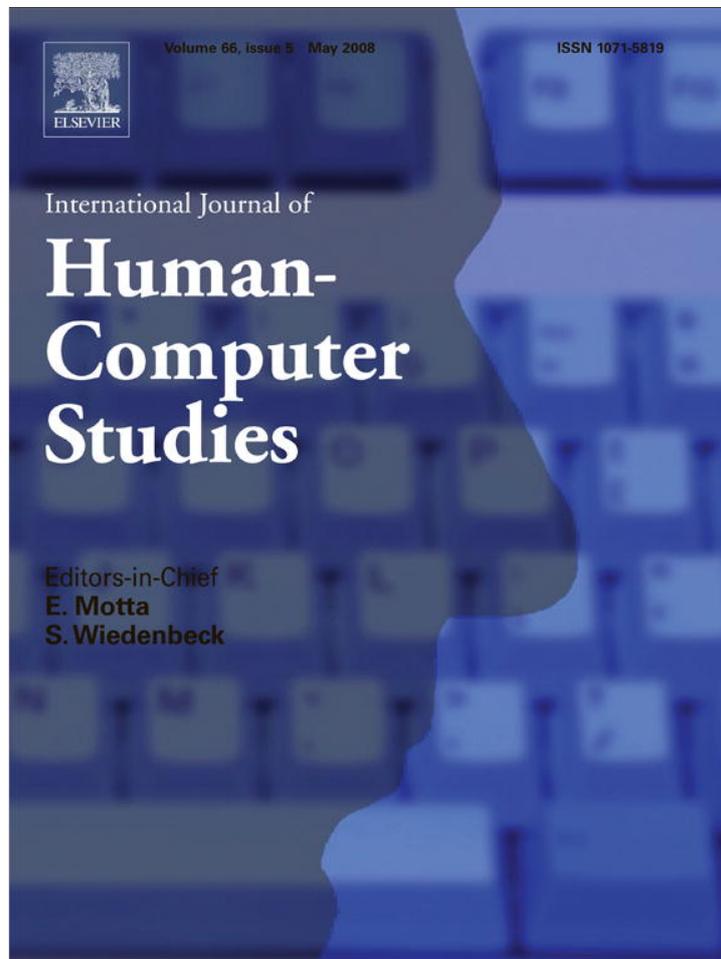


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Designing haptic icons to support collaborative turn-taking

Andrew Chan, Karon MacLean*, Joanna McGrenere

Department of Computer Science, University of British Columbia, 2366 Main Mall, Vancouver, British Columbia, Canada V6T 1Z4

Received 21 February 2006; received in revised form 28 August 2007; accepted 7 November 2007

Communicated by C. Schmandt

Available online 17 November 2007

Abstract

This paper describes research exploring the use of haptics to support users collaborating remotely in a single-user shared application. Mediation of turn-taking during remote collaboration provides a context to explore haptic affordances for background communication as well as control negotiation in remote collaboration: existing turn-taking protocols are rudimentary, lacking many communication cues available in face-to-face collaboration. We therefore designed a custom turn-taking protocol that allows users to express different levels of urgency in their request for control from a collaborator; state of control and requests are communicated by touch, with the intent of offloading visual attention. To support it, we developed a set of haptic icons, tangible stimuli to which specific meanings have been assigned.

Because we required an icon set which could be utilized with specified, varying levels of intrusiveness in real attentionally challenged situations, we used a perceptually guided procedure that consisted of four steps: initial icon set design, perceptual refinement, validation of learnability and effectiveness under workload, and deployment in an application simulation. We found that our haptic icons could be learned to a high degree of accuracy in under 3 min and remained identifiable even under significant cognitive workload. In an exploratory observational study comparing haptic, visual, and combined haptic and visual support for our protocol, participants overall preferred the combined multi-modal support, and in particular preferred the haptic support for control changes and the visual support for displaying state. In their control negotiation, users clearly utilized the option of requesting with graded urgency. The three major contributions in this paper are: (1) the introduction and first case study using a systematic process for refining and evaluating haptic icons for background communication in a primarily visual application; (2) the usability observed for a particular set of icons designed with that process; and (3) the introduction of an urgency-based turn-taking protocol and a comparison of haptic, visual and multi-modal support of our implementation of that protocol.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Haptic icons; Turn-taking; Remote collaboration; Background communication; Signaling; Workload; Design process

1. Introduction

Remote collaboration has become prevalent and is often essential in today's networked society, but current collaboration technology is rarely adequate in its ability to transmit the non-verbal and social cues that underlie smooth communication. At the same time, abstracted haptic signals are emerging as an underutilized medium which may be well-suited to transmitting background information in distributed situations; they are potentially less intrusive than other interaction modalities that by necessity make heavy

use of vision and speech. In this paper, we describe research which explores how haptic feedback could help to fill this communication deficit, using as a testbed the challenge of a remotely collaborating group's need to negotiate control in a single-user application.

Typical technical support for this type of distributed collaboration allows group members to share a view of an application and provides a turn-taking protocol to mediate who is in control of the single cursor. The primary tasks performed tend to be visual in nature; and yet the visual channel is also used to mediate control, meaning that prompts such as requests for control may be distracting if too intrusive, or missed if too subtle or transient. Current systems also lack support for users to express how urgently they wish to gain control, a need that

*Corresponding author. Tel.: +1 604 822 8169; fax: +1 604 822 4231.

E-mail addresses: ca.andrew.chan@gmail.com (A. Chan), maclean@cs.ubc.ca (K. MacLean), joanna@cs.ubc.ca (J. McGrenere).

is usually met in face-to-face collaboration. We hypothesize that haptic feedback could play a role in conveying urgency of a user's desire for control.

Haptic feedback has recently become a common element in consumer electronics, such as cell phone vibrators and gaming consoles that simulate driving over a rough road or a gun's recoil. However, the scope of "everyday" haptic feedback goes far beyond binary buzzes and recreated real-world forces: it can potentially communicate rich and graded information in situations where other senses are unavailable. Early examples include cell phones with distinctive "vibe-tones" to unobtrusively convey caller identity (Immersion Corporation, 2005), and the use of spatially delivered haptic cues to notify a driver when a vehicle is in a blind spot or following too closely, or to indicate directional cues (Burnett and Porter, 2001; Young et al., 2003; Enriquez and MacLean, 2004; Lindeman et al., 2005). Haptic icons, or tactons—haptic stimuli to which specific meanings have been assigned (MacLean and Enriquez, 2003; Brown et al., 2006a,b)—have been considered for use in mobile applications (Brown and Kaaresoja, 2006; Luk et al., 2006) and are central to the research we report here.

To move beyond simple uses of haptic feedback, we are hindered by the fact that the devising of usable and useful haptic representations is not yet supported by the archives of perceptual, practical and process data that underlie visual and auditory information display (Swindells et al., 2005). Haptic messages have been used in just a few applications with mixed success. Contemporary tactile display technology is comparable to a few monochrome graphic pixels. Users lack everyday experience with interpreting abstractions via touch, so their responses are often unpredictable; as are target environments—haptic feedback is generally suggested when other senses are already loaded, implying multitasking and disruptions. Designers of haptic icons clearly need an effective mechanism for predicting and fine-tuning how an icon set will function in a given situation, which is both grounded in basic perceptual data and informed by user reactions specific to the situation. This need became a practical focus of our work.

In this article, we describe our design of a turn-taking protocol by which distributed collaborators can share control of a single cursor; unlike existing approaches, it allows users to express the urgency of their need for involvement and can use either haptic icons or a graphical display to indicate this urgency. To support it, we created and refined a set of seven haptic icons, following an iterative and convergent four-step procedure consisting of initial *ad hoc* contextual icon set creation, perceptual refinement of the set, validation of learnability and effectiveness under workload, and evaluation of the icons in a simulation of the intended application. At each step, the haptic stimuli were adjusted as indicated by intermediate evaluation results; systematic iteration is critical because modification of any set member can affect how other members are perceived or distinguished.

In both abstracted and situated environments, we found that our haptic icons could be learned in under 3 min; they remained identifiable at 95% accuracy even under representative cognitive workload. This learning and identification performance seems viable in the context of the collaborative consumer application we envision. When we tested this protocol *in situ* and compared haptic, visual and combined multi-modal implementations, we found that participants successfully substituted the haptic signals for the more conventional graphical display. The availability of haptic signals influenced their interactions when performing a collaborative task, increasing both control turnover rate and equitability of time-in control. Participants overall preferred the combined multi-modal support, and in particular preferred the haptic support for control changes and the visual support for displaying state; in their control negotiation, users clearly utilized the option to request with graded urgency.

The three major contributions in this paper are: (1) the introduction and first case study using a systematic process for refining and evaluating haptic icons for background communication in a primarily visual application; (2) the usability observed for a particular set of icons designed with that process; and (3) the introduction of an urgency-based turn-taking protocol and a comparison of haptic, visual and multi-modal support of our implementation of that protocol.

1.1. Background and related work

This project draws on two bodies of past work which we briefly survey here: haptics, specifically in terms of haptic perceptual information transfer and the utility of haptic feedback in collaboration; and turn-taking in computer-mediated collaboration.

1.1.1. Haptics in collaboration and communication

There is evidence that the haptic channel may offer a means of non-intrusively informing users of their collaborator's desires, thus alleviating problems with current turn-taking protocols. A comprehensive understanding of our tactile psychophysical capabilities is emerging through the work of researchers such as Klatzky and Lederman, who have documented our exquisite sensitivity to texture felt through a probe (Klatzky et al., 2003). Tan et al. (1999) measured information transfer rates of 2–3 bits/s for vibrotactile stimuli independent of duration: appreciable content can be conveyed through this channel, and clarification of its perceived dimensions and their resolution continues to emerge (e.g. MacLean and Enriquez, 2003; Israr et al., 2006; Ternes, 2007). Attention has focused on the potential of vibrotactile displays for information display (van Erp, 2002; Pasquero, 2006), and there has been some multi-modal work for example in using haptic stimuli to orient attention in another sensory modality; for example, by using taps on the back to direct gaze (Young et al., 2003; Lindeman

et al., 2005). A more complete summary can be found in MacLean (2008).

Collaborating through force communication: Physically rendered virtual environment forces have also been used to aid collaboration, by offering a direct kinesthetic representation of a collaborator's movements. Sallnas et al. (2000) had users jointly arrange cubes in a virtual environment, and Basdogan et al. (2000) asked them to jointly move a ring along a wire without touching the wire. Both found that haptic feedback significantly improved task performance and aided cooperation. Oakley et al. (2001) added physical forces to telepointers in a shared editor so that users could push and pull each other. Force feedback can help under workload: guiding forces have been shown to positively and unintrusively influence behavior during an engrossing task (Steele and Gillespie, 2001; Feygin et al., 2002; Enriquez and MacLean, 2004; Forsyth and MacLean, 2006).

Auditory models for communicating abstractions: Other work has been directed towards building an *abstract* tactile language to support transmission of information to or among users, by associating haptic stimuli with assigned meanings rather than using forces or orienting vibrations to explicitly guide or direct. We precede this discussion by illustrating two possible avenues to designing this sort of abstract information display, by reference to the more mature auditory domain. The *semantic* approach, dubbed "auditory icons" (Gaver, 1993), uses representations of objects or notions that embody a literal, intuitive meaning: the sound of a paper being crushed indicates deleting a computer file. Others have used musical features (melodies, rhythms and crescendos) in a similar way using haptics (van Erp and Spapé, 2003; Brown et al., 2006a,b). However, with this approach there is no systematic basis for determining relative stimulus salience or differentiability, which can lead to problems. For example, the sound of an unimportant event might perceptually dominate the signal for an urgent event. Conversely, others have taken a *symbolic* approach focused on quantifying people's ability to perceptually differentiate "earcons" (Blattner et al., 1989; Brewster et al., 1993): sounds and rhythms with no intrinsic meaning, whose association has to be learned. They found that structuring bursts of sound aided in differentiation, as did varying musical timbre rather than using simple tones.

Haptic icons: In the haptic domain, initial work has focused on the latter, symbolic approach with haptic icons (MacLean and Enriquez, 2003; Tang et al., 2005; Luk et al., 2006) or "tactons" (Brown et al., 2006a,b; Brown and Kaaresoja, 2006). Following graphic icon usage (Horton, 1994; Calpin, 2001; Chen, 2003), these terms have been accepted by the haptics and HCI fields to mean most broadly a haptic stimulus to which a meaning has been associated, such that system display of the stimulus implies the associated abstraction. As with graphical icons, targets might include a linkage to an object or a function, a notification, a status indicator, or an indicator of identity

or content. A haptic stimulus (analogous to a graphic) requires an associated meaning or target to become an icon.

In the present work, we access an important subset of this large representational space: our turn-taking application makes use of icons which fill roles of status monitoring and change notification. Other types of targets for haptic icons have been proposed or implemented in, e.g. Allen et al. (2005), Brown and Kaaresoja (2006) and Luk et al. (2006).

Perceptual validation of icons in an application context has been rare and preliminary: for example, an attempt to use haptic icons to differentiate musical style of songs within a set was promising but indicated a need for further refinement of both stimulus structure and their associated meaning (Allen et al., 2005). Enriquez et al. (2006) demonstrated an encouraging ability of users to learn deliberately arbitrary associations between haptic stimuli and meanings, in a family-organized set of nine icons, but these were not tested in an application. Our group has also employed some steps of the process described in the current paper for a novel tactile display intended for mobile environments (Luk et al., 2006). In contrast to the present work, we focused there on exploring the basic sensations possible with the new display and of the tasks they might be suited for, but performed no *in situ* evaluation of informative haptic signals.

1.1.2. Turn-taking

Since Engelbart and English (1968) proposed the notion of Shared Screen Conferencing, researchers have sought ways to enable real-time collaboration among distributed users. One approach proposed by Lauwers and Lantz (1990) is to design *collaboration-aware* systems that support simultaneous input from multiple users. They also note that in a second approach, a traditional single-user application (which they term *collaboration-transparent*) can be overlaid with a *shared window system*, allowing simultaneous viewing of the application, but only single-user input. We also note a more recently reported third approach to supporting real-time collaboration, which adds technology to an existing single-user application to make it appear to be somewhat collaboration-aware (Xia et al., 2004; Sun et al., 2006).

While many prototypes of collaboration-aware applications have been developed (for example, Ellis et al., 1991; Pedersen et al., 1993), these applications have had little impact outside of research labs. Writing a collaboration-aware application, even with toolkit support (Tse and Greenberg, 2004), is still significantly more difficult than writing a single-user application. Corporations are reluctant to devote the time and money required to redevelop single-user applications to support multiple-user input, and even if they did, the multi-user features in these applications would be used less often than single-user features (Grudin, 1988).

In contrast, shared window systems have enjoyed a degree of commercial success, since they allow most

existing single-user applications to be used in a collaborative context without modification. Although users can simultaneously view the shared application, a turn-taking protocol is required to mediate writing access and control over the cursor. The two most commonly cited supporting protocols are *give* (a user who has requested control must wait for the user in control to voluntarily relinquish it) and *take* (a requestor is immediately given control).

There has been little empirical comparison of turn-taking protocols reported in the literature, and none have to our knowledge drawn conclusive results. We highlight two studies here. Inkpen et al. (1997) studied the effect of turn-taking protocols on pairs of co-located children who were solving Rube Goldberg-like puzzles. While girls did not show a difference, boys shared control more equally when using *take*; the amount of time boys had control was positively correlated with their ability to complete the same task on their own. McKinlay et al. (1993) compared the ability of subjects to reach consensus on a prioritization task using face-to-face communication and three computer-mediated communication protocols: *give*, *take*, and *free-for-all* (where users could work simultaneously). Face-to-face was the clear winner; and the data suggested that the others would be ordered as *give*, *free-for-all*, and *take*.

Commercial systems: Commercial web conferencing systems available on the market include WebEx (Cisco Systems Inc., 2007), Acrobat Connect (Adobe Inc., 2007), and MS Live Meeting (Microsoft Corp., 2005). All of these systems allow a user to share a single-user application with other users while a single user is in control, and all use *give* by default, although some can also be configured to *take* instead. Except for MS Live Meeting, they all include mechanisms for users to explicitly request control from the person in control. The applications use different methods for presenting requests: either “short-lived” dialog boxes, transient tool-tips, or persistent floating windows. Dialog boxes can be disruptive, while tool-tips may be overlooked or (conversely) obscure what the user is looking at. The methods the applications present requests differ in their degrees of persistence: in some the requests are transient, with the result that the request is gone when a user is ready to relinquish control after completing his task. A persistent floating window listing the requester’s name addresses this, but the windows can still be easily missed.

Expression of urgency in collaboration: Conversational analysis reveals that when physically co-located, our body language communicates information to our partners, such

as how urgently we wish to speak (Duncan and Niederehe, 1974); and that non-verbal communication is crucial for smooth and efficient information transfer (Boyle et al., 1994). We expect similar non-verbal communication needs for collaborative work; in particular, its use to express a desire to assume control with varying degrees of urgency. However, none of the turn-taking protocols we have observed provide flexibility in requesting application control, nor do their implementations support user management of the intrusiveness of a signal being sent—features we take for granted in everyday conversation. Visual elements such as dialog boxes and tool-tips are not ideal for this function: by sharing the medium of the primary task, they can grossly interfere when intrusive. Subtle status updates can still impinge undesirably on visual attention, and when they do not they are easily overlooked.

1.2. General approach and process

To support our turn-taking application, we needed a set of haptic icons that would function well in a particular environment. When designing haptic icons, it is desirable to maximize memorable icon set-size while simultaneously controlling at least four other quantities: ease of association of stimulus with target meaning, individual stimulus salience, discernibility of items in the set, and maintenance of desired salience level under cognitive workload. The present research demonstrates one way to manage these factors throughout the design sequence for a specific application, by adjusting stimulus values in a structured, iterative and user-centered process.

In developing our haptically mediated turn-taking application, we followed the four-step process illustrated in Fig. 1. In Step I *Initial Prototyping*, we designed our turn-taking protocol based on a combination of our own observations and the insights provided by past work on turn-taking in collaboration; this protocol included a set of seven signals to be communicated via haptic feedback. At the same time, we prototyped an initial set of stimuli to support the protocol’s haptic signals. In Step II *Perceptual Adjustment*, we used methodology based on multidimensional scaling (MDS) to perceptually refine our stimulus set, ensuring that families of stimuli were distinct from one another and that stimuli *within* a family were related yet distinguishable. Since our protocol will be used while users are engaged on a primary collaborative task, in Step III *Stress Test* we measured the ability of participants to learn the icons (stimulus-meaning

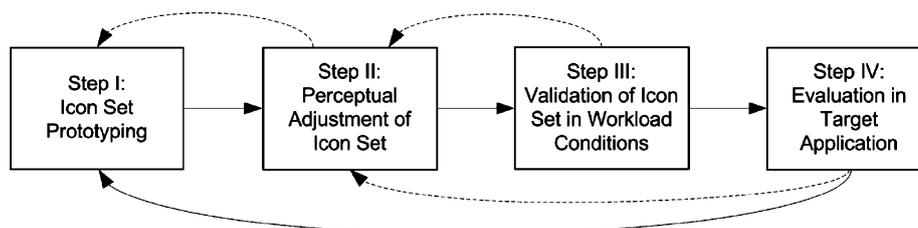


Fig. 1. A process for developing haptic icons. Solid lines are standard progression, and dotted lines indicate iterations that might be needed.

pairs), and to detect and identify them while performing distractor tasks that simulated anticipated application workload. In Step IV *Application Trial* we assessed the value of this perceptually validated set of icons in situ, with an observational study of our target application where we compared performance and subjective factors under different combinations of haptic and visual interfaces.

The method which we used to create our initial stimulus-meaning pairs borrows from both the semantic and symbolic approaches, with the aim of facilitating both learnability and usability. It is based on (a) encoding information in *families* of haptic signals and (b) devising tactile features through semantic reference to common application-related metaphors that we expected users would understand. Other approaches could be taken for initial creation of the stimulus-meaning pairs; while this one seemed well-suited to our purposes, a comprehensive exploration of this design space was not the goal of the present study.

Key hypothesized benefits of our process are (a) evaluative mechanisms at every step which indicate when to iterate on the haptic stimuli, by (b) considering the stimuli as a *set*, in recognition that users will learn and use them relative to one another, more so than for visual or auditory icons. Through examination of results at every step, we ascertain feasibility of the process itself and usability of the result achieved. Our observations build a case for the *need* of a systematic refinement process to ensure expected behavior in the face of complex and unpredictable target environments.

The remainder of this paper describes these four steps (Sections 2–5), with emphasis on Step IV *Application Trial*; the studies from Step II *Perceptual Adjustment* and Step III *Stress Test* are presented more completely in Chan et al. (2005) but summarized here to provide context for the entire process. In Section 6, we reflect on our haptic icon refinement process, usability of the icons produced with them, and insights into the expression of urgency and the impact of mediation modality obtained through our application simulation; and indicate next steps in each

area. Finally, in Section 7, we close by briefly highlighting the main conclusions of our work.

2. Step I: Turn-Taking Protocol and Initial Icon Prototypes

Based on a combination of insights from our own past work (Inkpen et al., 1997) and informal observations of a number of collaborative, floor-sharing example contexts, we felt that expression of urgency was a shortcoming in existing control-sharing protocols for distributed applications, and set about to devise a turn-taking protocol that supported control requests of graded urgency. Because these applications usually involve tasks that heavily utilize both vision and audition, we considered the touch sense as the best candidate for providing contextual information that needed to range between background and foreground in terms of intrusiveness.

2.1. An urgency-based turn-taking protocol

In our protocol, a user is always *in control* of the shared application, *waiting for control*, or simply *observing* the actions of his or her collaborators. A user obtains control by *gently requesting* control, *urgently requesting* control, or by *taking* control. If no one is in control at the time someone tries to obtain control, any of these methods result in the user immediately obtaining control. Otherwise, requests for control are queued, with one queue for gentle requests and one for urgent requests.

As users request control, the user in control is made aware of the requests and their urgency. When the user in control releases it, the first user in the queue of urgent requests is given control; if that queue is empty, the first user in the queue of gentle requests is given control. If both queues are empty when control is released, then no one is in control. This protocol always gives priority to urgent requests for control, but within a queue, temporal ordering of requests is maintained. A state-transition diagram is shown in Fig. 2.

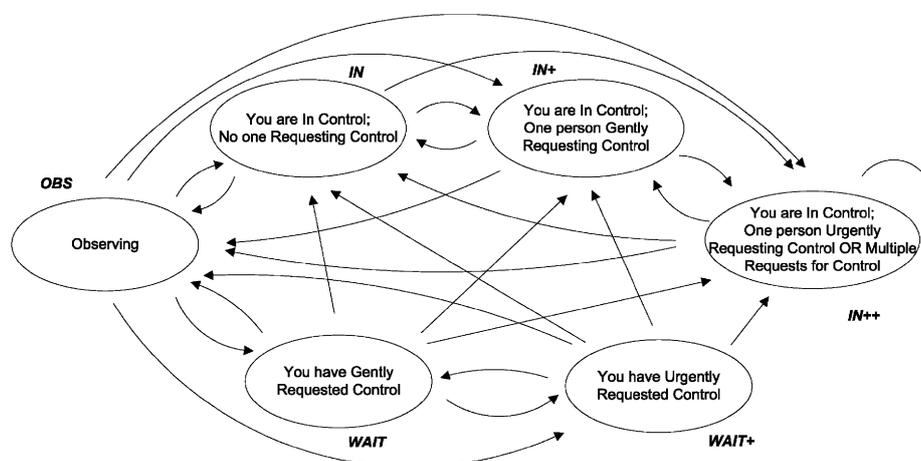


Fig. 2. State-transition diagram for the urgency-based turn-taking protocol. Transitions not possible in our implementation are not shown.

Table 1
Haptic icons used in the urgency-based turn-taking protocol

Family	State	Haptic sensation
Change of control	User has gained control of the shared application	0.4 s, 1000-magnitude, 100 Hz vibration, followed by a 0.1 s delay, followed by a 0.25 s, 8000-magnitude, 100 Hz vibration
	User has lost control of the shared application	0.25 s, 8000-magnitude, 100 Hz vibration, followed by a 0.1 s delay, followed by a 0.4 s, 1000-magnitude, 100 Hz vibration
In control	User is in control of the shared application	1 s, 500-magnitude, 60 Hz vibration; 1 s delay between iterations
	User is in control, but someone has gently requested control	1 s, 5000-magnitude, 60 Hz vibration; 1 s delay between iterations
	User is in control, but someone has urgently requested control	0.7 s, 5000-magnitude, 100 Hz vibration, followed by a 0.1 s pause, followed by a second identical vibration; 0.6 s delay between iterations
Waiting for control	User has gently requested control	Single pulse; 1 s delay between iterations
	User has strongly requested control	Two pulses, separated by a 0.15 s pause; 1 s delay between iterations

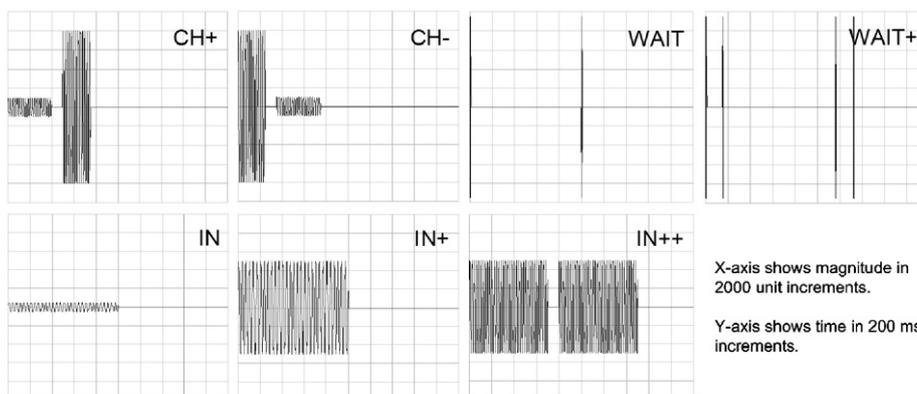


Fig. 3. Graphs of the haptic icons used in the urgency-based turn-taking protocol.

2.2. Families and metaphors

For haptic icons to be used in consumer applications they must be easy to learn. To increase learnable set size through cognitive chunking of subsets (Miller, 1956), we chose to create three families of haptic stimuli to match the logical signal groupings in our turn-taking protocol, rather than a perceptually flat set. Thus, icons in each family shared tactile features and had related meanings. The families, their icons (messages and associated stimuli), and the metaphors that inspired them are described below; Table 1 and Fig. 3 show the detailed configuration of the haptic icons as eventually deployed in Step IV Application trial, in the target turn-taking application.

2.2.1. Changes in Control

A pair of transient sensations, indicating that the receiving user has just gained (CH+) or lost (CH-) control of the application.

Metaphor: The haptic equivalent of the two-tone sound made when Windows XP detects that a USB device has been inserted into or removed from a computer. We

hypothesized that making the icons mirror opposites, like their auditory versions (beep-BEEP, BEEP-beep), would make them intuitive to learn. Both sensations were designed to be intrusive.

2.2.2. In Control

Three periodic, ongoing sensations delivered to the user with current floor control, indicating no outstanding requests (IN), or a collaborator has gently (IN+) or urgently (IN++) requested control.

Metaphor: A heartbeat. A gentle, pleasant, slightly prolonged periodic sensation was used for the no-outstanding-request icon, but progressively more intrusive sensations were used for the gentle and urgent request icons. The icon for an urgent request (IN++) was intended to be clearly noticeable and even slightly unpleasant.

2.2.3. Waiting for Control

Two periodic, ongoing sensations delivered to a user who has made a gentle (WAIT) or urgent (WAIT+) request.

Metaphor: A quicker, delicate “pulse” sensation suggesting a person lightly tapping (WAIT) or moderately tapping

(WAIT+) her fingers while waiting in line. The two sensations were distinguished by a doubling in rhythm and designed to be unintrusive.

Feedback was provided only to a user in control or waiting for control. A user simply observing the collaboration could determine his state by the absence of any sensation, thus both minimizing the chance of annoyance and highlighting more active roles.

2.3. Haptic display

We used Logitech iFeel mice to deliver haptic feedback in our studies. These are standard optical mice with the addition of an embedded vibrotactile display based on an inertial harmonic drive (technology licensed from Immersion Corp.). Although the range of frequencies and amplitudes the iFeel supports are limited, we wanted to find how much we could accomplish using off-the-shelf technology. This consideration also motivated the choice of a mouse rather than introduction of a new hardware element (e.g. a glove). We created haptic sensations using *Immersion Studio* (2004), a GUI-based haptic editor, and integrated the sensations into our test program using the API supplied by Immersion.

The iFeel has a frequency response of 0.01–500 Hz; however, the iFeel motor produced unmaskable confounding auditory noise above 100 Hz, and insufficient magnitude below 20 Hz, so our designs were limited to this range, despite the fact that peak tactile sensitivity occurs around 250 Hz (Shimoga, 1993). Immersion Studio also employs a “magnitude” from 0 to 10,000. It is difficult to interpret this quantity, since the perceived stimulus intensity is a function of both magnitude and frequency. In general, we observed that if two stimuli are the same frequency and above 10 Hz, they must have a magnitude difference of at least 1000 to be perceived as different. Our approach to icon design was thus to set frequency, duration and rhythm based on metaphor, and then adjust amplitude to obtain the desired overall intensity.



Fig. 4. Modified logitech iFeel mouse with thumb buttons.

Table 2
Button presses for acquiring and releasing control

Command	Action
Gently request control	Press front button once
Urgently request control	Press front button twice; if have already gently requested control, press front button once
Take control	Hold front button for 2s and release
Cancel request for control	Press back button
Release control	Press back button

We incorporated the ability to obtain and release control by adding two thumb buttons to the mouse, a feature found on high-end Logitech mice. Fig. 4 shows a picture of the modified iFeel; button presses required for control changes are described in Table 2.

3. Step II: Perceptual Adjustment of Icon Set

Our family-based, metaphor-driven approach yielded a preliminary set of haptic icons to support our turn-taking protocol. However, many subtle parameter variations were possible within the general initial “shapes”—e.g., pulse frequencies, durations and inter-stimulus intervals. We wanted to verify that the In Control and Changes in Control icons were both mutually distinguishable and perceptually clustered by family. A more detailed account of this study can be found in Chan et al. (2005).

3.1. Approach

For perceptual adjustment we used a technique based on MDS that provides an estimate of the perceptual similarity of haptic stimuli by analyzing participants' sorting of test stimuli into groups. This sorting technique, which has efficiency advantages over traditional paired comparison methods, is introduced in detail in MacLean and Enriquez (2003) and further analyzed and used in Pasquero et al. (2006) and Luk et al. (2006). The MDS algorithm then places the stimuli relative to one another in an n -dimensional space, where an appropriate n is revealed by characteristics of the data. Stimuli close to one another in this space are perceptually similar to each other. In addition, we asked participants to rate the intrusiveness and pleasantness of the stimuli on Likert scales.

Our input to the MDS algorithm came from a set of 26 stimuli. We created 24 candidates for the IN icons, varying along the parameters of frequency (20, 60, 100 Hz) and magnitude (500, 2000, 5000, 8000) as well as two levels of a temporal variable introduced by playing each of these 12 stimuli either singly for 1000 ms, or in two 700 ms bursts separated by a 100 ms delay. To these we added the CH+ and CH− icons we had prototyped, to verify their perceptual interactions with the IN icon candidates. We opted

not to test additional candidates of the CH+, CH−, nor the WAIT icons, as we felt they would be clearly distinct from the other stimuli. Since the MDS algorithm finds relative differences between stimuli, the presence of large qualitative family differences would reduce its ability to resolve more subtle differences between members of a family.

In Steps II *Perceptual Adjustment* and III *Stress Test*, participants used Pentium IV 2.67 GHz computers, running Windows XP Professional. To mask audible noise from the vibrotactile display that might influence their perception of the haptic feedback, participants wore Bose QuietComfort two acoustic noise-canceling headphones and listened to white noise.

3.2. Results

Ten participants (six males, four females; aged 21–31) were paid \$10 for a 1 h session. Six participants had little prior exposure to haptic feedback. Unlike MacLean and Enriquez (2003), we found that a 3D MDS analysis substantially improved the goodness of fit over 2D; this was not unexpected, due to the change both in display device and expanded number of design parameters used in creating the icons.

Based on this single-iteration analysis, we selected a set of IN icons with maximal perceptual distinctiveness and intrusive/pleasant ratings that conformed to our heartbeat metaphor: IN was considered quite pleasant and not at all intrusive, IN++ quite noticeable and slightly unpleasant, and IN+ in between. The CH icons appeared to be distinguishable from the IN icons and from each other.

In summary, we used this step of our process to refine the perceptual spacing of a set of seven haptic icons clustered into three metaphor-based families. If the MDS results had not revealed a set of desired perceptual spacings in the first pass, we would have iterated this step in conjunction with guided adjustment of our design parameters as was done in MacLean and Enriquez (2003), i.e. an internal iteration within the second block in Fig. 1.

4. Step III: Stress Test

The purpose of Step III is to ensure that users can learn and recall meanings for a candidate set of haptic icons in an abstraction of a realistic working context; in this case, a simulation of the type of multitasking load we expect users to be under in this sort of collaboration. In our target application, participants would have to identify haptic icons while performing a potentially engrossing collaborative task. Before embarking on the more complex collaborative setting, we first designed a user study with learning and evaluation phases to measure the effect of cognitive workload on the ability of individual participants to identify the seven haptic icons; including the goal of certain signals being detectable under higher workload levels than others. This study is described in more detail in Chan et al. (2005).

4.1. Approach

In the learning phase, participants were instructed to learn the meanings associated with the seven icons as quickly as possible. Since we felt participants would find it difficult to learn the turn-taking protocol meanings without an elaborate explanation, we substituted simpler labels that still reflected the icons' familial relations. Participants were presented with a simple GUI that allowed them to play back the icons as many times as they wanted in any order. To leave the learning phase, they had to pass a test by correctly identifying 19 out of 21 icons (each of the seven icons was presented three times, in random order); they were not penalized for attempting the test multiple times.

In the evaluation phase, participants' ability to recall the meanings they learned in the learning phase was tested under three conditions of successively increasing workload:

Control: Pairs of haptic icons were presented serially. Participants had to indicate when the transition from the first to the second icon in the pair occurred and identify the second icon. They did this by pressing the space bar, selecting the icon's label in a modal dialog box that appeared, and then dismissing the dialog box.

Visual distractor: In addition to the control task, participants solved picture puzzles on-screen by rearranging puzzle pieces on one image so that it matched a second one. When a puzzle was solved, a new puzzle was presented.

Visual+auditory distractors: In addition to the control and visual tasks, participants listened to an audio stream where different colors were spoken at random intervals, and pressed the "b" key whenever they heard the word "blue".

Since our turn-taking protocol is intended for use with a visual task, we did not include an auditory-only distractor condition in this study. The haptic, visual, and auditory tasks in our study were completely unrelated, whereas in our intended application and many others, the visual and auditory channels would be used in concert, with the haptic feedback used to mediate turn-taking. This decision meant that our study tasks were somewhat more difficult than they would be in our target application and thus provided a conservative estimate of performance.

For a given condition, all participants responded to five transitions to each of the seven icons. Transitions were randomly selected from all possible subsets in our turn-taking protocol (Fig. 2), and condition order was counter-balanced across participants. Thus, the evaluation phase used a three conditions × seven icons × five transitions design; all factors were within-subjects.

4.2. Results

Twelve participants (six males, six females; aged 17–28) were paid \$10 for a 1.5 h session, with a promise of a

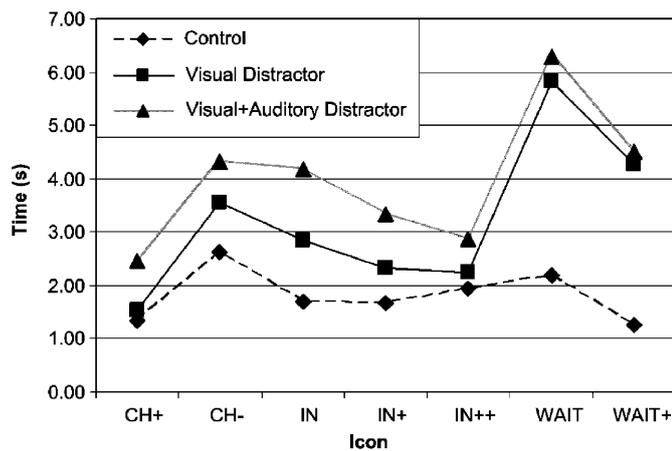


Fig. 5. Mean detection times by icon and condition ($N = 12$). Mean values for each condition are 1.8 s (control), 3.2 s (v distractor), and 4.0 s (v + a distractor).

\$10/individual bonus for the four best “scores.” Instructions explicitly directed participants to pay equal attention to the haptic, visual, and auditory tasks (when present) to maximize their scores. Participants had either no or minimal prior exposure to haptic interfaces (four and eight participants, respectively). After data collection we discovered that one participant had also participated in Step II *Perceptual Adjustment* (3 months previously) and was retained only after verifying similarity of this to other data. All results where $p < .05$ are reported here as significant; for precise statistical tests and significance levels, refer to Chan et al. (2005).

4.2.1. Learning time

We calculated learning time as the period participants spent exploring the icons, excluding the time for the learning test. Participants spent less than 3 min learning the haptic icons (mean 177 s, std. dev. 114 s, range 56–446 s). We believe these results are conservative, since participants were not told the metaphors underlying the icons and doing so would probably have speeded learning. They show that a modest-sized, well-designed set of haptic icons can be learned with relatively low effort—crucial if haptics are to be used in mainstream consumer applications.

4.2.2. Detection time

Calculated from the time the second icon in a pair began playing to when the participant pressed the space bar, mean detection times are shown in Fig. 5¹ by condition and icon. There was a significant main effect of condition, and all post-hoc comparisons were significant: the distractors did slow participants’ noticing of the transition. However, we considered the mean detection time of 4.0 s in the most

difficult (visual + auditory) condition to be acceptable for our intended turn-taking application.²

The results also revealed a significant interaction between condition and icon, indicating that the detection of some icons were more sensitive to workload than others. Specifically, post-hoc comparisons showed no significant difference in detection time as a function of workload condition for IN++: a desirable result, since this icon was intended to be the most intrusive, i.e. difficult to block out. In contrast, even a single distractor task significantly slowed detection for the WAIT family icons, suggesting that they are indeed more amenable to attentional backgrounding than the others, as planned. In general, icons designed to be more intrusive were detected more quickly than icons designed to be unintrusive. The only exception to our expectations was the CH– icon: although intended to be moderately intrusive, its detection times were comparable to the unintrusive icons. CH+, its exact mirror, did not have this problem.

4.2.3. Missed transitions

If the second icon in a pair had played for more than 10 s without detection, it was considered a missed transition. Overall, condition did have a significant impact on the number of missed transitions. Post-hoc comparisons showed that the visual + auditory distractor condition had significantly more missed transitions (18.8%) than the visual distractor (10.5%) and the control (1.7%) conditions.

A significant interaction between condition and icon indicated that transitions to some specific icons were missed more in some conditions than in others. Post-hoc comparisons again showed that transitions to the three icons designed to be most subtle (IN, WAIT and WAIT+) were often overlooked as workload increased.

4.2.4. Identification time

Measured from the appearance of a modal dialog box listing the seven icons to the participant clicking the OK button, identification time data revealed only a marginally significant ($p = 0.061$) main effect of condition (Fig. 6). This result suggests that participants’ ability to switch attentional focus and then identify the haptic icons is less affected by workload than is their ability to detect a change of icon. The average of 3.0 s under the heaviest workload is feasible for our application.³ There was also a significant

²We are unaware of specific benchmarks for detection time in collaborative applications, but we note that detection times in the range of 1–10 s are reported in the interruption/notification literature for studies investigating detection under workload of visual signals that differ by motion, shape, size and color (e.g. Bartram et al., 2003; Gluck et al., 2007).

³Again, we do not have a specific benchmark with which to compare our result. While Bartram et al. (2003) looked at visual signal identification in their interruption/notification work, and found mean identification times of approximately 1.5 s, identification in their study differed somewhat from that measured in our study. They measured the users’ ability to identify which icon, among a large set, had changed, rather than to identify how one icon changed; and workload conditions were different.

¹These figures differ slightly from those reported in Chan et al. (2005), as we discovered a minor error in our original statistical analysis. However, the corrected analysis still permits us to draw the same conclusions.

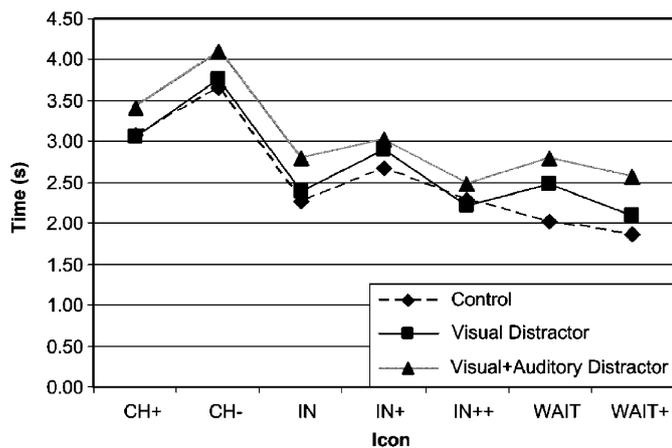


Fig. 6. Mean identification times by icon and condition ($N = 12$). Mean values for each condition are 2.6 s (control), 2.7 s (v distractor), and 3.0 s (v + a distractor).

main effect of icon, and post-hoc comparisons revealed that identification of the CH family icons, and in particular of CH–, took significantly longer than for the others. Indeed, CH– was mistaken for CH+ or IN+ four times more often than those icons were mistaken for CH–. Most importantly, however, the most deliberately intrusive icon (IN++) was the least affected by workload in terms of identification time.

4.2.5. Correct identification

On average, participants identified icons correctly in 95.7% of trials with no significant impact of condition, suggesting that participants did learn icon associations during the brief training period and maintained recall throughout the session.

4.2.6. Distractor task performance

On average, participants placed one puzzle piece every 4 s in conditions including the visual distractor, independent of condition; for the auditory distractor, participants also identified 91.1% of the keywords on average. Together, these results suggest that participants were attending to the distractor tasks.

4.2.7. Discussion of auditory effects

We were surprised at the degree to which auditory information influenced participants' performance. While piloting our Step III Stress Test study, participants reported using auditory noise from the iFeel to distinguish between the CH family icons, despite wearing noise-canceling headphones and listening to white noise. Since we wanted to measure the ability of participants to identify the icons using their haptic sense, we modified the icons after our piloting to reduce the auditory noise they generated. Unfortunately, the study's detection time and identification results (reported above) revealed that the modifications made the icons less distinguishable from one another. Had we repeated Step II *Perceptual Adjustment*

with the modified icons before Step III, it is likely that this oversight would have been detected. To compensate, we made minor changes to the CH icons again *after* Step III, and informally confirmed the changes with a few test participants.

4.3. Summary

The goals of Step III *Stress Test* were to determine how long it would take participants to learn the haptic icons, how quickly they could detect and identify them in isolation, and how this ability would be affected by the addition of workload. Participants on average were able to learn the icons using a substitute set of labels in less than 3 min, suggesting that learning the icons would be feasible in the target application. In the absence of other tasks, they detected and identified changes in the icons at 1.8 and 2.5 s, respectively. With the addition of visual and auditory distractor tasks, detection times increased significantly to 4.0 s, while identification times increased marginally to 3.0 s. Participants maintained 95.7% accuracy in haptic icon identification regardless of the workload condition in effect at time of detection. We are unaware of specific benchmarks against which to compare these results. However, intuitively they seemed reasonable and were sufficiently encouraging to move forward with testing the icons in our intended turn-taking application.

5. Step IV: Application Trial

Having created and validated a set of haptic icons to support our turn-taking protocol, we deployed them in a realistic simulation of the target context and conducted an exploratory observational user study. Below, we report a mix of qualitative and quantitative findings which will be elaborated and generalized in Section 6.

5.1. Approach

Groups of four users completed a sequence of collaborative furniture-layout tasks (using Microsoft Visio, a diagramming tool) under three conditions that differed in the amount of haptic and visual turn-taking information provided. The goal was to compare our protocol mediated by the traditional visual modality (costly in attention and screen space) with the potentially less disruptive haptic channel. To this end, we addressed the following research questions:

1. Can participants utilize the haptic icons in the context of our turn-taking protocol, after a modest learning time?
2. How will the different workload conditions impact collaborative style, equitability of control sharing and task performance?
3. Which modality (visual, haptic, or combined) will participants prefer for information about and control over turn-taking?

We could have made many other comparisons, for example contrasting our protocol with the more common give and take, and with purely verbal mediation of turn-taking (no protocol or technological support provided). These would have required an even more elaborate and lengthy study. We felt it was important to first investigate the urgency-based protocol on its own and secondly compare its haptic and visual instantiations, before making comparisons to other protocols. Without this first pass aimed at optimizing basic usability, it would be difficult to ascertain whether technology can add value here.

To implement our protocol, we modified an open-source software application called virtual network computing (VNC) (RealVNC Ltd., 2006). VNC consists of client and server applications; using the client, a user can control the desktop of a remote computer running the server. If multiple users are viewing the same desktop, the server by default simply handles keyboard and mouse inputs in the order in which they arrive, often causing unpredictable results. No support is provided to help users decide who is in control. We modified both the client and server to use our turn-taking protocol.

5.2. Conditions

All three conditions in our study used the same urgency protocol. To minimize confounds, the same information about who is in control or requesting control at a given urgency is available in each condition:

Visual (V): A User Window (implemented by modifying VNC) displays who is in control, who has gently and urgently requested control, and a list of the group members. This window is visible at all times, unlike the transient tool-tips or dialog boxes used in current commercial solutions, in order to provide information persistence comparable to H below. A custom Button Bar similarly allows users to request and release control. Both objects, shown in Fig. 7, persist beside the shared application window.

Haptic (H): Haptic inputs and feedback are provided as described in Section 2.3 and Tables 1 and 2. The User Window from the V condition can also be displayed by pressing the space bar, but in this condition the window has to be dismissed before any other actions can be taken; i.e., it does not persist as in the V condition. This allows the same information to be available as for V, but requires the haptic modality to be the first source of information; and additionally allows monitoring of demand for the additional information.

Haptic+Visual (H+V): This condition combines haptic input and feedback from the H condition with the persistent User Window and Button Bar from the V condition. Participants can use either the Button Bar or haptic inputs to request and release control.



Fig. 7. User window and button bar.

These conditions are not purely “Visual” or “Haptic,” since both kinesthetic activities and some visuals are necessarily involved in all cases. Rather, the condition names distinguish the primary source of turn-taking information and control. Nor was modality the only difference: because of sensory differences, it was not possible to assure that the two conditions were perfectly balanced in terms of perceptual salience, nor to match their interaction models (H uses symbolic icons, while V offers a literal, always-available and more informative view). Rather, each condition plays to the strength of its respective sense, with intrusiveness and information content made as similar as possible.

5.3. Study setup

The study was conducted in a graduate student laboratory. To simulate a distributed setting, participants were seated at workstations such that they could not easily see each other. Participants wore Sennheiser HD280 professional headphones and Sony ECM-T115 lapel microphones routed through a series of audio mixers for communication with one another. While the headphones did not have noise-canceling technology, their snug, closed-ear design made their ability to muffle external noise comparable to the Bose headphones used in the Step II *Perceptual Adjustment* and III *Stress Test* studies.

When working together, participants in each group used VNC clients running on Pentium III–IV computers (733 MHz–2 GHz, and 256–512 MB of RAM). The VNC

client is not computationally intensive, so task performance was similar across computers. Each computer had a 17" LCD display with 1280 × 1024 screen resolution. The clients connected via a 100 Mbps LAN to the VNC server, hosted on a Pentium IV 2.67 GHz, 512 MB computer. All computers ran Windows XP.

We automated data collection in several ways. The VNC client and server were instrumented to record all activity information in log files. A Linux open-source screen recording program called *vncrec* (Hayashi, 2002) captured groups' efforts at creating task solutions by accessing the VNC server. Another Linux open-source program called *transcode* (Bitterberg, 2004) was used to create MPEG4 movies of the data. Group conversations were recorded on audio tape.

5.4. Task

Our task was designed to closely approximate real-world group collaboration, and several task characteristics were deemed important. Groups should share a common body of knowledge, but each individual should possess specific, specialized knowledge. Also, while groups should work towards a well-defined set of goals, there should be constraints on how the goals can be achieved. Finally, group members should have conflicting interests, but collaboration should not be adversarial.

To satisfy these characteristics, we created three isomorphic tasks for our three conditions, all involving furniture layout in a communal work area using Microsoft Visio. To deepen the collaborative aspect, the three tasks shared eight firm constraints that *had* to be observed as specific task goals were met, and another eight soft constraints that *should* be observed; perfect solutions were impossible. For example, a goal in one task was to add 5–10 workstations to an existing room. One of the firm constraints (for all tasks) was to maintain three-foot-wide walkways to each piece of furniture in the room; a soft constraint was that noisy areas should be distant from workstations.

Groups were given 20 min to formulate a solution for each task. All members knew the complete set of goals they were to achieve, which were provided in written form; each member was personally responsible for two hard and two soft constraints. To maximize participation, participants were responsible for different sets of constraints for each of the three tasks. The tasks were designed such that creating a near-optimal solution would be very difficult in the time given, but that a sufficing solution would be possible.

5.5. Procedure

Participants individually completed a training phase before working together on the furniture-layout tasks. First, participants learned how to use Microsoft Visio. We simplified its interface to make it easier to learn, hiding all functions not needed in the study. Participants were given a

brief demonstration that showed how to add, move, rotate, and remove objects and groups of objects and then used Visio to complete a brief set of exercises, ensuring that they understood its use at this level. Next, participants were trained to identify the haptic signals we used, first reading descriptions of the stimuli and their protocol-based meanings. For this the GUI training application from Step III *Stress Test* was reused, with hints added at the bottom of the screen to help participants learn the associations. When participants felt they had learned the signals and their meanings, they proceeded to the learning test, which was identical to the test used in Step III. Participants again had to correctly identify 19 out of 21 icons to pass.

The group then completed the three conditions. Each condition was preceded by a 5-min warm-up period where the group completed a scripted set of actions to familiarize themselves with the user interface for that condition. Groups then spent 20 min working on a furniture-layout task. After each condition, participants individually completed a questionnaire and were required to take 5-min rest breaks. At the end of the study, participants individually completed an overall questionnaire. They were then interviewed and debriefed as a group.

The study required one 3-h session, for which participants were each paid \$25. As an incentive to concentrate on the tasks, groups were told that their task solutions would be evaluated, and the top 1/4 of the groups would each receive a \$40 bonus.

5.6. Design and participants

A within-subjects design was used so that participants could compare the modalities in the three conditions. We adopted a "2 × 2 + 1" design: the V and H conditions were counterbalanced, as were two out of the three layout tasks, but the remaining task was always paired with the H + V condition, and this was always the last condition presented. By placing this condition last, we were able to record and ask which modality participants relied on, having had equal exposure to the other two approaches, and at the same time reduce condition explosion.

To simulate real-world group composition and better understand the impact of the conditions on groups with different backgrounds, we recruited participants according to several constraints. Each group had to have at least one male and one female, and all members had to be acquainted with one another. Additionally, we required diversity in degree of technical background across groups. Participants were not screened with respect to Visio experience because our interface permitted only novice behavior. We rejected any individual familiar with haptic stimuli similar to those used in our study.

We limited experiment size to four groups, and thus our results are based on four groups × four users × three conditions × 20 min = 16 h of user data. While a larger sample would have provided greater statistical power, the cost of running each group was very high in terms of

recruitment effort and analysis time as well as participant compensation and equipment rentals. This was an initial exploration of a very new concept, and it seemed likely that the results would indicate adjustments prior to a more comprehensive commitment. Therefore, rather than striving for a homogenous sample we emphasized group diversity in technical background and gender, in order to maximize insight into how our protocol will fare in a real-world working environment. Given the small number of groups, we expected to rely exclusively on descriptive statistics. However, for completeness we did conduct analysis of variance (ANOVA) on our dependent measures where appropriate, and report the statistically significant results.

5.7. Dependent measures

We measured learning effort through the amount of time spent exploring the haptic stimuli and the number of attempts required to pass the learning test. We measured aspects of collaboration through the time spent in control before releasing or losing it, the time spent waiting for control after submitting a gentle or urgent request, and frequency data such as the number of requests for control. To gauge task performance, we evaluated the task solutions according to how well they satisfied the specified goals while respecting constraints.

We also collected qualitative subjective data relating to the dependent measures from questionnaires and post-study interview data. The questionnaires consisted of Likert-scale and open-ended questions, and requested participants to rank the conditions in order of preference.

5.8. Results

5.8.1. Group composition

Four groups of four participants participated in the study (eight male and eight female participants). Participants were university students aged 18–41, had normal tactile sensitivity, used the mouse with the right hand, and had not participated in our earlier studies. They exhibited a variety of haptics exposure: six reported none, six used game controllers with vibrotactile displays, three had used other haptic devices, and one did not respond to the question. Each group is described below:

- *The Engineers* (three males, one female) had taken over a year of undergraduate engineering courses together, participated jointly in extracurricular activities, and kept in touch even though one participant had changed faculties.
- *The Long-Time Friends* (two males, two females) were each majoring in a different area. They had known each other since secondary school and attended religious services, played sports, and took courses together.
- *The Teachers* (one male, three females) were completing their Education degrees, as part of a cohort of

approximately 40 people who took all their classes together for 1 year.

- *The Graduate Students* (two males, two females) consisted of two male–female pairs from different research labs in computer science and electrical engineering. The two pairs did not know each other; we permitted this due to recruitment difficulty.

5.8.2. Learning and using haptic stimuli

The learning component of the study was designed to ensure that participants could achieve a threshold level of performance within a reasonable amount of time, rather than (as in Step III *Stress Test*) test how quickly participants could learn the haptic signals. Therefore, participants were encouraged to learn “carefully” rather than “quickly.”

Learning time was calculated as the time spent exploring the haptic stimuli in the GUI training application. As in Step III *Stress Test*, the learning test itself was not included. Learning times ranged from 51 to 270 s (mean 135 s, std. dev. 64 s); the range, mean, and standard deviation all show an improvement over the Step III results, which suggests that the association process is facilitated by the provision of compatible icon meanings. Indeed, we observed that three of the five participants with the longest learning times required only one attempt to pass the learning test; this fact and our conservative task instructions suggest that learning times could have been even lower.

5.8.3. Collaborative style

We were interested in the impact of condition on collaborative indicators such as the distribution of the different methods for acquiring control, the frequency of control transfer, and the influence of those methods on wait times.

Methods of acquiring control: Fig. 8 shows the distribution of methods for acquiring control used for each condition. The figure shows a clear preference for *gentle requests* over *takes* and *urgent requests* in all conditions: *gentle requests* accounted for 60–73% of all acquisition requests made when control was held by another. A two-factor repeated-measures ANOVA (three acquisition methods \times three conditions) confirmed a significant effect of method ($F_{2,6} = 17.401$, $p = 0.003$, partial $\eta^2 = 0.853$)⁴ with Bonferroni corrected post-hoc pair-wise comparisons suggesting that *gentle requests* dominated *urgent requests* ($p = 0.054$) and *takes* ($p = 0.093$).

Besides *gentle requests*, *urgent requests*, and *takes*, participants could also directly *obtain* control. The *obtain* category includes any control acquisition when no one was already in control. We cannot determine from the data when participants knew someone was in control or not; for

⁴As appropriate, all analyses of variance of performance data in this experiment were conducted at the group level, treating each group as one unit of analysis.

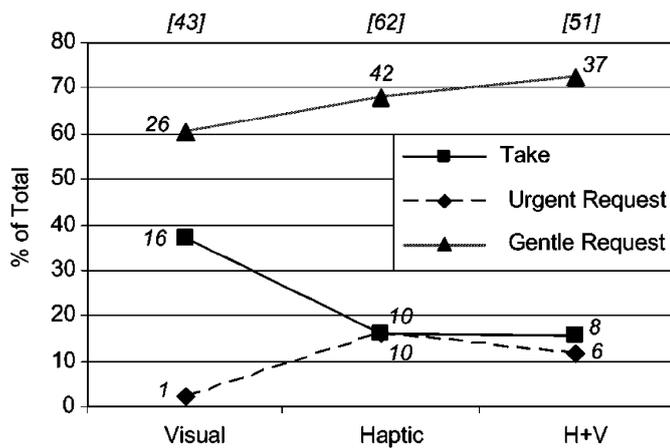


Fig. 8. Distribution of non-verbal methods for requesting control, shown by % of total with absolute values noted ($N = 16$). These do not include directly Obtained Controls, which are 12, 39, and 46 for V, H, H+V, respectively.

example, a *take* when one believes someone else is in control may represent an aggressive collaborative style, whereas in the latter situation it is simply one of three methods of assuming control when no one else has it. As a result, Fig. 8 lists values for *obtains* in its caption but does not plot them.

We expected that most of the time a participant would retain control until someone else requested it. However, the graduate students adopted a practice in their second condition whereby they each released control as soon as they had finished their task; this accounts for the large increase in *obtain* counts in the H and H+V conditions. Further study is needed to see whether this approach would be widely adopted in extended use.

The distribution of *takes* and *urgent requests* in Fig. 8 was also unexpected. In the V condition, *takes* dominated *urgent requests* (37% compared to 2%), whereas in the H and H+V conditions they were used more equally (16%/16% and 16%/12%, respectively). Despite its lack of statistical significance we did explore possible reasons for this intriguing behavior. For example, an *urgent request* might not have been salient enough in the V condition, where the requester's name moved into the box labeled Urgent Request. If that were the case, however, one would still have expected groups to attempt to use the *urgent request* before resorting to *take* or to verbal methods of acquiring control. However, only one *urgent request* was made across all groups in the condition.

In summary, the data show evidence of our protocol's key feature, the graded request: across all conditions participants made use of *gentle requests*, but also resorted to strong methods at a lower but still important rate (67.3%, 10.8% and 21.8%, respectively, for *gentle*, *urgent* and *take*). The different conditions did appear to affect the distribution of the stronger requests made, but more research will be required to clarify our results on that point.

Table 3

Number of changes in control by condition ($N = 16$)

	Visual	Haptic	Haptic + Visual
Engineers	13	22 ^a	17
Long-Time Friends	20	18 ^a	25
Teachers	8 ^a	20	18
Graduate Students	9 ^a	21	29
Average	12.50	20.25	22.25
Standard deviation	5.45	1.71	5.74

^aCondition seen first.

Frequency and timing of turnover: Across all groups and conditions, the number of control changes ranged from 8 to 29. As shown in Table 3, there were nearly twice as many changes in the H and H+V conditions compared to the V condition, but little additional change between H and H+V; a repeated-measures ANOVA approached significance ($F_{2,6} = 4.552$, $p = 0.063$, partial $\eta^2 = 0.603$). This suggests that haptics may facilitate more frequent turnover of control.

Increased turnover could indicate that participants found the haptic feedback irritating rather than informative. To examine whether this was the case, we analyzed the responses from a questionnaire conducted after the (final) H+V condition, where participants rated several statements on a Likert scale. Fourteen out of 16 participants disagreed or strongly disagreed with the statement, "The haptic feedback was too strong." Only two out of 16 agreed that "The haptic feedback was distracting" and none strongly agreed. In addition, 11 out of 16 agreed or strongly agreed with the statement, "I easily recognized what each haptic signal meant" with 4/5 of the remaining participants rating the statement as neutral. These findings, combined with the fact that none of our users mentioned in the post-study interview having experienced irritation when receiving a request for control, suggest that the haptic feedback was an effective method of facilitating control turnover rather than an irritant. However, we did not ask participants directly about their level of irritation on the questionnaires.

Considered by modality (Tables 4 and 5), *gentle requests* made in the H and H+V conditions were responded to more quickly than in V: a repeated-measures ANOVA showed a significant effect of condition ($F_{2,6} = 6.747$, $p = 0.029$, partial $\eta^2 = 0.692$) but Bonferroni corrected post-hoc pair-wise comparisons were not statistically significant. Only one *urgent request* was made in the V condition, making a statistical comparison across conditions impossible. Averaged across all conditions, control-holders were more responsive to urgent requests, releasing control in an average of 19.4s as opposed to 29.3s for gentle requests. While encouraging, this comparison was not statistically significant.

In summary, conditions including haptic signaling seemed to speed control turnover across request methods, which could indicate a generally more dynamic (quicker

Table 4
Gentle requestor's perspective: average time (s) from a *gentle request* until gaining control, by group and condition ($N = 16$)

	Visual	Haptic	Haptic + Visual
Engineers	111.4	39.5	8.6
Long-Time Friends	46.8	23.8	5.9
Teachers	24.3	28.0	6.4
Graduate Students	50.0	4.3	3.0
Column sum/number of groups	58.1	23.9	6.0
Standard deviation	37.3	14.6	2.3

Unweighted table average is 29.3 s.

Table 5
Urgent requestor's perspective: average time (s) from an *urgent request* until gaining control, by group and condition ($N = 16$)

	Visual	Haptic	Haptic + Visual
Engineers	–	7.2	22.0
Long-Time Friends	–	1.0	3.5
Teachers	67.0	33.0	2.0
Graduate Students	–	–	–
Column sum/number of groups	67.0	13.7	9.2
Standard deviation	–	17.0	11.1

“–” denotes no urgent requests made.

Unweighted table average is 19.4 s.

turnover) and equitable (uniform time-sharing) style than with V alone.

5.8.4. Equitability of sharing control

Although control turnover was higher in the presence of haptics, we further wanted to know whether haptics promoted *equitability* of control amongst team members. By examining how much time participants spent both in control and waiting for control, we can see that the collaboration dynamics changed across the conditions.

The average amount of time a participant spent in control of Visio in any one turn, before releasing or losing control, was noticeably larger in V than the other conditions (Table 6). A comparison of group means yielded a significant main effect of condition ($F_{2,6} = 5.849$, $p = 0.039$, partial $\eta^2 = 0.661$) but Bonferroni corrected post-hoc pair-wise comparisons were not significant.

Examination of the overall percentage of time each participant was in control under each condition revealed an interesting finding. For each group, a spread was calculated by subtracting the lowest time-in-control percentage (of the participants in the group) from the greatest (Table 7). We observed that the spreads were larger in V as compared to H + V and especially the H condition. Statistical analysis with a repeated-measures ANOVA yielded a main effect of condition ($F_{2,6} = 37.405$, $p < 0.001$, partial $\eta^2 = 0.926$), with post-hoc pair-wise comparisons showing that the spread was significantly larger in V than in H ($p = 0.011$) and H + V ($p = 0.014$).

Table 6
Control holder's perspective: average time (s) in control before releasing control, by group and condition ($N = 16$)

	Visual	Haptic	Haptic + Visual
Engineers	96.3	50.1	69.1
Long-Time Friends	57.8	63.9	49.6
Teachers	153.5	55.6	70.5
Graduate Students	109.8	33.4	35.8
Column sum/number of groups	104.3	50.8	56.3
Standard deviation	39.5	12.9	16.6

Unweighted table average is 70.5 s.

Table 7
Equitability of control time: spread of percentage of time in control, between group members most and least in control (% , $N = 16$, or 4 per group)

	Visual	Haptic	Haptic + Visual
Engineers	0.44	0.23	0.35
Long-Time Friends	0.39	0.25	0.24
Teachers	0.39	0.13	0.22
Graduate Students	0.47	0.25	0.35
Column sum/number of groups	0.42	0.22	0.29
Standard deviation	0.04	0.06	0.07

Smaller spreads indicate more even sharing.

Unweighted table average is 0.31 s.

Table 8
Condition preference: number of subjects who ranked a given condition first/second/third ($N = 16$)

	Visual	Haptic	Haptic + Visual
Obtaining control	1/7/8	4/5/7	11/4/1
Conveying state	4/6/6	1/5/10	11/5/0
Overall	2/8/6	2/4/10	12/4/0

Together these results suggest that control is shared more evenly in the presence of the haptic version of our turn-taking protocol, for this task.

5.8.5. Participant preferences

Condition rankings: At the experiment's end, participants ranked the three conditions in order of preference for acquiring control, for displaying the turn-taking state, and for overall preference. Table 8 shows that the H + V condition was strongly preferred overall (12 participants) as compared to the two single-modality conditions, which each received an equal number of overall first-order rankings (two participants each). While H + V was thus most favored for both acquiring control and displaying state (11 participants), the H condition was the next preferred for acquiring control (and indeed, was most used); whereas for displaying state, the V condition was next preferred. It thus appears that the H and V conditions may have differentially supported these two subtasks.

Table 9
Questionnaire results: Likert response results are shown averaged by condition

No.	V	H	H + V	Question text, as seen by participant
1.	4.0	3.9	4.2	I obtained control in a reasonable amount of time
2.	2.4	2.7	2.1	I didn't remain in control long enough
3.	2.8	2.7	2.2	Sharing control was frustrating
4.	4.0	4.1	4.3	My group shared control fluidly
5.	4.1	4.0	4.2	I was able to express my opinion
6.	4.3	4.3	4.0	My group members listened to my opinion
7.	4.0	4.1	4.4	When I moved from waiting for control to being in control, I noticed it quickly
8.	3.3	3.8	3.9	When someone gently requested control from me, I noticed it quickly
9.	3.3	4.1	4.1	When someone urgently requested control from me, I noticed it quickly
10.	4.1	3.5	4.0	When I was in control and lost control, I noticed it quickly
11.	3.3	2.3	3.3	I constantly monitored [opened] the User Window when someone asked me for control
12.	2.9	2.6	2.8	I constantly monitored [opened] the User Window when I was waiting for control
13.	2.8	2.5	2.8	I constantly monitored [opened] the User Window when I was neither in control nor waiting for control
14.		2.3	1.9	The haptic feedback was too strong
15.		2.6	2.6	The haptic feedback was too subtle
16.		3.4	3.9	The haptic feedback felt pleasant
17.		2.8	2.4	The haptic feedback was distracting
18.	3.6	4.0	4.1	My group successfully addressed the demands and constraints of the task
19.	3.1	2.9	2.9	If multiple people could access Visio at the same time, our solution would have been better
20.	3.8	3.7	3.9	If we were working face to face, our solution would have been better

After completing the indicated condition, participants were asked to rate the associated question on a scale of 1–5, where 1 = strongly disagree; 5 = strongly agree ($N = 16$).

The order of the H, V conditions was counterbalanced, while H + V was always administered last.

Cells are blank where the question was not applicable to that condition.

[] Indicate alternate wording used in different conditions.

When asked to justify their overall rankings, some participants who ranked H + V highest commented that the haptic feedback *notified* them of changes in state, while the User Window indicated *who* was in control or requesting control; due to this, for example, an inactive user (neither in control or waiting) cannot confirm haptically that *no one* is in control. One participant noted the lack of haptic feedback when no one was in control, and sometimes opened the User Window to confirm that fact. Another reported that difficulty arose occasionally in the H condition when the person gaining control was not who the person releasing control expected it would be, based on the audio dialog.

Self-reported use of user window: The questions asked after each condition and their results are shown in Table 9. Despite commenting in the interview that they desired the User Window to be visible, participants did not report using it heavily in the H condition, where it was not continually visible (Questions 11–13); and in fact, the number of times the User Window was *actually* opened in the H condition was on average fewer than three times/20 min per participant. Questionnaire results also suggest that participants monitored the User Window less in H when busy (in control with someone waiting) than in other states when they were more likely to be “idling”.

In summary, participants preferred H + V overall, and appeared to want access to some of the detailed information available only visually in the current implementation; and the pattern of H and V condition preference in

Table 10
Normalized task scores ($N = 16$)

	Visual	Haptic	Haptic + Visual
Engineers	0.69	0.28 ^a	0.84
Long-Time Friends	0.60	0.44 ^a	0.65
Teachers	0.22 ^a	0.70	0.73
Graduate Students	0.41 ^a	0.79	0.83
Average	0.48	0.55	0.76
Standard deviation	0.21	0.23	0.09

^aCondition seen first.

conjunction with self-reported and actual usage patterns suggest that that these conditions might differentially support monitoring and control acquisition functionalities.

5.8.6. Task performance

For completeness, we checked task performance across conditions, even though large effects were not expected out of 20 min segments. Task solutions were evaluated for how well they satisfied the task goals, while following the specified hard and soft constraints. Points were awarded for satisfying goals and deducted for failing constraints. We estimated maximum possible scores with previously generated reference solutions.

Each group's task solutions were scored (one per condition), and the resulting scores were divided by the reference solution scores (Table 10). Not surprisingly, a

learning effect was evident between the first and second conditions completed.

We also checked to ensure that there was no effect of task in the first two conditions, which were counter-balanced for task as well as condition order: the mean scores on those two tasks were 0.50 and 0.54, respectively. Thus, apart from learning effects, the different modalities did not appear to impact task performance in this study.

5.9. Summary

Our four-group study showed that modality impacted group behavior. In particular, haptics increased control turnover rate and resulted in more uniform time-in-control across group membership, potentially indicating a more dynamic collaborative style. Participants preferred combined modalities for mediation of turn-taking overall, and secondarily, H for acquiring control and V for displaying state. Participants learned the haptic icons in well under 3 min, and condition (modality) did not impact task performance.

6. Discussion and future work

In this section, we discuss the usability of the haptic icons we developed, our reflections on our haptic icon refinement process, and the turn-taking protocol we designed. We also comment on directions for ongoing and future work related to each of these three areas.

6.1. Usability of haptic icons

6.1.1. Learnability and performance

In our studies, participants were able to learn the set of seven haptic icons with a speed and accuracy which we found surprising. In the absence of hints to guide learning, participants in Step III *Stress Test* learned the icons in an average of 177 s; in Step IV *Application Trial*, with hints but also encouragement to learn carefully rather than quickly, participants only required an average of 135 s, a 23% decrease. While we are not aware of suitable benchmarks to provide a specification, intuitively this learning time compares positively with what users are often asked to commit to learning specific aspects of a new application or device. Further, in our own informal lab experience, learning time tends to improve with longer term exposure to haptic stimuli; so our results may be a lower bound.

In situ, task performance in Step IV *Application Trial* was not measurably affected by the modality used to mediate turn-taking; and subjectively, users in most ways accepted the icons and were able to utilize them for communication. These together support an overall conclusion that participants seemed able to easily adapt to using the haptic icons. This implies that the icons could be applied in a real-world setting with similar results, especially given that participants in Step IV had diverse

technical backgrounds and academic foci. A broader implication of these results is that deploying a well-designed set of haptic icons in consumer applications should be feasible.

6.1.2. Multitasking

Participants showed an impressive ability to detect and identify changes in the haptic icons under workload. While performing the visual and auditory distractor tasks in Step III *Stress Test*, participants still only required an average of 4.3 s to notice a change in the haptic icon being presented. They also maintained over 95% accuracy in identifying haptic icons regardless of workload. As with the learning time, it is reasonable to anticipate that with longer exposure, identification times and accuracy would also further improve. More generally, adequacy of identification timing will be related to specific task bandwidth, for example, the duration of non-interruptible task units and time-criticality of interactions: 4 s is probably fine for Visio, but not for an action video game. A means of modeling permissible latency using this approach would be a useful future contribution.

Icon identity was collected through a modal dialog box; therefore, the dependence of identification time on workload represented context switching rather than simultaneous multitasking. In some contexts, this may be a realistic approximation: by noticing the change, users have made this the primary focus of attention and “backgrounded” what they were previously occupied with. A different question which will be harder to test is whether with more use, users will be able to carry out and appropriately respond to stimuli without engaging attentive processes at all.

6.1.3. Subjective responses to our haptic feedback

When used in their application context, were our haptic signals useful, irritating, or distracting to recipients? Did sending a haptic message feel appropriately professional or perhaps too intimate? For a shy user, will pressing a button which s/he knows will gently ‘buzz’ the current floor-holder seem more or less intrusive and distasteful than interrupting the whole group with a soft-spoken verbal cue? Some of these possibilities did not emerge until analysis time, so we are left to build our inferences from the questions we asked as well as reference to other work.

Irritation and Intrusion: In their questionnaire responses (Table 9), users tended to agree that haptic feedback was pleasant but too subtle, and to disagree that it was too strong or distracting. This undermines the idea that the higher turnover rate when requests were made haptically derived from irritation: rather, it appears that users wished the feedback was *more* intrusive. This still leaves unclear whether “slightly unpleasant” was appropriate for intentionally intrusive stimuli: it is possible that senders and/or recipients may have perceived them as rude. A possible benefit of the training we used is that through repeated exposure, users could learn how to moderate their own use of the haptic stimuli.

Intimacy: Other works have demonstrated conveying intimate physical contact through a remote haptic link (Brave and Dahley, 1997; Yohanan et al., 2005, Smith and MacLean, 2007); but they have done this using continuous interactive forces designed to promote a sense of connection, rather than our asynchronous buzzes. Smith and MacLean (2007) further notes that an intimate mediating metaphor (holding hands) leads to a more intimate-feeling connection. The impersonal metaphors used here, combined with participants' rating of the haptic feedback as "pleasant," suggests that these interactions were not perceived as intimate. But this deserves a look in the future, with attention to individual reactions which might vary, for example, by personality, gender distribution and power level.

Clarity: Some participants indicated a preference for greater signal clarity and strength. This presents a challenge, since simply turning up the amplitude tends to jeopardize esthetics and be disruptive. The answer might lie in future haptic display technologies designed to be expressive without the "buzz", for example (Kaaresoja et al., 2006; Pasquero et al., 2007) and, more generally, devices capable of a wider range of sensations which can increase distinctiveness across a set. Looking forward, we anticipate the creation of novel display hardware that will eventually support the variety of sensations our haptic sense is able to capture from the real world. In parallel with such developments, it may also prove important to explicitly consider subjective factors such as esthetics in a step of the evaluation process with another or modified step in Fig. 1.

6.2. Reflections on our icon refinement process

Overall, our four-step process for icon refinement and evaluation under successively more specific and demanding experiment scenarios, summarized in Fig. 1, provided the type and level of guiding and/or confirmatory feedback we sought from it. While it is not possible to consider a quantitative measure of this process' success in the absence of a benchmark for comparison, our case study experience of its strengths and weaknesses leads to a number of concrete observations on its impact and value, as well as its generalizability and ways to improve it.

6.2.1. Evidence of process value

Using the current case study, we can evaluate process value from the perspectives of: (1) the quality of result the process delivered (i.e. final icon set usability) and (2) an examination of the impact of our navigation of the 'roadmap' provided by the process relative to a path not informed by the checks it lays out.

(1) As described in Section 6.1, we evolved a set of icons that were apparently usable (learnable, recognizable, detectable and utilized) in a realistic application context, with an amount of development and evaluative

effort which seemed reasonable given the novelty of both the media and the application. While not a validation of the process, it is encouraging. Moreover, the process provided interim evidence that our usability concerns were being addressed, and confirmation in the form of Step IV's *Application Trial* good usability results.

(2) The process provided guidance and showed convergence and thus it had an impact. The process supplied iteration loops with evaluative feedback points (Fig. 1, dashed arrows), and the evaluation results and any adjustments they suggested were generally consistent with later findings. At a number of points, evaluations indicated that stimuli needed to be changed; following these directives corresponded to an absence of later problems, whereas when we skipped a final evaluative step, a problem resulted which would certainly have been detected and repair of which unnecessarily delayed our progress. Generally, if in Step II *Perceptual Adjustment* one cannot get a good icon set, one may have to revert to Step I *Initial Prototyping* and choose different metaphors. The line from Step III *Stress Test* to Step II is for re-checking the perceptual differences in the icon set following an adjustment (this is the one we missed when we changed the *Change in Control* stimuli to make them less audible, see Section 4.2). Finally, there are lines from Step IV *Application Trial* to Steps II and I, for larger iterations when the icons are not as successful in the real world as hoped, but in this case study our icons seem to have converged on good values by Step IV.

What we *cannot* say, given this experience, is that this is the only or best possible means of obtaining these benefits (it is just the first); or that in terms of refinement, it can take an initial stimulus set to a 'global maximum' of performance. Its forte is clearly in local optimization. A designer still needs to *start* from roughly the right place, ever more so as vocabulary size increases, which is not easy.

In summary, we have evidence that this systematic icon refinement process contributes value. Our experience in developing the haptic icons for this application has been crucial for planned additional iterations on the process development, seeking ways to make it more efficient and effective.

6.2.2. Generalizability of process

To what extent can the steps we used in the present case study be generalized to other icon design challenges? We comment in turn on each step.

Step I (Initial Prototyping): We employed one possible vehicle for choosing the initial stimulus-meaning pairings for our turn-taking protocol, namely a metaphor-based, family-oriented approach. To formalize this step, there is a critical need for the haptics and HCI communities to jointly map the large space of haptic stimulus and icon

design, and support it with heuristics, such as those gleanable from (van Erp, 2002; Swindells et al., 2005; Brown et al., 2006a, b; Luk et al., 2006), as well as prototyping tools (e.g., Enriquez and MacLean, 2003; Swindells et al., 2006) and perceptual foundations (in the manner of Tan et al., 1999, 2003; Enriquez and MacLean, 2004, 2007; Pasquero, 2006; Dixon et al., 2007; Enriquez et al., 2007).

While our initial prototyping approach appears to have been successful in this case study, one specific avenue of future investigation will be to compare it to symbolic (as opposed to semantic) and/or perceptually flat (as opposed to clustered) set construction, examining how effectively each approach generalizes to larger sets and the learnability of the icons developed using each approach; as in Enriquez and MacLean, 2008.

Step II (Perceptual Adjustment): We employed a well-validated and efficient mechanism to ascertain how users will perceive the icon candidates as a set. This early step is application-independent and thus generalizable. A recently developed augmentation supports this technique's scaling to handle much larger icon set sizes (80–100), further broadening its applicability (Ternes, 2007).

Future work will involve investigating statistical measures that could be used in conjunction with MDS to determine whether two stimuli are indeed perceptually different and thus provide a clear indication of maximized set size and adequate perceptual spacing of families and family members. Such a metric will further enhance the generalizability of this step.

Step III (Stress Test): Multitasking, high workload environments are likely to be typical of the applications where haptic icons will provide value. While a workload evaluation will generally need to be customized to capture key aspects of a specific environment or task, we expect that an environmental context will often share important features with the one modeled in this case study, and thus this mechanism for “stress testing” can serve as a useful starting template. We note that one likely improvement to the generalizability of Step III will be to control workload at a finer resolution, for example by investigating more than the two levels of workload which we did.

Step IV (Application Trial): A realistic simulation of the sort executed here seems the best compromise between the options of an even more realistic field test, and a simpler, more abstracted and controlled laboratory setup. In a realistic simulation we encounter many of the unexpected subtleties of the actual environment; the effort involved in such a simulation, though considerable, is less than for a field trial, and offers more opportunities for manipulating variables of interest in a relatively controlled and observable environment. In the other direction, our level of simulation trades the solid conclusions available from simpler laboratory setups with their greater control and larger datasets, for more ecologically valid as well as nuanced observations.

While the setup will need to be customized to every new application context, as for Step III *Stress Test*, the general

structure can follow a common model. One ‘win’ in the overall process is that, as shown here, streamlined versions of the earlier steps can be efficiently used for training Step IV participants in the new tools prior to engaging in the real task.

6.2.3. Streamlining steps in process

As mentioned previously, when minor changes are made to an icon set anywhere downstream, an additional round of perceptual testing (Steps II *Perceptual Adjustment* and III *Stress Test*) may be needed to identify unintended set-wide consequences. It would therefore be useful to devise a means of combining Steps II *Perceptual Adjustment* and III *Stress Test* in order to reduce the II↔III iterations required. One approach could be to simulate varying workload during the Step II MDS task by limiting the time participants are given to categorize the stimuli, or engaging them with a distractor task, thus requiring them to make more instinctive judgments. This would be more conservative than the current process, wherein Step III places more emphasis on transitions that are important for a specific application.

Elsewhere, we note that Step IV *Application Trial* as executed here could jeopardize the goal of “reasonable effort.” Simulating and studying an application context to mine it for procedural, social and modality insights for a new medium is labor-intensive. However, in production use with the whole process validated and the basic premise of haptic icons more accepted and understood (where we hope we will soon be), this step might not be necessary, or it might be done differently. For example, if one needs simply to test signals in a real-world design scenario without much control, and access to just a few accepted observational or quantitative data channels, the context would not need to be simulated; a designer might proceed directly to a Step V field trial.

6.3. Multimodal support for an urgency-based turn-taking protocol

6.3.1. The utility of expressing urgency

Usage of graded-urgency requests: Our primary quantitative findings regarding turn-taking were that firstly, as predicted, users did choose *gentle requests* most often to request control, and secondly, that they also used stronger assertive methods when they felt it necessary: *urgent requests* and *takes* accounted for approximately one-third of all requests for control. This degree of strong requests is even more marked given that users made urgent requests even when these did not consistently speed up response (more below). Graded request urgency was clearly utilized.

An interesting further study would be to see whether one could manipulate the 1:2 ratio of strong to gentle requests found here by making the overall task more or less urgent; for example, for time-critical, higher stakes or higher entertainment value tasks we would expect a more balanced ratio. Labeling presents another possible variable.

If “urgent” had been described as “ask again” would participants have used it more broadly, would they have perceived it as nagging, or would they find their own usage patterns independently of initial coaching?

Response time: Self-report responses, consistent across conditions but small and statistically insignificant, showed a 0.2 pt increase in Likert scale agreement that control-holders *noticed* an urgent request quickly relative to a gentle request (Table 9). More data would be useful here, including a longitudinal view of responsiveness with increased protocol familiarity. We also conjecture that users might perceive they have successfully expressed urgency based on how it *feels* to ask as well as the consequence, akin to the way we are alleged to strike computer keys harder when writing emphatic words. The physicality of our request modality might have supported this.

Support of collaboration: In self-reports averaged across conditions, users reported a low level of frustration in sharing control, and agreed that their group shared control fluidly. On the whole, they felt they got control fast enough, disagreed that their time-in control was too short, and believed they quickly noticed when advancing off the queue. These observations and their degree, together with the absence of contradicting answers to the questions we asked, contribute to a sense that users were overall positive about the level of collaboration they were able to achieve in this task but that improvements are certainly possible.

In summary, our results altogether do support investing in a more extensive study of this approach to turn-taking. Extrapolating to expected behavior in real scenarios where users are vested in results, we hypothesize that the availability of choices other than silence or a strong interruption might allow higher-stakes collaboration to be intense without becoming adversarial, showing more marked impact of an urgency-based protocol. The next logical steps will be to compare our protocol both to other turn-taking protocols in more intense tasks, and to face-to-face collaboration, with the goal of the latter being the establishment of comparable benchmarks for performance in managing floor control in non-colocated situations *without* technological mediation.

6.3.2. Impact of mediation modality

Equality of conditions: The H and V conditions differed by primary modality (visuomotor for V, purely motor for H) and design; it was impossible to perfectly match their perceptual loudness, information content, and interaction techniques. To make our comparisons fair, we strove to make each condition as optimal as possible rather than the same. For example, we could have made the User Window blink or change color upon key request changes, but felt the increased distraction would be negative overall; the iFeel did vary haptic salience to express urgency level, because otherwise this difference could not be conveyed at all. The H condition involved an iconic notification model and minimal information, whereas in V, information was

provided literally, in detail and at all times. Further work will need to examine the impact of these factors; for example, in light of our observations, we speculate that a visual display showing urgent requests with more salience might have been acceptable, and reduced the need for monitoring.

H improves turnover and equity: In Step IV *Application Trial*, haptic mediation increased overall control turnover and seemed to promote equitability in control time across users; we view both as positive effects in the context of this “democratic” task. We hypothesize that users must have found the H method of control exchange either less cumbersome or in some sense more informative: since no condition *compelled* users to relinquish control faster or more often, an increase in either overall turnover frequency or use of a specific method was presumably by preference.

Preference for full information: Users wanted the User Window, with its details of the current request queue, visible but stated they used it less in haptic conditions; and in fact accessed it infrequently in H. Thus the added effort to open the User Window in the H condition was at least perceived as a usability impediment. This leaves open the question of whether in the long term, easy availability is positive (does no harm and contributes subjective benefits like confidence or general awareness), negative (a distraction and waste of screen space), or merely a “crutch” (it will simply slow down the rate at which users learn a novel interface and the haptic modality).

Similarly, users did report opening the User Window least in H relative to other conditions when in control with a queue (Table 9, Questions 11–13); i.e., when they were most busy, as opposed to idling. This might indicate either that the additional information was not really needed, or that the deficit was inconvenient but manageable. A longer, harder task (to reveal performance differences) might resolve both questions.

Synergistic functions: When asked which modality they preferred overall, participants consistently chose H+V; and secondarily they noted H was best for obtaining control, and V for conveying state. This result, together with User Window preferences, supports the notion that these two subtasks are different and that our H and V conditions are good for different things. As noted above, V affords *detailed* state monitoring, but has the potential downside of distraction during an engrossing task, particularly when in control. We had theorized that H would allow users to monitor the general queue state without visual competition, and that this would often be enough. Learning and familiarity might be at issue, as well as how the desire for detail relates to role or place in queue. Here, a longitudinal study would be informative; or conversely, exploring the impact of adding participant identification information to the haptic stimuli with more expressive display technology.

Control acquisition: What makes H better for control acquisition? In addition to modality, the interaction differed somewhat: the iFeel used two unlabeled custom

buttons to achieve five functions, most of which were accessed frequently. V's cursor-controlled Button Bar had a unique labeled button for each function, which on the surface should have made it more functionally clear. The most likely explanations are: (a) the H button interaction was faster (no mouse movement), and/or more natural (button functions were easy to use), and (b) H's manual nature and its specific button arrangement co-located with haptic state display eased users' visual loading. Both possibilities need to be taken forward.

In summary, we have observed concrete support for our theory that haptic mediation can enhance usability both perceptually and socially, because its use of a different sensory modality than the primary task makes it less disruptive. The details of how users used the different modalities and for what purpose give useful insights into future design iterations, as well as raising questions for further study. Participants preferred H + V overall, desiring access to the detailed information available only visually in the current implementation and wanting the more direct action request model offered by H. Individuals were able to consciously access haptic signal content under simulated task workload despite the short learning time, and evidence for other sensory modalities lead us to speculate that haptically presented content may also be acted upon without the user's awareness (Merikle et al., 2001; Rensink, 2004). Socially, we have identified issues of perceived intimacy and etiquette which need to be kept in mind when designing interactions in which unassertive users can be more comfortable in requesting control, and domineering users more aware of others' needs.

A logical next step is a longitudinal study, to neutralize learning effects and allow a more demanding and prolonged task with measurable outcomes for performance, fatigue and group dynamics.

7. Conclusions

Our goal in this research was to explore whether haptics could help to support non-verbal communication among distributed collaborators. We found a testbed in our development process in the problem of negotiating control of a shared single-user application, where we posited that using non-verbal means to express urgency of request would positively impact collaboration transparency and dynamics.

The haptic implementation of our turn-taking protocol benefited from a systematic user-based refinement process, which emphasized *perceptual adjustment* of the emergent prototype haptic icon set under conditions that moved progressively from abstract to realistic. We anticipate that this process, and others which will follow it with similar aim and hopefully broader capability, will have utility beyond its use here. It will be some time before haptic display technology, user familiarity and scientific knowledge of haptic perception will match graphic and auditory sophistication, and until then, haptic communication

systems must be created in a design cycle which centers tightly around user and context. The cycle described here needs most immediately to be expanded to encompass the initial icon set design, with heuristics and other tools to aid designers in starting at the right place.

The set of haptic icons developed using this process was usable, learned within 3 min and identified under workload with 95% accuracy and a worst-case mean of 4.0 s for intentionally subtle signals. In application simulation, users utilized and responded to the icons appropriately, and task performance was unaffected in this relatively brief task. We would like to see whether continued exposure in a longer and more intense task would reveal further impact, implying a difference in collaboration effectiveness among conditions. In parallel work, we are developing the means of creating much larger icon sets and testing longitudinally whether they can be fully used. The question of whether this modality can, over time, come to feel like a "natural" communication channel for abstract information is central to the success of this type of application.

Users utilized both haptically mediated turn-taking and graded-urgency requests for control, encouraging further investigation of this approach. Comparison of primarily visual and haptic implementations of our turn-taking protocol suggests that people used our conditions in different ways, obtaining detailed information visually but preferring the haptic condition for action. Haptic mediation appears to influence collaborative style, with increased control turnover and more equitable time-in-control. Users chose to make strong requests half as often as gentle ones, a ratio which seems natural and which might (we speculate) be fed by the user's desire to *express* urgency as much as by the experienced response from collaborators. Our simulation did not generate the high intensity which often characterizes real communication-critical situations, and we look forward to seeing how such increased realism would impact the effects we saw here.

The potential utility of lightweight haptic signals and our urgency-based protocol extends to other categories of distributed collaboration and shared control over resources or a group's attention, including co-located settings where it can be awkward to get the attention of a dominating participant. Possible applications include meeting floor control, team document editing and design reviews. In more specialized examples, air traffic controllers need to smoothly transfer responsibility and information as airplanes cross jurisdictions, as do emergency personnel when coordinating a response to a massive forest fire. The use of interpersonal haptic signaling might non-intrusively facilitate other kinds of group communication: for example, students may learn better if they can confidentially ask instructors to slow down. Finally, this deployment of haptic icons can also serve as an indicator of their potential use in non-collaborative situations such as computer-user communication, where issues like graded levels of interruption and assignment of meaning to background signals also apply.

References

- Adobe Inc., 2007. Adobe acrobat connect (formerly Macromedia Breeze). Accessed August, 2007. <<http://www.adobe.com/products/acrobatconnectpro/>>.
- Allen, M., Gluck, J., MacLean, K.E., Tang, E., 2005. An initial usability assessment for symbolic haptic rendering of music parameters. In: Proceedings of the Seventh International Conference on Multimodal Interfaces (ICMI '05), Trento, Italy, pp. 244–251.
- Bartram, L., Ware, C., Calvert, T., 2003. Moticons: detection, distraction and task. *International Journal of Human-Computer Studies* 58 (5), 515–545.
- Basdogan, C., Ho, C.H., Srinivasan, M., Slater, M., 2000. An experimental study on the role of touch in shared virtual environments. *ACM Transactions on Computer-Human Interaction (TOCHI) (Special Issue on Human-Computer Interaction and Collaborative Virtual Environments)* 7 (4), 443–460.
- Bitterberg, T., 2004. *transcode*. <<http://directory.fsf.org/transcode.html>> (accessed August, 2007).
- Blattner, M.M., Sumikawa, D.A., Greenberg, R.M., 1989. Earcons and icons: their structure and common design principles. *Human-Computer Interaction* 4 (1), 11–44.
- Boyle, E., Anderson, A., Newlands, A., 1994. The effects of visibility on dialogue and performance in a cooperative problem solving task. *Language and Speech* 37 (1), 1–20.
- Brave, S., Dahley, A., 1997. inTouch: a medium for haptic interpersonal communication. In: *The Conference on Human Factors in Computing Systems (CHI'97)*. ACM Press, New York, pp. 363–364.
- Brewster, S.A., Wright, P.C., Edwards, A.D.N., 1993. An evaluation of earcons for use in auditory human-computer interfaces. In: *ACM Conference on Human Factors in Computing Systems (CHI'93)*, Amsterdam. ACM Press, New York, pp. 222–227.
- Brown, L.M., Kaaresoja, T., 2006. Feel who's talking: using tactons for mobile phone alerts. In: *CHI '06 Extended Abstracts on Human Factors in Computing Systems*, Montréal, Canada. ACM Press, New York, pp. 604–609.
- Brown, L.M., Brewster, S.A., Purchase, H.C., 2006a. Multidimensional tactons for non-visual information presentation in mobile devices. In: *Proceedings of the Eighth Conference on Human-Computer Interaction with Mobile Devices and Services*, Helsinki, Finland.
- Brown, L.M., Brewster, S.A., Purchase, H.C., 2006b. Tactile crescendos and sforzandos: applying musical techniques to tactile icon design. In: *CHI '06 Extended Abstracts on Human Factors in Computing Systems*, Montréal, Canada. ACM Press, New York, pp. 610–615.
- Burnett, G., Porter, M., 2001. Ubiquitous computing within cars: designing controls for non-visual use. *International Journal of Human-Computer Studies* 55, 521–531.
- Calpin, S., 2001. *Icon Design: Graphic Icons in Computer Interface Design*. Watson-Guptill Publications.
- Chan, A., MacLean, K.E., McGrenere, J., 2005. Learning and identifying haptic icons under workload. In: *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC 2005, IEEE-VR2005)*, Pisa, Italy, pp. 432–439.
- Chen, P.P., 2003. Toward a methodology of graphical icon design. In: *IEEE Fifth International Symposium on Multimedia Software Engineering (ISMSE '03)*. IEEE Computer Society Press, Silver Spring, MD, pp. 120–121.
- Cisco Systems Inc., 2007. WebEx. <<http://www.webex.com/>> (accessed August, 2007).
- Dixon, S.R., Wickens, C.D., McCarley, J.S., 2007. On the independence of compliance and reliance: are automation false alarms worse than misses? *Human Factors* 49 (4), 564–572.
- Duncan, S., Niederehe, G., 1974. On signaling that it's your turn to speak. *Journal of Experimental Social Psychology* 10, 234–247.
- Ellis, C., Gibbs, S., Rein, G., 1991. Groupware: some issues and experiences. *Communications of the ACM* 34 (1), 39–58.
- Engelbart, D., English, W., 1968. A research center for augmenting human intellect. In: *AFIPS Conference Proceedings of the 1968 Fall Joint Computer Conference*, San Francisco, CA, pp. 395–410.
- Enriquez, M., MacLean, K., 2003. The Hapticon Editor: a tool in support of haptic Communication Research. In: *Proceedings of the 11th Annual Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, IEEE-VR2003*, Los Angeles, CA, pp. 356–362.
- Enriquez, M.J., MacLean, K.E., 2004. Impact of haptic warning signal reliability in a time-and-safety-critical task. In: *Proceedings of the 12th Annual Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, IEEE-VR2004*, Chicago, pp. 407–415.
- Enriquez, M.J., MacLean, K.E., 2007. Backward and common-onset masking of vibrotactile stimuli. *Robotics and Neuroscience in Brain Research Bulletin (special issue)*, in press.
- Enriquez, M., MacLean, K.E., 2008. The role of choice in longitudinal recall of meaningful tactile signals. In: *Proceedings of the 16th Annual IEEE Symposium on Haptic Interfaces (HAPTICS-08)*, June.
- Enriquez, M., MacLean, K.E., Chita, C., 2006. Haptic phonemes: basic building blocks of haptic communication. In: *Proceedings of the Eighth International Conference on Multimodal Interfaces (ICMI 2006)*, Banff, Canada. ACM Press, New York.
- Enriquez, M., MacLean, K., Neilson, H., 2007. Interplay of tactile and visual guidance cues under multimodal workload, UBC Department of Computer Science.
- Feygin, D., Keehner, M., Tendick, F., 2002. Haptic guidance: experimental evaluation of a haptic training method for a perceptual motor skill. In: *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Los Alamitos, CA. IEEE Computer Society, Silver Spring, MD.
- Forsyth, B., MacLean, K.E., 2006. Predictive haptic guidance: intelligent user assistance for the control of dynamic tasks. *IEEE Transactions on Visualization and Computer Graphics* 12 (1), 103–113.
- Gaver, W., 1993. Synthesizing auditory icons. In: *Conference on Human Factors in Computing Systems (CHI'93)*, Amsterdam. ACM Press, New York, pp. 228–235.
- Gluck, J., Bunt, A., McGrenere, J., 2007. Matching attentional draw with utility in interruption. In: *the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, San Jose, CA, USA. ACM Press, New York, pp. 41–50.
- Grudin, J., 1988. Why CSCW applications fail: problems in the design and evaluation of organization of organizational interfaces. In: *Proceedings of the 1988 ACM Conference on Computer-Supported Cooperative Work*, Portland, OR, USA. ACM Press, New York, pp. 85–93.
- Hayashi, Y., 2002. *vncrec*. <<http://www.sodan.org/~penny/vncrec/>> (accessed August, 2007).
- Horton, W., 1994. *The Icon Book-visual Symbols for Computer Systems and Documentation*. Wiley, New York, NY.
- Immersion Corporation, 2005. Vibetonz system. <<http://www.immersion.com/mobility>> (accessed August, 2007).
- Immersion.com., 2004. *Immersion Studio for Gaming*. Immersion Corporation.
- Inkpen, K., McGrenere, J., Booth, K., Klawe, M., 1997. Turn-taking protocols for mouse-driven collaborative environments. In: *Graphics Interface '97*, Kelowna, Canada, pp. 138–145.
- Israr, A., Tan, H.Z., Reed, C.M., 2006. Frequency and amplitude discrimination along the kinesthetic-cutaneous continuum in the presence of masking stimuli. *Journal of the Acoustical Society of America* 120 (5), 2789–2800.
- Kaaresoja, T., Brown, L.M., Linjama, J., 2006. Snap-Crackle-Pop: tactile feedback for mobile touch screens. In: *Eurohaptics*.
- Klatzky, R., Lederman, S., Hamilton, C., Grindley, M., Swendsen, R., 2003. Feeling textures through a probe: effects of probe and surface geometry and exploratory factors. *Perception and Psychophysics* 65 (4), 613–631.
- Lauwers, J.C., Lantz, K.A., 1990. Collaboration awareness in support of collaboration transparency: requirements for the next generation of

- shared window systems. In: SIGCHI Conference on Human Factors in Computing Systems, pp. 303–311.
- Lindeman, R.W., Sibert, A.J.L., Mendez-Mendez, A.E., Patil, A.S., Phifer, A.D., 2005. Effectiveness of directional vibrotactile cuing on a building-clearing task. In: SIGCHI Conference on Human Factors in Computing Systems (CHI '05), Portland, OR, USA. ACM Press, New York, pp. 271–280.
- Luk, J., Pasquero, J., Little, S., MacLean, K., Levesque, V., Hayward, V., 2006. A role for haptics in mobile interaction: initial design using a handheld tactile display prototype. In: ACM Conference on Human Factors in Computing Systems (CHI '06), Montreal, Canada. ACM, New York, pp. 171–180.
- MacLean, K., 2008. Haptics in the wild: Interaction design for everyday interfaces. In: Carswell, M. (Ed.), *Reviews of Human Factors and Ergonomics*, HFES, Santa Monica, CA, to appear.
- MacLean, K., Enriquez, M., 2003. Perceptual design of haptic icons. In: *Proceedings of the Eurohaptics*, Dublin, Ireland.
- McKinlay, A., Arnott, J., Procter, R., Masting, O., Woodburn, R., 1993. A study of turn-taking in a computer-supported group task. In: *People and Computers VIII, Proceedings of the HCI '93 Conference*, pp. 383–394.
- Merikle, P., Smilek, D., Eastwood, J.D., 2001. Perception without awareness: perspectives from cognitive psychology. *Cognition* 79, 115–134.
- Microsoft Corp., 2005. Microsoft Office Live Meeting. <<http://office.microsoft.com/en-us/livemeeting/>> (accessed August, 2007).
- Miller, G., 1956. The magical number seven, plus or minus two: some limits on our capability for processing information. *Psychological Review* 63, 81–97.
- Oakley, I., Brewster, S., Gray, P., 2001. Can you feel the force? An investigation of haptic collaboration in shared, editors. In: *Proceedings of the Eurohaptics*, Birmingham, UK.
- Pasquero, J., 2006. Survey on Communication Through Touch. Center for Intelligent Machines, McGill University, Montreal, Canada.
- Pasquero, J., Luk, J., Little, S., Maclean, K., 2006. Perceptual analysis of haptic icons: an investigation into the validity of cluster sorted MDS. In: *14th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems (IEEE-VR'06)*, Alexandria, VA. IEEE, New York.
- Pasquero, J., Luk, J., Levesque, V., Wang, Q., Hayward, V., MacLean, K.E., 2007. Haptically enabled handheld information display with distributed tactile transducer. *IEEE Transactions on Multimedia* 9 (4), 746–753.
- Pedersen, E., McCall, K., Moran, T., Halasz, F., 1993. Tivoli: an electronic whiteboard for informal workgroup meetings. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Amsterdam, The Netherlands. ACM Press, New York, pp. 391–398.
- RealVNC Ltd., 2006. Virtual Network Computing. <<http://www.realvnc.com/>> (accessed August, 2007).
- Rensink, R., 2004. Visual sensing without seeing. *Psychological Science* 15 (1), 27–32.
- Sallnas, E., Rasmus-Grohn, K., Sjostrom, C., 2000. Supporting presence in collaborative environments by haptic force feedback. *Transactions on Computer-Human Interaction (TOCHI)* 7 (4), 461–476.
- Shimoga, K.B., 1993. A survey of perceptual feedback issues in dexterous telemanipulation. II. Finger touch feedback. In: *Virtual Reality Annual International Symposium*. IEEE, New York, pp. 271–279.
- Smith, J., MacLean, K.E., 2007. Communicating emotion through a haptic link: design space and methodology. *International Journal of Human Computer Studies (IJHCS)*, Special issue on Affective Evaluation—Innovative Approaches 65 (4), 376–387.
- Steele, M., Gillespie, B., 2001. Shared control between human and machine. In: *Proceedings of the Human Factors and Ergonomics Society*, Minneapolis, MN. HFES.
- Sun, C., Xia, S., Sun, D., Chen, D., Shen, H., Cai, W., 2006. Transparent adaptation of single-user applications for multi-user real-time collaboration. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13 (4), 531–582.
- Swindells, C., Smith, J.D., MacLean, K.E., 2005. An exploration of representations to aid design of haptic behaviours. In: *CHI 2005 Workshop-Hands-on Haptics: Exploring Non-Visual Visualisation Using the Sense of Touch*, Portland, OR. ACM, New York, pp. 5–8.
- Swindells, C., Maksakov, E., MacLean, K.E., Chung, V., 2006. The role of prototyping tools for haptic behavior design. In: *IEEE Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems (HAPTICS 2006)*, Washington, DC, USA.
- Tan, H., Durlach, N., Reed, C., Rabinowitz, W., 1999. Information transmission with a multifinger tactual display. *Perception and Psychophysics* 61 (6), 993–1008.
- Tan, H.Z., Reed, C.M., Delhorne, L.A., Durlach, N.I., Wan, N., 2003. Temporal masking of multidimensional tactual stimuli. *Journal of the Acoustical Society of America* 114 (6), 3295–3308.
- Tang, A., McLachlan, P., Lowe, K., Chalapati, R.S., MacLean, K.E., 2005. Perceiving ordinal data haptically under workload. In: *Seventh International Conference on Multimodal Interfaces (ICMI '05)*, Trento, Italy, pp. 244–251.
- Ternes, D., 2007. Building large sets of haptic icons: rhythm as a design parameter, and between-subjects MDS for evaluation. M.Sc. Thesis, University of British Columbia, 200pp.
- Tse, E., Greenberg, S., 2004. Rapidly prototyping single display groupware through the SDGToolkit. In: *Proceedings of the Fifth Australasian User Interface Conference*, pp. 101–110.
- van Erp, J., 2002. Guidelines for the use of vibro-tactile displays in human computer interaction. In: *Proceedings of EuroHaptics*, Edinburgh, UK.
- van Erp, J.B.F., Spapé, M.M.A., 2003. Distilling the underlying dimensions of tactile melodies. In: *Eurohaptics*, Dublin, Ireland, pp. 111–120.
- Xia, S., Sun, D., Sun, C., Chen, D., Shen, H., 2004. Leveraging single user applications for multiuser collaboration: the cword approach. In: *ACM Conference on Computer Supported Cooperative Work (CSCW'04)*, Chicago, USA, pp. 162–171.
- Yohanan, S., Chan, M., Hopkins, J., Sun, H., MacLean, K.E., 2005. Hapticat: exploration of affective touch. In: *Seventh International Conference on Multimodal Interfaces (ICMI '05)*, Trento, Italy, pp. 244–251.
- Young, J.J., Tan, H.Z., Gray, R., 2003. Validity of haptic cues and its effect on priming visual spatial attention. In: *11th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Los Angeles, CA. IEEE Computer Society, Silver Spring, MD, pp. 166–170.