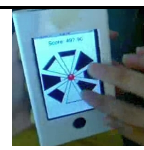	159	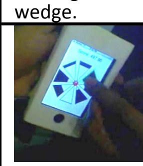
Scrubbing the bottom-right wedge, but the tactile feedback happened at the top-right.		Scrubbing the top-right wedge, but P2 noticed the tactile feedback coming from the bottom-right.
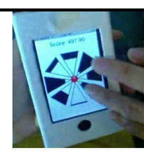	160	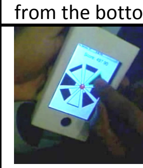
Still scrubbing the bottom-right wedge.		Moving to the bottom-right wedge.
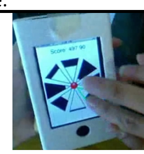	161	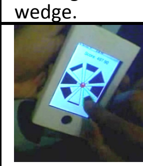
Scrubbing the bottom-right wedge. The black filling started to disappear.		Holding the bottom-right wedge.
	163	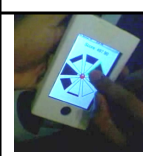
Moving to the top-right wedge which P2 was touching before.		Moving to the top-right wedge.

Figure 8. Collaboration observed in the *Tactile-NoAudio* condition. This pair made different moves initially, but PA3-2 followed PA3-1 based on the tactile feedback. After the bottom-right wedge, both participants moved to the top-right wedge which PA3-2 wanted to work on before.

STUDY1 SUMMARY

We found that a system with visual feedback generally supports coordination better than one with tactile feedback or without any feedback. Our analysis also revealed that the number of *Targeting* utterances affect game performance only in the *Audio-only* condition. The results imply that an additional feedback channel, either visual or tactile feedback, could help the participants perform spatial coordination without explicit audio communication.

The results also highlight different benefits of visual and tactile feedback, and suggest that these two feedback channels could complement each other well: visual feedback is beneficial to provide precise spatial information, and tactile feedback can address occlusion issues. We thus decided to conduct another user study to measure performance when users have both feedback channels.

STUDY2: EFFECTS OF COMBINED FEEDBACK

This user study focused on comparing user collaboration with a system combining both visual and tactile feedback against systems using either feedback channel. We used the same system and experimental procedure as the first user study. We included three audio-enabled feedback conditions: *Visual*, *Tactile*, and *Visual+Tactile*. We excluded any audio-disabled conditions because the results above already indicate that the audio channel has a substantial impact on the game performance. The presentation order of the conditions was counter-balanced across participant pairs. We recruited another twelve pairs of participants for this study (PB1–PB12). Their demographics were similar to the ones in our first user study. This study took 50 minutes on average, and each participant was compensated with $30 after the study.

STUDY2 RESULTS

Figure 9 shows the game scores across the three feedback conditions. Unpaired Welch's t-tests did not show any significant difference in the scores of *Visual* and *Tactile* between our first and second study ($t_{(21.0)}$=0.81, p>.05 for *Visual*; and $t_{(19.4)}$=0.28, p>.05 for *Tactile*). This implies that the scores in the second user study were comparable to the ones in our first study. A Mauchly's test of sphericity did not reveal a violation; thus, we once again report the results of the ANOVA F-tests. A one-way repeated-measure ANOVA revealed a significant difference in the scores by condition ($F_{(2,22)}$=6.01, p<.01). Tukey's HSD revealed that the scores in *Visual+Tactile* were significantly higher than the scores in *Visual* and *Tactile* conditions (p's <.05).

Similar to the first study, we examined the communication processes that took place to better understand the coordinating role of each feedback condition. We used the same coding scheme presented in Table 2 and the average numbers of utterances by content type for the second study are shown in Table 4. We then compared differences in the number of utterances between the first and second study. As a result, unpaired Welch's t-tests did not show any significant difference in the scores of *Visual* and *Tactile* between our first and second study ($t_{(16.1)}$=1.79, p>.05 for *Visual*; and $t_{(17.6)}$=0.80, p>.05 for *Tactile*).

We analyzed the data using the same random effects linear regression model described in the first study with one

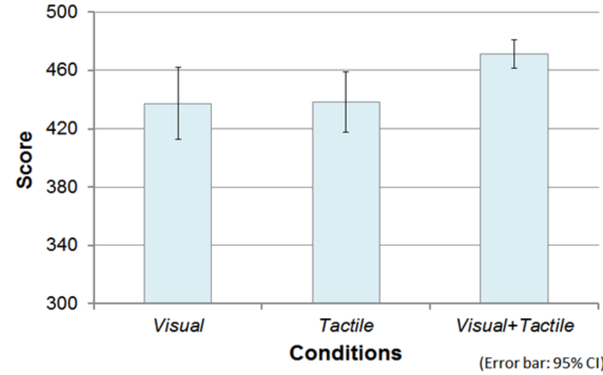

Figure 9. The average scores for the conditions tested in the second user study.

Theme	*Visual*	*Tactile*	*Visual+Tactile*
Targeting	**72.4** (26.0)	**77.3** (30.0)	**86.4** (12.2)
Confirmation	**10.5** (9.9)	**7.5** (6.0)	**5.8** (3.8)
Clarification	**0.9** (1.3)	**1.6** (2.0)	**1.8** (2.3)
Strategy switching	**1.8** (2.2)	**1.0** (1.6)	**0.9** (1.4)
Awareness	**4.3** (4.9)	**2.1** (2.6)	**1.6** (2.0)
Planning	**1.0** (1.2)	**1.7** (2.3)	**1.6** (2.1)
Prompting	**0.4** (1.1)	**0.5** (1.0)	**0.1** (0.3)
Total	**91.3** (27.0)	**91.8** (31.0)	**98.3** (11.0)

Table 4. The average number and standard deviation of

exception being that interactions were removed from the model due to the fact that there were no significant higher order interactions. The resulting model fit was relatively high (R^2=.70, $Adj\text{-}R^2$=.63).

Controlling for language, we still see a significant effect of *Condition* ($F_{(2, 21.37)}$=4.93, p<.05) where performance for the *Visual+Tactile* condition was better than both the *Visual* and *Tactile* conditions. Of the content types examined, there is a positive main effect of *Targeting* content whereby increases in the use of targeting comments was associated with higher scores ($F_{(1,26.33)}$=14.03, p<.01), controlling for condition. The other main effects and interaction effects were not found to be significant at the 95% confidence level.

While these findings suggest a link between the content of the discussions and the performance of the pairs, we found no evidence of differentiated influence of various discourse strategies in the conditions examined in the second study. This suggests that there may not be a difference in the need for additional information across the visual and tactile channels. Alternatively, it could be that our discourse coding scheme was not sensitive to the coordinating differences that are afforded by the visual and tactile channels; thus, we further analyzed the transcripts for evidence of differential use.

Most participants agreed that visual feedback was easy to understand and showed the location of a game partner accurately. But four pairs explicitly mentioned that occlusion by hands was a problem. In contrast, participants expressed different opinions on tactile feedback and used it in different ways. Four pairs mentioned that they needed more effort in associating tactile feedback to a particular wedge than visual feedback. However, tactile feedback could subtly guide a player's attention to the area of interest as one of the participants commented:

> *"The good thing [about vibration] for me was it was a more subconscious cue; wasn't something I had to pay attention to. But with color, I've got to pay attention to where the color is going, I had to process it. But vibration subconsciously pulled attention to the area."* [PB2-2]

Two pairs commented that tactile feedback offered a separate channel to maintain awareness about a game partner without disrupting the visual information space.

> *"The vibration is a much better safety net than visual. I think because I could feel that or maybe because so much was happening on the screen already. Tactile touch didn't add something extra on the screen... Changing colors was just a distraction, versus vibration wasn't another visual distraction."* [PB10-1]

Participants overall liked the *Visual+Tactile* condition, and often preferred it over the other conditions. One participant explained to us that she used tactile feedback as a redundant cue to ensure that she and her partner were contacting the same wedge.

> *"Vibration is confirmation of if you are doing the right thing. I'm hearing where to go, and vibration confirms me that we are going to the same spot."* [PB2-1]

STUDY2 SUMMARY

Our second user study revealed that a system using both visual and tactile feedback outperformed systems using only either type of feedback when verbal communication is available. Its qualitative results also support the notion that visual and tactile feedback can complement each other to support users' spatial coordination.

DISCUSSIONS

The results show that the occlusion caused by the participant's hand often impeded smooth collaboration in the *Visual* conditions. Occlusion is a well-known problem in touch-screen devices. To address this issue, Vogel and Balakrishnan demonstrated an interface which changes the locations of the objects depending on the position of the hand or arm over the screen [21]. This technique could solve some occlusion problems we observed in our user studies. But, rearranging the objects might introduce additional complexity on spatial coordination. For example, an object on the right side in one person's device might appear at a different location in the partner's device with an occlusion-aware interface. Tactile feedback can directly mitigate some occlusion problems, and provide awareness for coordination particularly when the visual components are complex (*e.g.*, online network multiplayer games).

The system had delays in tactile feedback during our experiments due to our hardware and system limitations. Gergle *et al.* studied the effects of delays in visual feedback for collaborative tasks [9]. They found that in rapidly changing dynamic environments, delays on the order of 200ms can cause performance deficits in visual piece arrangement tasks. Jay *et al.* conducted a study examining the effects of delay in haptic and visual feedback on collaborative tasks which require strict spatial and temporal coordination [12]. Their task was to move two cursors towards the target in a graphical user interface while maintaining the relative distance of the cursors within a threshold. Their results show that even a very small delay (25 msec in tactile feedback, and 50 msec in visual feedback) can impact the collaboration when strict spatial and temporal coordination is necessary.

The tasks studied in our experiment include higher temporal demand (*i.e.*, the system fills wedges constantly) than those in Gergle *et al*.'s study; thus, even a small delay might have caused a significant impact. This may be one reason why *Tactile-NoAudio* was the weakest condition in terms of the performance score because participants were unable to explicitly coordinate through the audio channel to compensate the delay in tactile feedback. However, as Jay *et al*. discussed in [12], supporting strict spatial and temporal coordination is challenging in both visual and tactile feedback, and thus future research is necessary on how to overcome the delay on the feedback.

We had the limited output resolution of spatial tactile feedback. Israr and Poupyrev have developed a method to create virtual vibration points where no physical vibration motors exist [11]. With their algorithm, a system would have an improved output resolution, enabling the design of combined feedback channel to support more precise spatial coordination. Our results, which highlight the different advantages for each type of feedback and benefits for combined feedback in spatial coordination, can inform the designs of collaborative handheld systems with higher output resolution than our system.

LIMITATIONS

There are several limitations in our study that should be considered. First, we abstracted spatial coordination tasks into a game that required participants to perform frequent spatial coordination. As a result, we were able to perform an in-depth analysis of spatial coordination processes. While the game provides a controlled environment for doing this, it is unclear the extent to which these processes will translate to more contextualized everyday interactions. In addition, there are different types of spatial coordination from what we studied, such as avoiding the same object or location. Our results suggest that the combination of visual and spatial tactile feedback can enhance the user's awareness about her collaborator's actions, and could help collaboration for other spatial coordination tasks. But, we note that the results may require careful interpretation when applied to more natural spatial coordination tasks. As such, this game may not relate directly to any specific application, and researchers should carefully consider the role that spatial coordination plays in their given task environment.

Second, the mapping of spatial tactile feedback between vibration locations and wedges is one-to-one in our game design. When this mapping becomes one-to-many, the user would need visual cues to correctly identify the precise location which both users are targeting. Spatial tactile feedback may not be beneficial in this case, but our results suggest that it can still help the user focus her attention towards a part of the workspace, and can reduce visual distraction. This can be important when the users' workspace shows a great deal of visual information.

CONCLUSIONS AND FUTURE WORK

We investigated how visual and tactile feedback affects synchronous collaborative tasks, in particular spatial coordination tasks, on mobile devices. Our results highlight the following findings in the context of collaborative handheld systems:

- Visual feedback can provide precise spatial information about collaborators, but can hamper collaboration when it is occluded and sometimes distracts the user's attention.
- Spatial tactile feedback can provide spatial information about collaborators as well, but improvements are necessary to convey precise spatial information. It can reduce the overload of information in visual space and can gently guide the user's attention to an area of interest.
- Visual and spatial tactile feedback can complement each other, and systems with both feedback channels can offer better spatial coordination support than systems using only one form of feedback.

Our results also imply that spatial tactile feedback could be used to support collaboration with visually impaired users. For example, sighted users can demonstrate to visually impaired users how to interact with interfaces on a mobile touch-screen device. The interaction will be conveyed through tactile feedback, and visually impaired users will be able to feel how to interact with a device remotely.

REFERENCES

1. Brave, S., Ishii, H., Dahley, A. Tangible interfaces for remote collaboration and communication. In *Proc. of CSCW,* ACM Press (1998), 169-178.
2. Brown, L. M., Sellen, A., Krishna, R., Harper, R. Exploring the potential of audio-tactile messaging for remote interpersonal communication. In *Proc. of CHI,* ACM Press (2009), 1527-1530.
3. Chan, A., MacLean, K., McGrenere, J. Designing haptic icons to support collaborative turn-taking. *Int. J. Human-Computer Studies 66*, (2008), 333-355.
4. Chang, A., O'Modhrain, S., Jacob, R., Gunther, E., Ishii, H. ComTouch: design of a vibrotactile communication device. In *Proc. of DIS*, ACM Press (2002), 312-320.
5. Dourish, P., Bellotti, V. Awareness and coordination in shared workspaces. In *Proc. of CSCW*, ACM Press (1992), 107-114.
6. Fogg, B., Cutler, L. D., Arnold, P., Eisbach, C. HandJive: a device for interpersonal haptic entertainment. In *Proc. of CHI*, ACM Press (1998), 57-64.
7. Fukumoto, M., Sugimura, T. Active click: tactile feed-back for touch panels. In *Extended Abstracts of CHI*, ACM Press (2001), 121-122.
8. Gergle, D., Kraut, R.E., Fussell, S.R. Language efficiency and visual technology: minimizing collaborative effort with visual information. *Journal of Language and Social Psychology 23*, 4 (2004), 491-517.
9. Gergle, D., Kraut, R.E., Fussell, S.R. The impact of delayed visual feedback on collaborative performance. In *Proc. of CHI*, ACM Press (2006), 1303-1312.
10. Hoggan, E., Brewster, S.A., Johnston, J. Investigating the effectiveness of tactile feedback for mobile touchscreens. In *Proc. of CHI*, ACM Press (2008), 2019-2028.
11. Israr, A., Poupyrev, I. Tactile Brush: Drawing on Skin with a Tactile Grid Display. In *Proc. of CHI*, ACM Press (2011), 1573-1582.
12. Jay, C., Glencross, M., Hubbold, R. Modeling the effects of delayed haptic and visual feedback in a collaborative virtual environment. *ACM Trans. Comput.-Hum. Interact. 14,* 2 (2007), 8.
13. McGookin, D., Brewster, S. An initial investigation into non-visual computer supported collaboration. In *Extended Abstracts of CHI*, ACM Press (2007), 2573-2578.
14. Oakley, I., Brewster, S., Gray, P. Can you feel the force? An investigation of haptic collaboration in shared editors. In *Proc. of EuroHaptics*, EuroHaptics Society (2001).
15. Palmer, C. I., Gardner, E. P. Simulation of motion of the skin IV responses of pacinian corpuscle afferents innervating the primate hand to stripe patterns on the optacon. *J. Neurophysiol. 64*, 1, (1990), 236-247.
16. Pielot, M., Krull, O., Boll, S. Where is my team? Supporting situation awareness with tactile displays. In *Proc. of CHI*, ACM Press (2010), 1705-1714.
17. Poupyrev, I., Maruyama, S., Rekimoto, J. Ambient touch: designing tactile interfaces for handheld devices. In *Proc. of UIST*, ACM Press (2002), 51-60.
18. Plimmer, B., Crossan, A., Brewster, S. A., Blagojevic, R. Multimodal collaborative handwriting training for visually impaired people. In Proc. of *CHI*, ACM Press (2008), 393-402.
19. Sahami, A., Holleis, P., Schmidt, A., Hakkila, J. Rich Tactile Output on Mobile Devices. In *Proc. of AmI*, Springer (2008), 210-221.
20. Stark, R., Carlstedt, T., Hallin, R. G., Risling, M. Distribution of human pacinian corpuscles in the hand. *J. Hand Surg. Eur. 23B*, 3, (1998), 370-372.
21. Vogel, D., Balakrishnan, R. Occlusion-aware interfaces. In *Proc. of CHI*, ACM Press (2010), 263-272.
22. Yatani, K., Truong, K., N. SemFeel: A user interface with semantic tactile feedback for mobile touch-screen devices. In *Proc. of UIST*, ACM Press (2009), 111-120.