Feeling Aware: Investigating the Use of a Mobile Variable-Friction Tactile Display for Awareness Information

Jeremy Birnholtz
Northwestern University
Evanston, IL USA

Darren Gergle
Northwestern University
Evanston, IL USA

Noah Liebman
Northwestern University
Evanston, IL USA

Sarah Sinclair
Cornell University
Ithaca, NY USA

{jeremyb, dgergle}@northwestern.edu, nliebman@u.northwestern.edu, sjs334@cornell.edu

ABSTRACT

Providing the initiator of a conversation with awareness information about their target's availability can improve interruption timing. Providing this information, however, can be particularly challenging with mobile devices, which are carried by users in a range contexts and have limited audible and visual display possibilities due to possible background noise and small screen size. Haptic or tactile displays present a potentially useful alternative, as interaction with mobile devices is often by touch. This paper reports on a dual-task comparison of awareness information usage on a mobile variable-friction tactile display vs. a visual display under conditions that varied in task type and task workload. Participants using the tactile display performed marginally better on the primary task, but were less accurate and slower on the awareness task. However, there is evidence that some of the awareness task differences dissipated over time. Self-report data also suggests that people's experience with the tactile display was generally positive and improved over time.

Author Keywords

Awareness displays; availability; interruption; tactile; variable friction; CSCW

ACM Classification Keywords

H.5.2. User Interfaces: Haptic I/O

INTRODUCTION

Unscheduled and spontaneous interactions have repeatedly been shown to play a key role in maintaining social and work relationships [31], coordinating [26], and answering timely questions [23]. On the one hand, the ubiquity of mobile devices and connectivity today can help facilitate these interactions, as almost anybody can be reached at almost any time or place. On the other hand, the always-on, always-connected world can also leave people vulnerable to interruption. This has spawned concerns that people are attending more to their devices than to each other [36]. It also highlights the importance of strategies for managing when one is available [e.g., 7], and for being sensitive to others' context when interrupting [e.g., 20].

There is evidence, however, that today's tools for managing availability and interrupting sensitively are inadequate. Many people resort to relatively extreme tactics, such as turning off their device altogether [40], or downloading tools that disable social applications to minimize distraction [34]. Given that interruptions can be distracting [2,12,27], it is unsurprising that people seek to avoid them.

One likely reason that so many interruptions are avoided is that few common tools effectively provide availability information to the interrupter about their target. There is evidence, however, that when this information is available people will use it to time their interruptions so they occur during less busy periods and are less distracting [13]. One potentially important area of mobile HCI research, therefore, is considering mechanisms for providing availability information to initiators of interaction.

A common scenario in both work and social contexts (and that has been replicated in experimental studies such as [5,13,28]) is one in which the initiator of interaction needs to ask the interruption target a quick, timely question. They do not want to distract the target from whatever he or she is doing if it is important, but do want to ask the question at the earliest opportune moment. It can be useful here to have an ongoing sense of the target's status or availability, so that the initiator can know when this status changes.

Mobile devices provide unique attributes and challenges in providing availability or awareness information. Screen space is scarce, so it may not make sense to use a visual display that will take up space or be easily occluded. Audible displays [e.g., 9,21] can be distracting to others near the device so may be inappropriate in public or quiet contexts; alternatively, in noisy environments such as a New York City street corner, they may not be heard at all.

One display modality that has not been explored extensively is providing availability information via tactile feedback (i.e., a "tactile" display). Tactile sensation is particularly appropriate in the mobile context, as people

MobileHCI '15, August 25 - 28, 2015, Copenhagen, Denmark Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3652-9/15/08…$15.00
DOI: http://dx.doi.org/10.1145/2785830.2785857

interact with mobile devices primarily via touch. There is also some evidence that tactile display can help convey emotional [22] and ludic [29] information; and that tactile information can aid in coordination tasks [e.g., 10,39]. Given the wide array of available display mechanisms coupled with the variety of ways in which users interact with a surface device (e.g., varying touch interaction profiles such as repeated taps for typing a text message vs. dragging to scroll or play a game), there are many open questions about whether, how and under what circumstances it may be effective to use tactile display for providing availability information.

In what follows, we present an exploratory laboratory experiment using a tactile display to provide availability information in a collaborative dual-task scenario. We focus in particular on whether this is a viable approach at all, and also compare it to a visual display across three different tasks with two levels of workload.

BACKGROUND

Research on awareness displays has focused on helping people benefit from useful informal interactions in work and social settings [23,31], while avoiding the drawbacks of distraction [2]. This is typically achieved by providing information to the interrupter about the availability of their target. This involves three questions.

First is the question of what information to share and how. Many common tools (e.g., Skype) rely on the individual user to manually set their status to "available" or "busy," but many people do not regularly do so. With this in mind, Fogarty et al. [15] and Begole et al. [3] experimented with gleaning contextual information via sensors (e.g., door open vs. closed) from people's workspaces and using this to predict availability. The details of this problem are beyond the scope of this study, but we assume here that some reliable indicator of target availability is available.

Second is the question of how much detail to share, which raises issues both of target privacy [11] and distraction. Dabbish and Kraut [13] compared interruptions using an awareness display that provided a full screen shot of the target's activities with an abstracted summary. They found that the abstracted summary was just as effective in terms of timing interruptions, but less distracting.

Third is the question of how to display the information. Most explorations of awareness have focused on visual displays (see [35] for a review) on the interrupter's primary screen. Birnholtz et al. [5] also experimented with using a peripheral visual display to provide information in a less intrusive way by extending beyond the primary screen.

Non-Visual Displays

As we noted earlier, displaying awareness information on mobile devices is constrained by small screens and often limited visual attention as the device may be used while engaged simultaneously in other activities. This raises the possibility of using non-visual displays, but this space has not been extensively explored. There is some evidence, however, that the addition of non-visual feedback to visual displays can improve task performance [9].

Some early work in desktop contexts experimented with audible awareness information [e.g., 18,21], but this has not been explored extensively in the mobile context. However, as previously discussed, the mobile environment may yield contexts such as overly noisy or quiet social environments that limit the utility of audible awareness information.

Another type of display is haptic or tactile, where information is sensed via touch. Given that touch is the primary way people interact with their mobile devices, tactile display could be useful in addressing the problem of providing availability information [33]. For example, Chan and colleagues developed a set of haptic icons using a vibrotactile mouse that people used to request control of a single-user application. They found that haptic requests for control resulted in faster responses than visual requests [10]. Yatani and colleagues [39] focused on spatial awareness by targeting vibration at a location on a device where a remote collaborator is working. In addition to improving performance, this approach was received positively by participants. Pielot and de Oliveira [30] found that a vibrating pulse on a pocketed device can provide awareness with minimal distraction. Other uses of haptic technology in communication have included supporting emotional connection in long-distance relationships [22] or playfulness in everyday interactions [29].

Variable Friction and Touch Profile

One novel non-visual display technique that has potential for providing awareness information in the mobile context is the variable friction display [38]. These are touchscreen displays where the perceived tactile friction can be electronically manipulated. Variable friction is importantly distinct from vibrotactile technology, which vibrates an entire device, and is commonly used for notifications. The sensation of texture on the screen is subtler than vibration, such that one can use the device normally without being overwhelmed or distracted by the texture.

By varying the friction of a display surface such that there are regions of high friction and low friction, it is possible to create the sensation of textures on a smooth surface [38]. Variable friction displays can aid in basic touchscreen tasks like target acquisition and drag-based selection [25]. There is also some evidence that variable friction can add an element of interpersonal connection to multiuser apps [29].

One potential problem with this approach, however, is that friction only occurs when adjacent objects are in contact and in motion. Thus, manipulations in friction are detectable only when the user's finger is in motion and in contact with the display. The appropriateness of this technology may vary as different tasks have different patterns of touches [24].

Workload

In addition to display characteristics, the level of user engagement required to complete the task can affect how much attention is available to perceive displays of awareness information [13,14]. For our question, this presents two possibilities. On the one hand, multiple resource theory suggests that input from one sense may not interfere significantly with input from other senses, because processing the two inputs uses different cognitive resources so they are processed in parallel [37]. If this is the case, people should be able to process a tactile awareness display separately from a simultaneous visual task, resulting in a performance advantage. Hopp and colleagues [19] showed evidence of this, as haptic interruptions were minimally disruptive to participants engaged in a visual task.

On the other hand, cognitive resources are finite and variable friction displays can be subtle in implementation. It is possible that people who are sufficiently engaged in a visual task may not notice variations in the display surface. This question is important, as there is little sense in displaying information people will not notice.

THE PRESENT STUDY

We ran a controlled laboratory experiment comparing a variable-friction tactile notification display to a more conventional visual display on the same device. In addition to display type, primary task type and workload were manipulated. In all conditions, participants engaged in a primary task, and were simultaneously asked to respond to notifications presented via tactile or visual display. Participants used the TPad (see Figure 1), a variable-friction display that augments a Nexus 7 tablet [38].

Figure 1. TPad variable-friction display used for the study.

Our exploration was guided by three broad research questions. We were first interested in the effects of display type (visual vs. tactile), task type and task workload on participants' performance. As in prior work on awareness and notification displays [e.g., 13], there are two dimensions of performance to consider: whether the participants notice the awareness display and whether it distracts them from their primary task in ways that are detrimental to their performance. We asked,

RQ1: What are the effects of display type, task type and task workload on participant primary task performance and their ability to notice incoming notifications?

Second, we were interested in understanding the underlying process dynamics of effects on performance and response to incoming notifications. Mobile devices are used in a wide range of contexts, and different tradeoffs between performance, noticeability of notifications and other key factors may be appropriate in these different contexts. For example, people may be willing to tolerate more errors in situations where they must subtly feel their phone out of view (e.g., under a table during a meeting) if the alternative is not receiving the notifications at all. Understanding the dynamics of task effects and response can help us better understand when these display techniques may be appropriate, and can inform the design of subsequent systems by providing the details of why and when tactile notification may be appropriate. We asked,

RQ2: What are the effects of display type, task type and task workload on the underlying dynamics of task performance and response to incoming notifications?

Finally, we were interested in participant experience with the tasks and a variable-friction tactile display, as this is not yet a common technology. We wondered about their impressions of the technology, its effects on their performance, and whether it could be used more broadly. We asked,

RQ3: What do participants think of variable-friction tactile display and what effects do they perceive it having on their performance and experience?

Method

Our experiment used a 2 × 2 × 3 within-participants design including three games in which task type and workload (see Table 1) were manipulated, and simultaneously responding to "interruptions" from a fictional partner.

Independent Variable	Levels	Operationalization
Display type	Tactile	Texture using variable friction
	Visual	Colored glow around the edge of the display
Task type	Sorting game	Task with shorter touch profiles and more frequent touches
	Shuffleboard game	Task with medium length touch profiles and more frequent touches
	Catching game	Task with longer touch profiles and less frequent touches
Task workload	Low	Fewer objects to be dealt with
	High	More objects to be dealt with

Table 1. Experimental conditions manipulated in the study.

Participants

Participants (N = 56, 54% female; M = 23 years old) were students, staff, and members of the surrounding community of a mid-sized Midwestern US university. They were recruited using paper fliers and word of mouth, and compensated $15 for their time.

Independent Variables and Experimental Manipulations

Display Type: Notification display type in each condition was either tactile or visual. Display in the tactile condition consisted of varying the perceived friction of the entire display surface. When there was no notification, the display was configured to feel smooth (as it normally would if not manipulated). To notify participants of a change in awareness state, the display was configured to feel rough in texture. The specific texture was chosen based on pilot testing and prior work suggesting that shorter spatial periods are more easily detected [4]. We opted for a spatial period of 0.16 cm, meaning that the perceived friction changed from high to low and back every 0.16 cm as the participant's finger moved on the display. It is not strictly necessary to alter the texture of the entire display, but we did so because our tasks used the entire display surface.

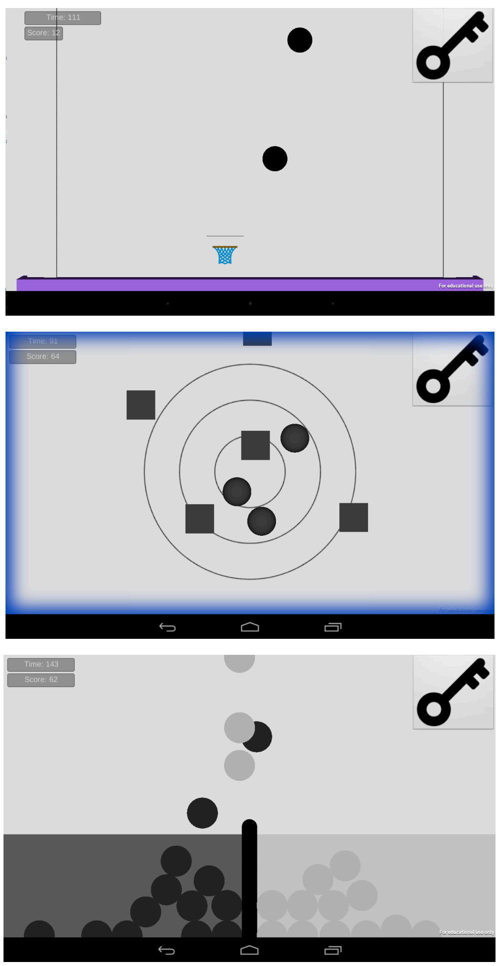

Figure 2. Three games developed to highlight different task characteristics and touch profiles: Catching (top panel), shuffleboard (middle panel), and sorting (bottom panel).

The visual notification display was a colored glow around the edge of the screen (see Figure 2, middle). We aimed to design a visual awareness display that, like the tactile awareness display, was noticeable, yet did not occupy the same physical space as the primary tasks.

Task type: Task type was manipulated via three different games that aimed to capture the different ways that touch on mobile devices can occur, including long touches/continuous drags, short-to-medium touches/quick drags and short–to-medium touches/flicks, as well as changes in the frequency of touches across tasks. Points were awarded in each game to incentivize performance.

The catching game (see Figure 2, top) involved catching falling shapes in a basket that could be dragged from right to left, with the goal of catching as many as possible. This game was modeled roughly on tasks used in experiments by McFarlane [28] and Dabbish and Kraut [13]. Points were awarded for each shape successfully caught.

The shuffleboard game (see Figure 2, middle) involved shapes that emanated from the center of the screen that had to be dragged back to the center before they reached the edge of the screen. More points were awarded for shapes that were caught closer to the center and dragged back.

The sorting game (see Figure 2, bottom) involved sorting shapes by color such that black shapes fell to the left side of the screen and gray shapes fell to the right side.

A manipulation check revealed that the average duration of touch events in the catching game (16.6 sec) was very different from the shuffleboard and sorting games, whose average touch durations were closer (0.337 sec and 0.307 sec, respectively). The number of touches also varied across task types with the catching game exhibiting the fewest average touches (35), followed by the sorting game (237 touches) and the shuffleboard game (254 touches).

Workload: Workload (low or high) was manipulated by varying the number of on-screen shapes that required interaction. To keep a consistent number of objects on screen in the low and high workload conditions (i.e., to control for the number of visual objects), we took advantage of the "pop-out" effect where a group similar objects can be attended or ignored with relatively little effort [17]. In our case, we used two shapes (i.e., circles and squares) in the low-workload condition and told participants to interact only with one type. In the high-workload condition, all shapes were the same (i.e., only circles) and participants interacted with all of them. While this may seem counterintuitive, the effect was to lower the workload without increasing the difficulty of visual search or reducing the number of shapes on the screen. A manipulation check based on participant assessment of workload via the NASA Task Load Index (TLX) confirmed this with the high workload condition resulting in increased perceptions of task load and decreased perceptions of performance when compared to the low workload condition, even when controlling for game type ($p < 0.01$).

Procedure

Upon arrival in the laboratory, participants were told they would be playing a series of games with a partner who was located in another room, with the objective of cooperatively earning as many points as possible. In actuality, the partner was fictional, but we included this in the design to replicate real-world scenarios in which timely interruptions can be helpful, but poorly timed interruptions are disruptive. Participants were told:

"You, and not your partner, have a special job that will allow you to score more points as a team. Your partner is playing a similar game, except their rounds aren't quite as long as yours. In order to advance to the next round, they either need a key from you, which you can send to them by pressing a button, or they need to wait 20 seconds. This is where you have to cooperate. If your partner is in between two rounds when you send them the key, they get to go on and keep earning points for both of you. If you send the key while your partner is in the middle of a round, though, it disrupts them for several seconds and this will end up costing you both points. And if you don't send the key at all, your partner is missing opportunities to earn points while waiting for the next round to start."

Depending on the condition, one of the awareness displays described above (tactile or visual) indicated to participants if their partner was available to accept the key. They key was sent by pressing the 'key' button visible in the upper right corner of the games (see Figure 2).

After each round, participants saw their own score, their partner's score, and—to emphasize cooperation—the sum of the two scores. The participant's own score was the number of objects they correctly interacted with (e.g., catching a ball and missing a distractor) minus the number of objects they incorrectly interacted with (e.g., missing a ball and catching a distractor). Their partner's score was based on their own score, but with points subtracted for each time the participant pressed the key button at the wrong time, and with a smaller number of points added each time the participant pressed the key button at the right time. The partner score was then shifted by a random margin within 10% to obscure the exact relationship between it and the participant's score.

To allow participants to experience the games and display types, they then played each of the three game types for a one-minute practice round. During these practice rounds, the availability state toggled on and off every five seconds.

Following the practice rounds, participants provided demographic information in a brief questionnaire and experienced each of the 12 experimental conditions (2 awareness display types × 2 workload levels × 3 task types), which lasted 3.5 minutes each. The order of conditions was counterbalanced to ensure that any learning effects would even out across conditions. During these rounds, availability would shift to the "available" state every forty seconds, and stay on for ten seconds, or until the participant sent the key to their partner.

During all rounds, participants wore headphones playing white noise to eliminate any possibility of a confound from the user detecting the small, high-pitched sound that is made when they make contact with a textured screen. To minimize fatigue, participants were allowed to take up to a five-minute break after completing the sixth round.

Measures

To answer our research questions, we examined outcomes related to performance, process and experience.

Performance Measures. We measured performance by examining participants' scores in each of the three games. As each game type involved different numbers of objects, scores were normalized so they could be compared.

Process Measures. We examined whether or not people timed their interruptions (i.e., key button presses) according to their partner's availability state as indicated on the display. We measured the detectability of changes in availability as the time between the availability state change and the pressing of the key button. Participants' attempts to send the key to their partners were also categorized as *good* or *bad* based on whether or not they occurred while their partner was available. A *missed* button press is when the key button was not pressed during the 10-second availability window. In the language of binary classification, good presses are true positives, bad presses are false positives, and missed presses are false negatives. An F_1 score, ranging from zero to one, is a common way to combine these into a single measure [32]. A score of one indicates that all interruptions were correctly timed relative to availability.

Experience Measures. After each round, participants completed a questionnaire to assess their workload, their own and their partner's effectiveness, their perceptions of the display and the perceived connection between partners.

Workload was evaluated using the NASA Task Load Index (TLX) [16]. On a scale from very low (0) to very high (10), the NASA-TLX measures users' perceived mental demand, physical demand, temporal demand, performance, effort, and frustration in engaging with a task. We used the NASA-TLX in raw form, meaning that we aggregated it into a single scale without weights ($\alpha = 0.85$).

Partner connection and performance were measured using an instrument adapted from [5]. It asked questions like "My partner and I cooperated on this task effectively," and "I was able to determine when it would be a good time to interrupt my partner," and was measured using a seven-point scale from strongly disagree (1) to strongly agree (7).

Statistical Analysis

Our analytical approach used mixed effects models. The independent variables included were display type, task type, workload, and all two-way and three-way interactions. Because each participant performed in every condition and observations were repeated and not independent, participant was modeled as a random effect.

RESULTS

Our experiment was designed to examine the potential effectiveness of using a tactile display to provide availability information during a collaborative task. To evaluate this potential, we measured three primary dimensions: performance (game score and interruption timing), process (reaction time), and self-reports capturing aspects of the participant experience.

Task Performance

One measure of task performance is the game score achieved by the participants for each of the various conditions (see Figure 3). Participants exhibited marginally better performance when they received notifications via the *tactile display* (M = 1.85, SE = .04) versus the *visual display* (M = 1.80, SE = .04), $F_{(1, 59.35)} = 3.06$, $p = .09$.

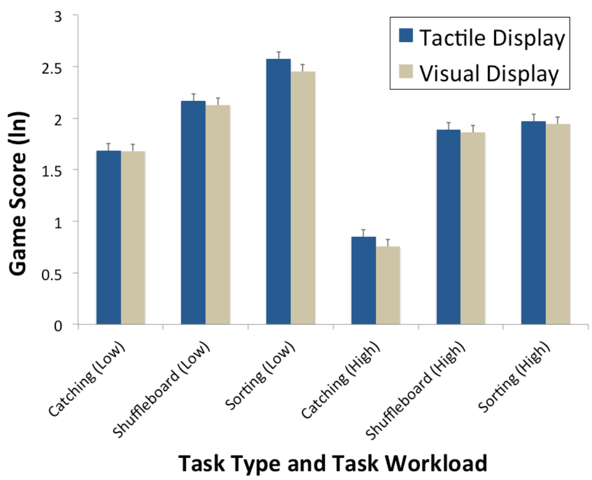

Figure 3. Game score by display type, task type and task workload (error bars represent +1 SE).

The task workload manipulation had a significant influence on performance whereby the *high workload* (M = 1.54, SE = .042) was associated with lower performance scores than the easier *low workload* trials (M = 2.11, SE = .042), $F_{(1, 59.18)} = 203.59$, $p < .0001$. There was no evidence of an interaction between display type and workload, $F_{(1, 59.13)} = .010$, $p = .92$.

Participants achieved a higher score in the sorting game with shorter duration touches (M = 2.34, SE = .046) compared to the shuffleboard game with longer touches (M = 2.01, SE = .046; for the comparison $F_{(1, 116.3)} = 20.56$, $p < .0001$), which in turn was better than the catching game with longer touches (M = 1.24, SE = .046), $F_{(1, 116.6)} = 235.10$, $p < .0001$. There was no evidence of an interaction between display type and task type ($F_{(1, 59.13)} = .01$, $p = .92$).

Another aspect to consider when evaluating the potential effectiveness of an availability display is whether or not the participants accurately shift their interruptions to times when their partner is available. In our task, this means that participants would select more appropriate times to send the "key" to their partner (i.e., "good presses," see above).

Overall, performance on this metric was quite good with successful presses occurring over 75% of the time in both the tactile and visual display conditions. The number of good presses per round was non-normally distributed and skewed such that the median was 5 (out of 5 possible - with a M = 4.72, SD = 0.664) in the *visual display* condition. In the *tactile display* condition, performance was, however, somewhat lower, with a median of 4 (M = 4.04, SD = 1.31).

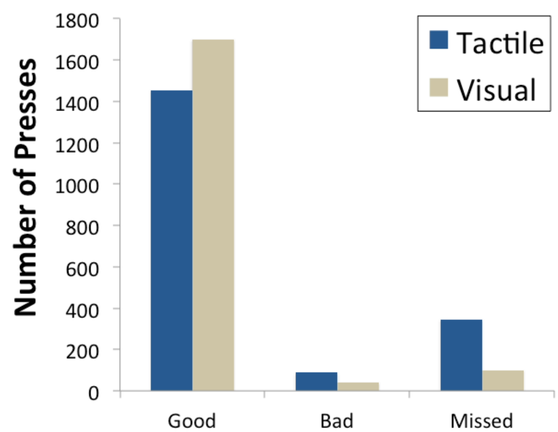

Figure 4. Histogram of "good," "bad," and "missed" presses by display type.

While the majority of presses were good with both the visual display and the tactile display, the tactile display trials appeared to produce a slightly higher proportion of bad presses and many more cases of missed presses (as shown in Figure 4). We modeled the F_1 score, described above, as a function of display type using a Poisson regression because the distribution of scores reflected that it was derived from count data. On this measure, participants performed better on the *visual display* trials (M = 0.96) than on the *tactile display* trials (M = 0.88), ($\chi^2_{(1, N = 699)} = 12.6$, $p < .001$), supporting the notion that interruptions were, in the aggregate, better timed with the visual display than with the tactile display.

We also examined if the effect of display type on interruption timing varied as task types changed and workload increased. However, we found no evidence of an effect of task type (p = .318) or workload (p = .794), and no evidence of any two-way or higher order interactions.

Task Process

The analyses thus far suggest that tactile displays perform marginally better than visual displays on the primary task performance measure, but are worse than visual displays on interruption timing. The main difference on interruption timing may be due to an inability to detect the awareness indicator.

To further explore this, we examined the reaction time, or the speed with which participants responded to a notification, as one way of operationalizing detectability: someone should be able to respond more quickly to something that is more readily noticeable.

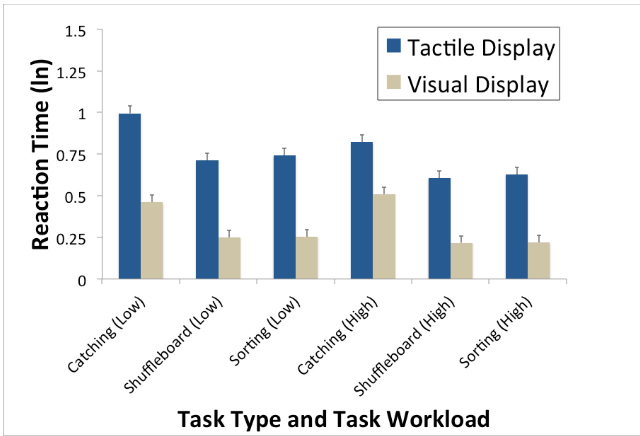

Figure 5. Reaction time by display type, task type and task workload (error bars represent +1 SE).

The overall average reaction time was 1.87 seconds (SD = 1.04), and Figure 5 shows the average reaction times broken out across the various conditions. (The data were transformed with the natural log to correct for increasingly large residuals and to minimize skew; however, the same pattern of results was found with non-transformed data).

The participants were faster to respond to the notification in the *visual display* (M = .317 seconds (ln), SE = .029) trials than they were in the *tactile display* (M = .757 seconds (ln), SE = .029) trials, $F_{(1,\ 57.5)} = 176.29$, $p < .0001$. In other words, participants were about 758 milliseconds slower to respond in the tactile display condition than in the visual display condition.

The reaction time also varied depending on the task type and particular features of the game. Participants were quickest to respond to notifications when playing either the sorting game (M = .460, SE = .028) or shuffleboard game (M = .450, SE = .028), in comparison to the catching game (M = .700, SE = .028; ps < .0001 for both contrasts). It is important to note that reaction times appear slower for the catching game which requires less frequent but longer duration touches, which highlights the role that the general interaction properties of the task can have on performance. There was no evidence of any other two- or three-way interactions that included task type.

Workload also had a significant influence on reaction time whereby the *high workload* trials (M = .502, SE = .026) were associated with quicker responses than the trials performed with a *low workload* (M = .572, SE = .026), $F_{(1,\ 57.69)} = 11.93$, p = .001. It is also important to note that there was an interaction (see Figure 5) where the higher workload condition led to improved performance for the tactile display, but made little difference for the visual display, $F_{(1,\ 57.1)} = 14.58$, $p < .001$. We further discuss this seemingly counter-intuitive result in the discussion section where we suggest that it may be an attentional mechanism at play. In other words, the participant's finger must be in contact with and moving over the surface in order to be "attending" to the notification – something that is increasingly likely with a high workload. As a result, the tactile display is associated with faster reaction times when put under a high workload. In fact, if we only examine the cases where the participant's finger was in contact with the screen, the average difference in reaction time drops to 499ms.

Another important process aspect to investigate for a novel interaction technique is whether differential learning takes place over the course of the study. Recall that each participant played twelve experimental rounds during the experiment. We tested the differences in interruption timing during the first third and last third of rounds of the experiment, and found that the difference between the tactile and visual display gets smaller over time. During the first four rounds, interruption timing was significantly better with the visual display than the tactile display ($\chi^2_{(1,\ N\ =\ 233)} = 5.37$, p = .02). During the last four rounds, this difference is lessened ($\chi^2_{(1,\ N\ =\ 234)} = 3.16$, p = .08). While not conclusive, it may suggest that participants are learning to recognize the notification feedback better over the course of the study.

Participant Experience with the Tactile Display

The participants also reported on perceived effectiveness of the tactile and visual displays. When examining the results in Table 2, the first thing to note is that for all measures the responses are significantly above the neutral sentiment midpoint for both the tactile and visual display (i.e., the 95% CI's are on the positive side of the response scale and do not cross the midpoint). We take this as an indication of the potential promise of the novel tactile display technology.

When comparing the responses between the tactile and visual displays, the results favor the visual display. As seen in Table 2, there were clear aspects of the interaction upon which the visual display was rated higher than the tactile display: feeling that the cooperation was effective ($p < .05$), knowing when to interrupt ($p < .05$), paying attention to their partner's availability ($p < .01$), and the effort expended to figure out when to interrupt a partner ($p < .01$).

The most promising results for the tactile display are the interaction effects of display type and round number (i.e., time). These interactions reveal diminishing differences over the duration of the study for feeling that the cooperation was effective (p < .05) and marginally for paying attention to their partner's availability (p < .10). In other words, while the ratings on these questions differed in the earlier rounds they converged in the later rounds. Similar to the reaction time results, this may suggest that participants gained confidence with the tactile display as the experiment progressed.

			F Ratios	
	Visual M(SE)	*Tactile M(SE)*	*Display Type*	*Display Type X Round*
"My partner and I cooperated on this task effectively"	5.63 (.165)	5.44 (.165)	6.3*	5.83*
"I was able to determine when it would be a good time to interrupt my partner"	5.27 (.23)	4.82 (.23)	10.86**	.685
"I paid attention to my partner's availability"	5.73 (.206)	5.56 (.207)	8.02**	3.69†
"I spent a lot of time figuring out when to interrupt my partner" (reversed)	5.81 (.184)	5.35 (.183)	10.33**	.291

† *p* < .10, * *p* < .05, ** *p* < .01

Table 2. Self-report data on aspects of the interaction experience. The table presents means and differences by display type as well as changes in the difference over time.

DISCUSSION

In this exploratory study we set out to answer two major questions: (1) is it viable to use variable-friction tactile displays to provide availability information in collaborative tasks? (2) If this approach is feasible, how does it compare to the more common approach of providing visual availability information? Overall we feel the evidence suggests the answer to the first question is a definitive yes; while the answer to the second question is less definitive and ultimately depends on the task and design space.

Overall, we found that primary task performance (i.e., game score) was marginally better with the tactile display. However, on the secondary task (i.e., sending the key to their partner) the participants were faster and more accurate with the visual display. Looking across these two results, it appears there is a tradeoff whereby the tactile notification display was less disruptive to the primary task, but may have been too subtle and as a result led to slower responses and more missed opportunities for notification.

Digging deeper, the first thing to note is that while performance on the secondary task with the tactile display was not as good as the visual display, the participants still produced successful presses more than 75% of the time (with the bulk of the errors being "missed" opportunities as opposed to "bad" interruptions). They also reacted in under 2.5 seconds to the tactile notifications. Keep in mind that in the tactile display condition the participant's finger must be in contact with the screen and moving in order for them to be "attending" to the notification. This requirement will present a significant challenge when applying the tactile feedback technique to anything that requires constant attention since there will inevitably be times during a task when moving contact with the surface is not occurring. That said, the response speeds exhibited are likely to be quick enough for many notification and awareness-based tasks, especially for more loosely coupled collaborations where extremely fast reaction times may matter less.

Another aspect to consider is the extent to which participants' use of the tactile display improved over time, as—unlike the visual display—this was a novel technology that most of them had not experienced. Our results showed that reaction times with the tactile display improved as the experiment progressed, suggesting that familiarity may improve performance even further than we observed in this relatively short trial. However, additional research and is necessary to substantiate this possibility.

In terms of user experience, participants generally responded positively to both the tactile and visual displays. While they judged the visual display to be more effective for knowing when to interrupt and attending to their partner's availability, there was some evidence that the differences on dimensions of coordination narrowed as the experiment progressed.

On the whole, while this was admittedly not a clear win for the tactile display, we nonetheless believe our results to be encouraging. The tactile display appeared to be less disruptive to the primary task, yet there were several areas where the visual display was preferred. As previously discussed, performance was relatively strong overall and participant experience was generally positive.

From a design standpoint, our results suggest the viability of a tactile awareness display, particularly in situations where a visual or audible awareness display on a mobile device is simply not practical or possible. Such situations might include those: where the full screen is required (e.g., certain collaborative gaming or simulation scenarios), where visual attention is sporadic but one is touching the device regularly (e.g., using the device while walking or otherwise multi-tasking), or where visual attention is not possible due to visual impairment of the user or very bright lighting conditions. While further research is needed to assess the details of design and feasibility, our results suggest that these scenarios are worth exploring in that it may be better to have a tactile display than no display at all.

LIMITATIONS AND FUTURE WORK

As with any study of this nature, this work has limitations that urge cautious interpretation and provide substantial opportunities for future work.

One possible limitation is our operationalization of task type and task manipulations. On the one hand, these manipulations were strong enough to show clear differences on the TLX scale and in how participants interacted with the display surface. On the other hand, it is possible that these manipulations were not strong enough to cause the type of significant sensory interference that might impede game performance or reaction to an awareness display. This does not discount our findings from a practical standpoint, as many real-world tasks do not involve high levels of sensory interference, but it does limit our ability to make claims about the effects these factors might have on performance more generally.

A second important limitation relates to an inherent difference between the display technologies. Changes in visual display can be perceived any time the user is in the vicinity of the display and attending to it. Tactile display, on the other hand, requires active touching of the surface to be perceived. Thus, it could be that participants more easily perceived the visual display, particularly when they were not touching the display. Additional research could use eye tracking and more detailed touch logging with similar tasks to better tease apart these effects.

Finally, this work considered visual and tactile awareness displays separately, but multimodal feedback in the form of a joint visual-tactile display could provide performance improvements over either one individually [9,39].

CONCLUSION

We have presented an experiment comparing tactile awareness displays using variable-friction technology to a more conventional visual display on the same mobile device. Participants using the tactile display performed marginally better on their primary task, but were somewhat less accurate and slower in responding to awareness information than when using the visual display. However, there is evidence that some of the awareness task differences dissipated over time. Self-report data also suggests that people's experience with the tactile display was positive and improved over time. This suggests that tactile displays of awareness information present a viable technique for providing awareness information, but that there are clear tradeoffs in terms of accuracy and detectability that should be considered by designers.

ACKNOWLEDGEMENTS

We thank the reviewers along with members of the Social Media Lab and CollabLab for their insightful feedback on this work. We also thank Prof. Ed Colgate, Joe Mullenbach and Craig Schultz for the T-Pad hardware used in this study and their assistance with it; and Jonah Ruffer for software development. This work was funded in part by National Science Foundation grants #0953943 and #1217143.

REFERENCES

1. Bailey, B.P. and Iqbal, S.T. Understanding Changes in Mental Workload During Execution of Goal-directed Tasks and Its Application for Interruption Management. *ACM Transactions on Computer Human Interaction (ToCHI)*, *14*, 4 (2008), 21:1–21:28.
2. Bailey, B.P. and Konstan, J.A. On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. *Computers in Human Behavior 22*, 4 (2006), 685–708.
3. Begole, J. "Bo," Matsakis, N.E., and Tang, J.C. Lilsys: Sensing Unavailability. *Proceedings of CSCW 2004,* ACM Press (2004), 511–514.
4. Biet, M., Casiez, G., Giraud, F., and Lemaire-Semail, B. Discrimination of Virtual Square Gratings by Dynamic Touch on Friction Based Tactile Displays. *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, (2008), 41–48.
5. Birnholtz, J., Bi, N., and Fussell, S. Do you see that I see?: effects of perceived visibility on awareness checking behavior. *Proceedings of CHI 2012,* ACM Press (2012), 1765–1774.
6. Birnholtz, J., Reynolds, L., Luxenberg, E., Gutwin, C., and Mustafa, M. Awareness Beyond the Desktop: Exploring Attention and Distraction with a Projected Peripheral-vision Display. *Proceedings of Graphics Interface 2010*, Canadian Information Processing Society (2010), 55–62.
7. Birnholtz, J., Reynolds, L., Smith, M.E., and Hancock, J. "Everyone Has to Do It:" A joint action approach to managing social inattention. *Computers in Human Behavior 29*, (2013), 2230–2238.
8. Birnholtz, J., Schultz, J., Lepage, M., and Gutwin, C. A Framework for Supporting Joint Interpersonal Attention in Distributed Groups. *Proceedings of Human-Computer Interaction – INTERACT 2011*. Springer Berlin Heidelberg, 2011, 295–312.
9. Burke, J.L., Prewett, M.S., Gray, A.A., et al. Comparing the Effects of Visual-auditory and Visual-tactile Feedback on User Performance: A Meta-analysis. *Proceedings of ICMI,* ACM Press (2006), 108–117.
10. Chan, A., MacLean, K., and McGrenere, J. Designing haptic icons to support collaborative turn-taking. *International Journal of Human-Computer Studies 66*, (2008), 333–355.
11. Clement, A. Considering privacy in the development of multi-media communications. *Computer Supported Cooperative Work 2*, 1-2 (1994), 67–88.
12. Czerwinski, M., Horvitz, E., and Wilhite, S. A Diary Study of Task Switching and Interruptions. *Proceedings of CHI 2004*, ACM Press (2004), 175–182.
13. Dabbish, L. and Kraut, R.E. Controlling interruptions: awareness displays and social motivation for coordination. *Proceedings CSCW 2004,* ACM Press (2004), 182–191.

14. Dabbish, L. and Kraut, R.E. Awareness Displays and Social Motivation for Coordinating Communication. *Information Systems Research 19*, 2 (2008), 221–238.
15. Fogarty, J., Hudson, S.E., Atkeson, C.G., et al. Predicting Human Interruptibility with Sensors. *ACM Transactions on Computer Human Interaction (ToCHI)*, *12*, (2005), 119–146.
16. Hart, S.G. NASA-task load index (NASA-TLX); 20 years later. *Proceedings of HFES*, Sage Publications (2006), 904–908.
17. Healey, C.G., Booth, K.S., and Enns, J.T. High-speed Visual Estimation Using Preattentive Processing. *ACM Transactions on Computer Human Interaction (ToCHI)*, *3* (1996), 107–135.
18. Hindus, D., Ackerman, M.S., Mainwaring, S., and Starr, B. Thunderwire: A Field Study of an Audio-only Media Space. *Proceedings of CSCW 1996*, ACM Press (1996), 238–247.
19. Hopp, P.J., Smith, C. a. P., Clegg, B.A., and Heggestad, E.D. Interruption management: the use of attention-directing tactile cues. *Human Factors 47*, 1 (2005), 1–11.
20. Inbar, O., Joost, G., Hemmert, F., Porat, T., and Tractinsky, N. Tactful calling: investigating asymmetric social dilemmas in mobile communications. *Behaviour & Information Technology 33*, 12 (2014), 1317–1332.
21. Isaacs, E., Walendowski, A., and Ranganthan, D. Hubbub: A Sound-enhanced Mobile Instant Messenger That Supports Awareness and Opportunistic Interactions. *Proceedings of CHI 2002,* ACM Press (2002), 179–186.
22. Kowalski, R., Loehmann, S., and Hausen, D. Cubble: A Multi-device Hybrid Approach Supporting Communication in Long-distance Relationships. *Proceedings of TEI 2013,* ACM Press (2013), 201–204.
23. Kraut, R.E., Fish, R.S., Root, R.W., and Chalfonte, B.L. Informal communication in organizations: Form, function, and technology. *Human reactions to technology: Claremont symposium on applied social psychology*, (1990), 145–199.
24. Leftheriotis, I. and Chorianopoulos, K. User Experience Quality in Multi-touch Tasks. *Proceedings of EICS 2011*, ACM Press (2011), 277–282.
25. Levesque, V., Oram, L., MacLean, K., et al. Enhancing Physicality in Touch Interaction with Programmable Friction. *Proceedings of CHI 2011*, ACM Press (2011), 2481–2490.
26. Ling, R. *New Tech, New Ties: How Mobile Communication Is Reshaping Social Cohesion.* The MIT Press, Cambridge, Mass, 2008.
27. Mark, G., Gudith, D., and Klocke, U. The Cost of Interrupted Work: More Speed and Stress. *Proceedings of CHI 2008*, ACM Press (2008), 107–110.
28. McFarlane, D. Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction 17*, 1 (2002), 63–139.
29. Mullenbach, J., Shultz, C., Colgate, J.E., and Piper, A.M. Exploring Affective Communication Through Variable-friction Surface Haptics. *Proceedings of CHI 2014,* ACM Press (2014), 3963–3972.
30. Pielot, M. and Oliveira, R. de. Peripheral Vibro-tactile Displays. *Proceedings of Human-Computer Interaction with Mobile Devices and Services*, ACM Press (2013), 1–10.
31. Rabby, M.K. and Walther, J.B. Maintaining on-line relationships. In *Maintaining relationships through communication: Relational, contextual, and cultural variations*. Lawrence Erlbaum Associates, Mahwah, N.J., 2003, 141–162.
32. Rijsbergen, C.J.V. *Information Retrieval.* Butterworth-Heinemann, Newton, MA, USA, 1979.
33. Saket, B., Prasojo, C., Huang, Y., and Zhao, S. Designing an Effective Vibration-based Notification Interface for Mobile Phones. *Proceedings of CSCW 2013,* ACM Press (2013), 149–1504.
34. Stutzman, F. Productivity in the Age of Social Media: Freedom and Anti-Social. In R.T. Scholz, ed., *Learning Through Digital Media: Experiments in Technology and Pedagogy*. 2011.
35. Tang, J.C. Approaching and Leave-taking: Negotiating Contact in Computer-mediated Communication. *ACM Transactions on Computer Human Interaction (ToCHI)*, *14*, (2007).
36. Turkle, S. *Alone Together: Why We Expect More from Technology and Less from Each Other.* Basic Books, 2012.
37. Wickens, C.D. Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science 3*, 2 (2002), 159–177.
38. Winfield, L., Glassmire, J., Colgate, J.E., and Peshkin, M. T-PaD: Tactile Pattern Display through Variable Friction Reduction. *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, (2007), 421–426.
39. Yatani, K., Gergle, D., and Truong, K. Investigating Effects of Visual and Tactile Feedback on Spatial Coordination in Collaborative Handheld Systems. *Proceedings of CSCW 2012,* ACM Press (2012), 661–670.
40. Zickuhr, K. Location-Based Services. *Pew Research Center's Internet & American Life Project*, (2013), 25.