

Designing Haptic Icons to Support an Urgency-Based Turn-Taking Protocol

by

Andrew Chan

B.Sc., University of British Columbia, 2002

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE STUDIES

(Department of Computer Science)

We accept this thesis as conforming
to the required standard

The University of British Columbia

October 2004

© Andrew Chan, 2004

Abstract

Collaboration is taking place increasingly between individuals living in different cities, countries, or continents. Instead of relying on time-consuming, expensive, and exhausting business travel, companies are turning to web conferencing systems, Internet-based systems that support distributed meetings, training, and collaboration. These systems support *view-sharing*, where an individual can share an application with his or her collaborators, allowing them to view and interact with the application in real-time. While flexible, these systems only permit one user to control the application at a time, necessitating a turn-taking protocol.

Current web conferencing systems depend heavily on visual elements like dialog boxes or tool-tips to deliver messages such as requests for control. However, the collaborative tasks being performed are typically highly visual in nature themselves, meaning that messages can either intrude or be missed. Another shortcoming of current systems is that they fail to support flexibility in requesting control, something we take for granted in face-to-face collaboration.

In this thesis, we introduce a novel urgency-based turn-taking protocol, where users can request control with two levels of urgency or immediately take control. Haptic icons, touch-sense stimuli that have been assigned a meaning, are used in this protocol to periodically inform a user of the current turn-taking state. Our research was conducted in three phases. First, we designed the protocol and selected a set of haptic icons. Next, we evaluated the ability of subjects to learn the haptic icons and identify them under different amounts of cognitive workload. Finally, we recruited groups of subjects to use the protocol in a collaborative environment and evaluated their performance.

Our results show that haptic feedback is a viable channel for communicating turn-taking information. The haptic icons can be learned in a reasonable amount of time and recalled with high accuracy. As well, users in control are more responsive to requests for control and control is shared more equally among group members when haptic feedback is present. The urgency-based protocol also shows promise when used with haptic feedback.

Contents

Abstract	iii
Contents.....	v
List of Tables.....	ix
List of Figures	xi
Acknowledgements	xiii
1. Introduction	1
1.1 Motivation.....	1
1.2 Research Approach and Overview.....	3
2. Background and Related Work	5
2.1 Groupware	5
2.1.1 Multi-user synchronous groupware	6
2.1.2 Single-user synchronous groupware.....	8
2.1.3 Current Systems.....	11
2.2 Haptics	17
2.2.1 Psychophysical properties of haptics.....	17
2.2.2 Haptics in Teleoperated and Virtual Environments.....	18
2.2.3 Haptic Communication.....	19
2.3 Summary.....	20
3. Designing an Urgency-Based Turn-Taking Protocol	21
3.1 An Urgency-Based Turn-Taking Protocol	21
3.1.1 Turn-Taking in Conversation	22
3.1.2 Protocol Description	23
3.2 Icon Delivery and Control Input Device: A Haptic Mouse	23
3.3 Prototyping Haptic Icons to Support the Protocol	28
4 Study 1: Selecting Haptic Icons for the Urgency-Based Turn-Taking Protocol	31
4.1 Multidimensional Scaling	32
4.2 Method.....	33
4.3 Analysis	36
4.4 Results.....	36

4.4.1	Analysis of MDS Graphs.....	37
4.4.2	Likert Scale Responses.....	42
4.4.3	Selecting the In Control Icons	43
4.4.4	Confirming the Change in Control icons.....	44
4.5	Haptic Icons	44
4.6	Summary	45
5	Study 2: Learning and Using Haptic Icons in the Presence of Workload	47
5.1	Experiment Procedure.....	47
5.1.1	Learning Phase	49
5.1.2	Evaluation Phase.....	50
5.2	Performance Metrics.....	53
5.3	Hypotheses.....	53
5.4	Results.....	54
5.4.1	Learning Time	54
5.4.2	Detection Time Hypothesis	55
5.4.3	Identification Time Hypothesis	57
5.4.4	Correct Identification Hypothesis.....	58
5.4.5	“Mistake” Hypothesis.....	58
5.4.6	Distracter Task Performance	60
5.5	Discussion.....	60
5.6	Summary	62
6	Study 3: Evaluating the Urgency-Based Turn-Taking Protocol.....	65
6.1	Research Questions.....	65
6.2	Conditions.....	66
6.3	Study Setup	68
6.4	Task.....	72
6.5	Study Procedure.....	73
6.6	Study Design and Subjects	74
6.7	Dependent Measures.....	74
6.8	Results.....	74
6.8.1	Learning and Using Haptic Stimuli	75
6.8.2	Collaborative Style	76
6.8.3	Equitability of Sharing Control	80
6.8.4	Subject Preferences.....	82
6.8.5	Task Performance	83
6.9	Discussion.....	84
6.10	Summary	86
7	Conclusions and Future Work.....	89
7.1	Using MDS to Categorize Haptic Icons.....	89
7.2	Learning Haptic Icons.....	90
7.3	Identifying Haptic Icons while Engaged in Other Tasks	90
7.4	Haptic Feedback for Mediating Turn-Taking	91
7.5	Value of the Urgency-Based Turn-Taking Protocol	92
7.6	Future Work.....	92

Bibliography.....	95
Appendix A: Study 1 Materials.....	99
A.1 Consent Form.....	100
A.2 Study 1 Instructions	102
A.3 Pre- and Post-Study Questions.....	103
A.4 MDS Graphs	104
A.5 Likert Scale Responses	117
Appendix B: Study 2 Materials.....	119
B.1 Consent Form	120
B.2 Study 2 Instructions.....	122
B.3 Post-Study Interview Questions.....	126
Appendix C: Study 3 Materials.....	127
C.1 Consent Form	128
C.2 Study 3 Instructions.....	130
C.3 Questionnaires.....	140
C.4 Post-Condition Likert Scale Responses	143
C.5 Floor-layout Problems and Solutions.....	159

List of Tables

Table 2-1 - Groupware Taxonomy	6
Table 2-2 - Methods for releasing, assigning, and handling requests for control	9
Table 3-1 - Description of the state transitions in Figure 3-1.....	25
Table 3-2 - Haptic input for obtaining and releasing control	27
Table 4-1 - Haptic Stimuli evaluated in Study 1	34
Table 4-2 - Partitioning of Study 1 subjects into groups for MDS analysis.....	39
Table 4-3 - Observations from analysis of MDS graphs.....	41
Table 4-4 - Haptic icons selected after Study 1.....	45
Table 5-1 – Haptic icon set used in Study 2.....	48
Table 5-2 - Mean detection times for each condition	55
Table 5-3 - <i>p</i> -values for differences in detection times across conditions	56
Table 5-4 - Mean identification times for each condition	58
Table 5-5 - Number of false alarms in each condition	59
Table 5-6 - Percentage of missed transitions for each condition.....	59
Table 6-1 - Haptic icons used in Study 3	67
Table 6-2 - Number of changes in control for each condition	77
Table 6-3 - Gentle Requestor’s Perspective. Average time from a gentle request until gaining control, by group and condition.....	81
Table 6-4 - Urgent Requestor’s Perspective. Average time from an urgent request until gaining control, by group and condition.....	81
Table 6-5 - Control-Holder’s Perspective. Average lengths of periods in control before releasing or losing control, by group and condition	81
Table 6-6 - Equitability of Control Time. Spread of percentage of time in control, between group members most and least in control	81
Table 6-7 - Condition Preference	82
Table 6-8 - Normalized task scores	84

List of Figures

Figure 2-1 - Request for control in Microsoft NetMeeting	12
Figure 2-2 - Request for control in WebEx	13
Figure 2-3 - Giving control to a remote user in WebEx	13
Figure 2-4 - Giving control in Microsoft LiveMeeting	14
Figure 2-5 – Transient window in Breeze showing a request for control	15
Figure 2-6 - Persistent window in Breeze showing a request for control	15
Figure 3-1 - State transition diagram for the <i>urgency-based</i> turn-taking protocol.....	24
Figure 3-2 - Logitech <i>iFeel</i> mouse with thumb buttons	26
Figure 3-3 - Immersion Studio screenshot	28
Figure 4-1 - Visual Basic application for sorting haptic stimuli	35
Figure 4-2 - MDS Graph for all 10 subjects, perspective projection	39
Figure 4-3 - MDS Graph for all 10 subjects, looking down Dimension 1	39
Figure 4-4 - MDS Graph for all 10 subjects, looking down Dimension 2	39
Figure 4-5 - MDS Graph for all 10 subjects, looking down Dimension 3	39
Figure 5-1 - Screen for exploring haptic icons in Study 2.	49
Figure 5-2 - Study 2 learning test.....	50
Figure 5-3 - Visual distracter task in Study 2.....	52
Figure 5-4 - Mean detection times for each condition	55
Figure 5-5 - Mean identification times for each condition.....	57
Figure 6-1 - User Window and Button Bar	66
Figure 6-2 - Layout of experiment room in Study 3	69
Figure 6-3 - Audio setup of experiment room in Study 3	71
Figure 6-4 - Distribution of verbal methods of requesting control	78
Figure 6-5 - Distribution of non-verbal methods for requesting control.....	79

Acknowledgements

I was fortunate to have not one, but two wonderful supervisors who provided ideas, insight, and guidance during my research: Dr. Joanna McGrenere and Dr. Karon MacLean. Thank you both for your encouragement and for always being an email away. I also wish to thank Dr. Kellogg Booth, the third member of my supervisory committee, and Dr. Gail Murphy, my third reader. Gail introduced me to academic research while I was an undergraduate, hiring me as a research assistant. The marvelous experience I had factored into my decision to pursue this degree.

I also had the pleasure of being part of two research labs. Many of my lab-mates in the Imager Graphics, HCI and Visualization Lab participated in my studies and provided valuable feedback. The friendly disposition of the Sensory Perception and INteraction (SPIN) Lab members made up for the lack of windows in the lab. The SPIN members also graciously gave up their workstations while I piloted and ran my third study. I wish to thank Colin Swindells, who always was willing to lend an ear and offer advice, Mario Enriquez, who answered a lot of multidimensional scaling and haptics questions, and Jocelyn Smith, who proofread this thesis.

Several people assisted me in the preparations for the third study. Glen Lee and Luc Dierckx from the CS tech staff reconfigured, reinstalled, and debugged problems on the Windows and Linux machines respectively. As busy as they are, they almost always had time to help when I dropped by their offices. Bruce Dow took great care in adding extra buttons to the haptic mice we used. Jason Harrison taught me how to set up and use the audio equipment, and loaned me some of his own.

My thanks also go to my parents, who taught me to seize every opportunity to explore and to learn. Their unfailing love and support were a tremendous encouragement to me. I also wish to thank my friends for their many words of encouragement. Finally, I wish to thank Maggie Wong for being at my side.

ANDREW CHAN

*The University of British Columbia
October 2004*

Chapter 1

Introduction

1.1 Motivation

Collaboration is taking place increasingly between individuals living in different cities, countries, or continents. Instead of relying on time-consuming, expensive, and exhausting business travel, organizations are turning to solutions such as videoconferencing to allow dispersed teams of individuals to meet together. In the last decade, the rapid proliferation of high-speed Internet access has led to the rise of web conferencing systems, Internet-based systems that support distributed meetings, training, and collaboration. These systems support *view-sharing*, where an individual can share an application with his or her collaborators, allowing them to view and interact with the application in real-time. A major advantage of these systems is that they are flexible: almost all applications can be shared without modifications to the software, and no special hardware or networking infrastructure is required. As a result, ad-hoc conferences can be created quickly. However, these systems are also limited in that only one user can interact with the shared application at a time.

To illustrate this, suppose members of a vehicle design team are meeting to review computer-generated models of a proposed sedan. Senior management has asked for some changes to be made immediately before a vote is cast on whether to proceed with the project. Three members of the design team are at the company's West coast studio, and one member is on the East coast, attending meetings with management. The lead designer is attending a

vehicle launch at the Paris Auto Show. They agree on a meeting time, with one of the West coast members acting as the host.

When the time arrives, the members connect to the host's computer and also join a telephone conference call. The host immediately starts Computer-Aided Design software and loads a 3D model for review. Once she has done this, the member on the East Coast begins rotating the model, pointing out management's concerns. The lead designer suggests changing the height of the windows, manipulating the model to demonstrate his idea. Another member objects, rotating the model to the rear-quarter view and pointing out a consequence of the change. In this manner, the team members continue to modify the model and discuss their changes until they are satisfied.

Since only one user can interact with the shared application at a time, there must be a means for group members to take turns controlling it. Typically, a user who wants control requests it by selecting a menu item or pressing a GUI button. The user in control is then notified by a tool-tip, dialog box, or message window and can choose to accept or deny the request.

Our research focuses on improving system support for collaboration between groups of distributed individuals in situations such as the one described above. In particular, we are interested in ways of facilitating turn-taking between collaborators. There are four key shortcomings in current systems. First, current systems depend heavily on visual elements to deliver information, such as requests for control. Given that computer-mediated collaboration tends to involve tasks that are highly visual as well the visual elements for changing control may distract or impede a user who is working; if the user is deeply engrossed in the task, the notification may be missed altogether. Second, current systems also tend to assume the user in control will immediately address a request for control. While a user will likely agree to release control in response to a request, the user will probably do so only once the change he or she is making is complete. In this case, the user must remember that a request was made.

A third shortcoming of current systems is that they usually require the host to act as a moderator, receiving requests for control and deciding when and to whom to grant control. While this may be useful in meetings where the host is the main presenter, and other participants form an audience, in collaborative situations where all members are expected to actively contribute to the object of interest, this becomes a bottleneck. Finally, current systems fail to support flexibility in requesting control, something we take for granted in

face-to-face collaboration. In face-to-face conversation, nonverbal cues are used to indicate a desire to speak, and the urgency with which the individual wishes to speak [24]. In distributed collaboration, the lack of nonverbal cues necessitates a means of explicitly requesting control, but none of the systems we have studied permit the level of urgency to be expressed.

1.2 Research Approach and Overview

In this thesis, we describe our design and implementation of an *urgency-based* turn-taking protocol, where users can request control with two levels of urgency, or immediately take control. Instead of using visual elements, haptic feedback is used to deliver messages such as requests for control, allowing users to focus their visual attention on the task at hand. Haptic feedback also permits users to receive information periodically, so that they can tell at a “haptic glance” [41] whether they are in control, waiting for control, or simply observing the actions of their collaborators. In particular, a user in control can be intermittently reminded of his or her collaborators’ (possibly changing) intentions. While haptic feedback has been used to recreate real-world physical forces in virtual and teleoperated environments, research into conveying messages through haptic feedback is just beginning.

Our research was divided into three phases. In the first phase, we designed the *urgency-based* protocol and prototyped a set of *haptic icons* [28] in support of it. A haptic icon is a brief haptic stimulus to which a meaning has been associated. We then conducted Study 1 to select an optimal set of haptic icons. In the study, we asked subjects to repeatedly sort a set of stimuli, including the prototyped icons, into different numbers of categories. Our goals were to:

1. Select a set of mutually distinguishable haptic icons
2. Select the icons such that icons with related meanings would also feel similar
3. Ensure the icons have appropriate levels of noticeability and pleasantness
4. Determine the parameters subjects use to categorize different haptic stimuli

While the first study yielded a candidate set of haptic icons, we had no way to ascertain whether they could be learned quickly and their meanings recalled easily. If subjects struggled to associate the meanings with the haptic stimuli, this would pose a serious barrier to real-world use. As well, if subjects consistently failed to notice changes in the haptic icon being presented (indicating, for example, a request for control), or subjects misidentified the

icons, the utility of this approach would be diminished. This is especially true in our intended application, where users must identify the icons while engaged in the primary task of collaborating with the other members of their group. As a result, for the second phase of our research, we designed a study where we evaluated subjects' ability to learn the haptic icons and identify them under three levels of workload. The goals of Study 2 were to determine:

1. How the time required to detect a change in haptic icons would be affected by workload
2. How the time required to identify a new haptic icon and the accuracy of identification would be affected by workload
3. How the number of mistakes committed by subjects would be affected by workload

In the third phase, we conducted an observational user study, Study 3, where groups of 4 subjects used our protocol to collaborate on furniture-layout problems. Groups compared three implementations of our protocol, one with haptic feedback, one with visual feedback, and one that combined both modalities. Our goals in Study 3 were to:

1. Observe the effect of modality on collaboration between group members
2. Determine the effect of modality on task performance
3. Learn which modality subjects preferred

This thesis is composed of seven chapters, starting with this introduction. In Chapter 2, we review relevant literature in the areas of groupware and haptics and describe current view-sharing systems. Chapter 3 motivates and introduces our *urgency-based* turn-taking protocol, along with our requirements for prototyping a set of haptic icons to support the protocol. Chapter 4 describes Study 1, where we used a technique called Multidimensional Scaling (MDS) to select an appropriate set of haptic icons to use. Chapter 5 describes Study 2, where we evaluated subjects' ability to learn the icons, and then identify them under different amounts of cognitive workload. In Chapter 6, we describe Study 3, where groups of subjects used our protocol in a collaborative task. Finally, we present the conclusions and future work in Chapter 7.

Chapters 5 and 6 are based on papers the author and his supervisors prepared for submission to conferences [19, 20].

Chapter 2

Background and Related Work

In this chapter we review literature relevant to our research in the areas of groupware and haptics. We begin by defining groupware and identifying its different types. We then discuss research that has taken place within the subcategory of real-time distributed groupware. After this, we narrow our focus and describe research on turn-taking protocols. We also examine current commercial systems that implement some of these protocols, and the shortcomings the systems exhibit. Next, we turn our attention to the area of haptics; we present some of the traditional uses of haptics and review recent work on haptic communication. We close by summarizing the current state of research in turn-taking protocols and haptic communication.

2.1. Groupware

Groupware has been defined as “Computer-based systems that support groups of people engaged in a common task (or goal) and that provide an interface to a shared environment” [25]. This broad definition includes many systems that are commonly used today, such as e-mail, chat programs, and electronic bulletin boards. Johansen [39] created the classic 2x2 taxonomy shown in Table 2-1, as reprinted in Baecker, Grudin, Buxton, and Greenberg [11].

Our research focuses on synchronous, distributed groupware. In other words, our target is to support groups of individuals who want to work together simultaneously on a task, but who are geographically dispersed. The first system to do this was presented by Engelbart at the 1968 Fall Joint Computer Conference in San Francisco. In what has become known as “The Mother of all Demos,” Engelbart used the oNLine System (NLS) to demonstrate how he

		Location	
		Co-located	Distributed
Time	Synchronous	Face to Face Interactions - Public computer displays - Electronic meeting rooms - Group decision support systems	Remote Interactions - Shared view desktop conferencing - Desktop conferencing with collaborative editors - Video conferencing - Media Spaces
	Asynchronous	Ongoing Tasks - Team rooms - Group displays - Shift work groupware - Project management	Communication and Coordination - Vanilla e-mail - Asynchronous conferencing, bulletin boards - Structured messaging systems - Workflow management - Version control - Meeting schedulers - Cooperative hypertext, organizational memory

Table 2-1: Groupware Taxonomy. From [39].

could collaborate with a colleague at the Stanford Research Institute in Menlo Park. His paper at the conference also discussed a co-located version of the system [27]. At the time, NLS did not have the ability to pass control among users, but users could access a shared telepointer, allowing them to point at, but not manipulate, objects on the screen. Later work by Engelbart on the AUGMENT system [26] enabled a user to pass control to another user with whom the screen was shared; that user could then manipulate the contents as if they were his or her own.

In implementing synchronous groupware systems (whether distributed or not), a major question has been whether to support multi-user or single-user interaction in a shared application. Systems have been implemented that demonstrate each approach. Below, we describe efforts to date in developing multi-user, synchronous groupware and the challenges in moving such systems from research labs into widespread usage. We then discuss research into single-user, synchronous groupware.

2.1.1 Multi-user synchronous groupware

Multi-user synchronous groupware applications enable multiple users to simultaneously collaborate on a task, and are often referred to as collaboration-aware applications [45]. Two examples from the early 1990's are GROVE [25] and Tivoli [53].

GROVE was a shared text editor that could be used either in co-located or distributed situations. It permitted multi-user input without any locking; social protocols were used to mediate who could edit a certain part of the document. Tivoli implemented an electronic whiteboard where co-located users could simultaneously write and modify information.

Building a multi-user synchronous groupware application is not trivial, particularly when users are distributed. The first decision must be whether to build a centralized or replicated system. A centralized system is simpler to build because a single node coordinates the activity of all users, but the node can become a performance bottleneck. A replicated system, where each node is responsible for keeping itself consistent with the other nodes, scales better but is also more difficult to build. Most multi-user synchronous groupware applications are highly replicated, with minimal reliance on centralized services. These applications must ensure that all nodes execute instructions in the same order (when the instructions may be received out of order) and provide concurrency control so that two users cannot simultaneously change the same object.

Collaboration-aware toolkits provide a layer upon which application designers can build. They abstract away many of the technical difficulties described above in building these applications, allowing designers to rapidly prototype and refine interfaces. One example was DistEdit [44], a toolkit designed to facilitate the modification of text editors into multi-user group editors. More general toolkits include LIZA [31] and Rendezvous [52]. In the last decade, the GroupLab at the University of Calgary has released several groupware toolkits with support for collaboration-aware user-interface widgets. These toolkits include: GroupKit, a toolkit for building groupware applications using Tcl/Tk [55]; SDGToolkit, a toolkit for building single-display (co-located) groupware [62]; and GroupLab Collabrary, a toolkit for building multimedia groupware [16].

The clear advantage of multi-user synchronous groupware is that users can interact with a shared application in parallel. However, with the exception of whiteboard tools offered in web-conferencing systems [3, 8, 9], there are few commercial examples of collaboration-aware applications. Not only are these applications technically difficult to build, but also designing usable multi-user user interfaces is non-trivial. Multi-user systems often relax the WYSIWIS (What You See Is What I See) principle to allow each user to view any portion of the shared object, rather than sharing a single view. As a result, a multi-user system must provide awareness by communicating where all the users are in a shared object [36] and what

they are doing, so that collaborators' actions do not conflict with one another. Simple actions, such as selecting an item from a menu, do not translate seamlessly to a multi-user environment. For example, in a single-user application, seeing the menu confirms a user's own action. In a multi-user environment, the other users may not understand the intentions of a user selecting a menu simply by watching it occur, particularly if the users are not looking at the same view of the shared object as the selector. Another disadvantage of multi-user synchronous groupware is the high cost of redeveloping existing single-user applications to support multi-user interaction, a cost that commercial software vendors have thus far been unwilling to incur. This is exacerbated by the likelihood that even when available, multi-user features will be used less often than those that support a single user [35].

2.1.2 Single-user synchronous groupware

In contrast to multi-user synchronous groupware, single-user synchronous groupware provides an effective "stepping stone" between traditional single-user applications and collaboration-aware applications ([40], as quoted in [34]). Also known as view-sharing systems, these systems allow a user to share a view of a running application with other users; in most systems, the remote users can also control the shared application. The greatest advantage of these systems is that they can be used with nearly all existing software without requiring modifications to the software.

Single-user synchronous groupware systems are often less complex and easier to implement than multi-user systems. They use a centralized architecture to coordinate activity; since only one user can interact with the system at a time, scalability is less of an issue than in a multi-user system. This architecture greatly simplifies concurrency control because, at a minimum, a consistent (if not always fair) protocol can be enforced based on the temporal order in which messages are received. These systems also use a strict WYSIWIS protocol, meaning all users see the exact same view. Thus, the intentions of the user in control are more likely to be understood by the other users.

Since single-user synchronous groupware systems do not support multi-user input, access to the shared application is mediated through a floor-control policy, also known as a turn-taking protocol. Many different policies have been proposed. A taxonomy of possible protocols by Myers, Chuang, Tjandra, Chen, and Lee [50] lists the possible ways of releasing control, assigning control upon its release, and requesting control. Table 2-2 summarizes the

possibilities. The three most-commonly cited protocols are described below with respect to this taxonomy:

- *Give*, which only specifies that an explicit release must occur before control is transferred; different methods may be used for requesting and assigning control.
- *Take*, where an explicit request is followed by an explicit loss and an immediate grant
- *Free-floor*, where an implicit request is followed by an explicit loss and an immediate grant

A fourth protocol typically used to describe interaction in multi-user synchronous groupware systems is *Free-for-all*, where users can work in parallel. There have been many research prototypes that support one or more protocols; a review of early work was done by Greenberg [33]. We highlight several systems, including more recent work.

The MMConf system provided an architecture for building shared multimedia applications [21]. It consisted of a conference manager and a toolkit for building conference-aware applications. Both single-user and multi-user interaction were supported; users could choose one of the four turn-taking protocols described above: *give*, *take*, *free-floor*, and *free-for-all*. MMConf also had limited support for a telepointer, in that the active user's cursor could be displayed to all users. The authors noted that allowing others to control the

	Method	Description
Releasing Control	Explicit Release	User in control must release control before someone else can acquire it
	Implicit Release	System releases control automatically, such as when the user in control has not used system for a period of time
	Explicit Loss	Control is given to another user regardless of whether user in control is finished
Assigning Control	Moderator	One of the users decides who gets control
	Explicit Request	User requests control by pressing a button or equivalent means
	Implicit Request	System interprets input from the user (such as typing or mouse movement) as a request for control
	Rule-Based	Algorithm used to decide who receives control
Request Handling	Immediate Grant	User's request is granted immediately. This only works with Explicit Loss
	Queued	User's request is queued; when user in control releases, person at front of queue gets control
	Ignored	Requests are ignored unless the floor is available

Table 2-2 - Methods for releasing, assigning, and handling requests for control. From [50].

telepointer would be beneficial in certain situations.

A finer-grained turn-taking protocol was proposed by Boyd called *fair dragging* [14]. In this system, the floor would only need to be controlled when repositioning an object. The floor would be obtained, if possible, when a mouse button was depressed, maintained while the object was dragged, and released when the mouse button was released.

While many protocols have been proposed and implemented, few studies have been conducted that formally compare different protocols. We report three comparative studies here. As part of the Pebbles PDA project, Myers, Chuang, Tjandra, Chen, and Lee compared five different single-user turn-taking protocols to a *free-for-all* protocol to solve jigsaw puzzles [50]. Subjects worked in pairs and were co-located. Subjects performed significantly better on the *free-for-all* protocol compared to the other turn-taking protocols, and there were no significant differences between the single-user protocols. The performance benefit of the *free-for-all* protocol was likely due to the nature of the task, which lent itself to a high degree of parallelism.

Inkpen, McGrenere, Booth, and Klawe studied the effect of turn-taking protocols on pairs of co-located children [38]. The children played a game where they had to solve a variety of Rube Goldberg-like puzzles using one of three protocols: sharing a single mouse, *give*, or *take*. They found that boys shared control more equally when using the *take* protocol, and that the amount of time boys had control was positively correlated with their ability to complete the same task on their own. Both results were significant. Girls solved significantly more puzzles using the *give* protocol. Although the result was not significant, boys solved more using the *take* protocol.

A study by McKinlay, Proctor, Masting, Woodburn, and Arnott examined the effectiveness of face-to-face communication and computer-mediated communication using three turn-taking protocols: *free-for-all*, *give*, and *take* [47]. Subjects were randomly assigned to groups of three or six, and completed all four conditions. They were given a hypothetical situation in which they were stranded in the Arctic and told to rank a set of items in order of importance. Besides the face-to-face condition, the only way subjects could communicate was through a chat application, using one of the three protocols. The degree of consensus reached by the group was used as a dependent measure. Although a significant effect of condition was found, no post hoc comparisons were reported to ascertain which protocols (if any) were superior to others. *Face-to-face* was ranked highest, followed by *give*, *free-for-all*,

and *take*. They reported that three-person groups were able to achieve a greater degree of consensus than the six-person groups, although the result was not statistically significant. It is possible that their dependent measure was not sensitive enough to capture differences between protocols; as other research has shown, humans are quite adept at accomplishing tasks even in less-than-ideal circumstances [15].

In summary, view-sharing systems can effectively turn traditional single-user applications into collaborative ones without any modifications. Although many turn-taking protocols have been proposed and implemented, there has been a dearth of evaluation to establish which protocol(s) are superior under what conditions. As Lauwers and Lantz note, it is unlikely that a single protocol will suffice for all groups in every situation [45].

2.1.3 Current Systems

Today, view-sharing software is a standard component in *web conferencing systems*, Internet-based systems that support distributed meetings, training, and collaboration. We begin with a description of Microsoft NetMeeting [9], which was first released in 1996 and represents one of the earlier commercial view-sharing systems. Here, we describe the last version of NetMeeting, version 3.01, released in 2000. Then, we describe three systems that represent the current state-of-the-practice: WebEx [8], Microsoft LiveMeeting [4], and Macromedia Breeze [3]. WebEx is considered the industry leader in web conferencing. LiveMeeting (formerly known as Placeware Conference Center) and Breeze are both designed to challenge WebEx. All of the systems enable a host to select any application on their computer and share it over the Internet with others. As well, they all use a hybrid of the *give* and *take* protocol by default. In each system, the host gives control to others, although the exact ways in which requests for control and granting of control are handled differ between systems. When other users are in control, the host always has the option of taking control back. Each system is described in turn.

Microsoft NetMeeting - In NetMeeting, the host of the shared application must allow remote users to request control. The host also has the option of automatically accepting requests for control, in effect enabling a *take* protocol, but by default, a *give* protocol is used. A remote user requests control by selecting a menu item, and then the host is presented with a dialog box, shown in Figure 2-1, asking whether to accept or reject the request. If the host accepts,

control is immediately passed to the requester, although the host can regain control at any time by pressing a key or a mouse button. If the host rejects, the requester is informed that the request was rejected. The host can also ignore the dialog; it remains visible, but allows the user to continue to work. After approximately 30 seconds, the dialog is removed and the requester is told that the host did not respond to the request. While one user is requesting control, if additional users attempt to request control, they are told that the host is busy. As well, when a remote user is in control of the shared application, other users cannot request control until the host regains control. Instead, the remote user in control has the option of forwarding control to a third user, as long as the host and the third user agree. Thus, the host plays a central role.

WebEx - WebEx uses a slightly different model than NetMeeting for turn-taking. Like NetMeeting, the host of a shared application controls whether requests should be automatically accepted, and remote users select a menu item to request control. Instead of dialog boxes, WebEx uses tool-tips to inform users of requests and changes in control. When a remote user requests control through a menu item, a tool-tip is displayed near the cursor of the user in control, stating that “Attendee: X requests remote control.” This tool-tip is

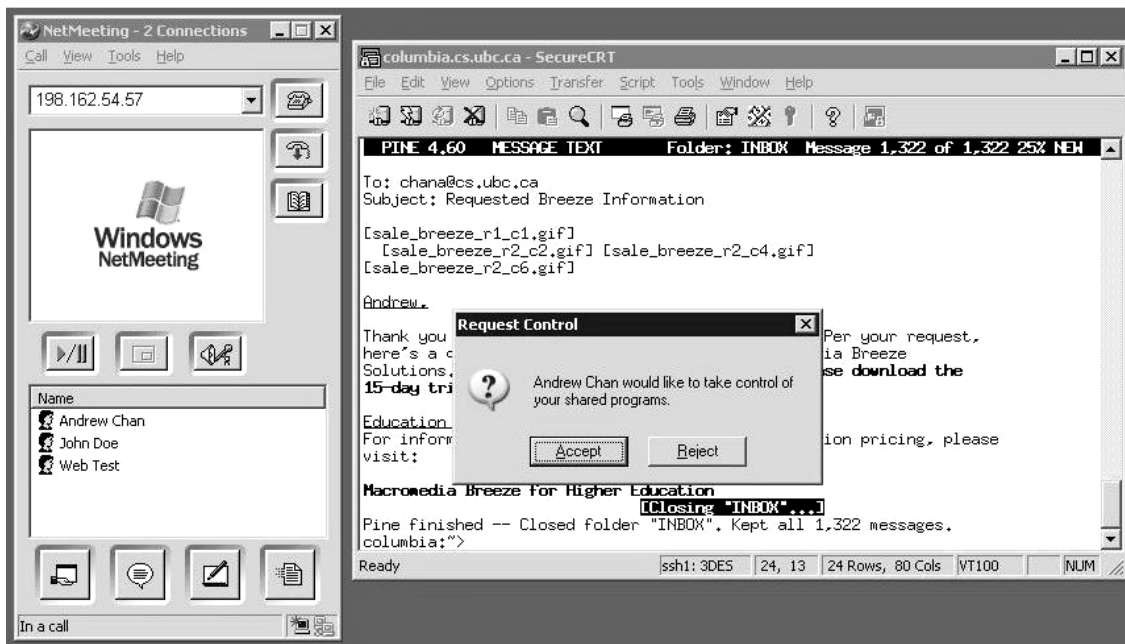


Figure 2-1 - Request for control in Microsoft NetMeeting.

displayed for a few seconds then disappears. The host has the option to give control to any of the remote users, not just those requesting control. Figure 2-2 and Figure 2-3 show this process. When the host gives control to a user, the user gaining control is told by a tool-tip to click to gain control. As in NetMeeting, the host always can regain control by clicking a mouse button. When a remote user requests control and another remote user is in control, both the in-control user and the host see the request, but the in-control user can only release control to the host; there is no facility to automatically transfer control from one remote user to another. Thus, the host fully moderates control in this system.

Microsoft LiveMeeting - The host is also the moderator in LiveMeeting. A button-bar allows the host to start, stop, and pause sharing. It also enables the host to give control to a remote user, as shown in Figure 2-4; the user is informed through a dialog box that he or she has gained control. At any time, the host can regain control by pressing a button on a button-bar.



Figure 2-2 - Request for control in WebEx.

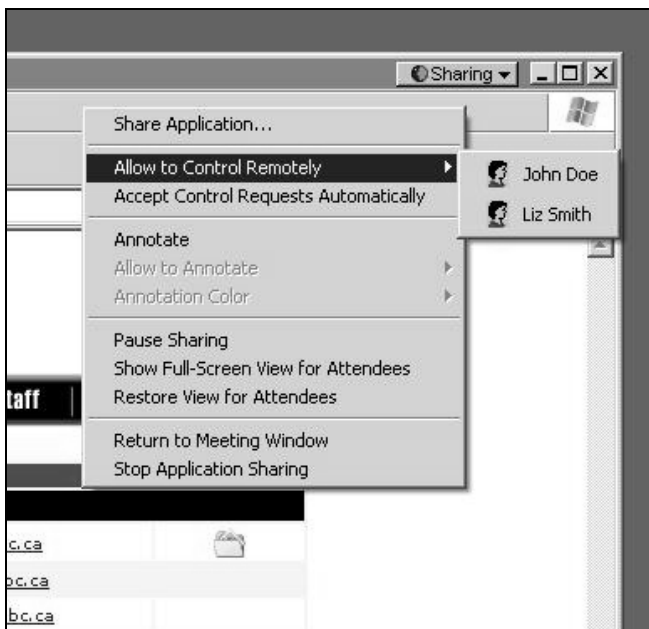


Figure 2-3 - Giving control to a remote user in WebEx.

LiveMeeting has no facility for requesting control; remote users must request control verbally (assuming an audio link is available) or through a text chat message.



Figure 2-4 - Giving control in Microsoft LiveMeeting.

Macromedia Breeze - Breeze has several features that distinguish it from the other products. When a remote user requests control, two floating windows appear on the host's display. One window (Figure 2-5) is displayed for 30 seconds at the top of the Breeze screen, if it is visible. The other window (Figure 2-6) appears in the lower right-hand corner of the display, with buttons for the host to accept or decline the request. The latter window is always visible, and persists until the host accepts or declines the request, or the requester cancels the request. If multiple users request control, a floating window is created for each request, and the windows are stacked on top of one another, with the most recent request on top. The host can process requests in any order he or she chooses. An unusual and somewhat counter-intuitive feature is that *any* user can immediately return control to the host when a remote user is in control. The host acts as a mediator in Breeze; requests for control are only shown to the host, even when a remote user is in control. Remote users cannot forward control to other users.



Figure 2-5 – Transient window in Breeze showing a request for control.




Figure 2-6 - Persistent window in Breeze showing a request for control.

We make several observations about these systems:

- Different methods are used to give control to a remote user. In Breeze and NetMeeting, the host can only give a remote user control in response to a request for control from that user. In LiveMeeting, which has no request mechanism, the host can give any user control. WebEx enables remote users to request control, but also allows the host to give any user control.
- Systems also have different assumptions as to when requests for control will be handled. NetMeeting encourages an immediate response by raising a dialog box that the user is virtually compelled to attend to. The tooltips used by WebEx do not convey the same sense of urgency and thus provide the in-control user with greater flexibility in handling the request, but if a user chooses to retain control to finish what he or she is doing, there is no reminder of the request once the tooltip disappears. Breeze, with its persistent floating windows, provides the best support.
- Only Breeze provides explicit support for requests from multiple users. WebEx and LiveMeeting support this (LiveMeeting verbally), but require the user to remember the requests and their order. NetMeeting provides the least support; it does not

support requests from more than one user at a time, and even if the host declines the requests but remembers the order in which requests were made, he or she cannot give control to the correct user later without asking that user to re-request control.

- All of the view-sharing systems rely on visual cues to convey information. They range from being intrusive to the point of being disruptive (e.g., the NetMeeting dialog boxes) to easily missed (e.g., Breeze and WebEx, when the user's attention is focused on a different part of the screen than where the messages appear). The visual cues can also obscure information on the screen, again disrupting the user's work.
- None of the systems give remote users flexibility in requesting control without resorting to verbal means. When a user requests control in NetMeeting, WebEx, or Breeze (recall that LiveMeeting does not have a mechanism for requesting control), the user cannot convey how urgently he or she wants control to the host. This may affect the quality of collaboration, as a host may interrupt his or her own efforts to respond to what was intended as a low-priority request for control, or the host may choose to finish what he or she is doing, despite the fact that the requestor wants control immediately. Although users can resort to verbal means to indicate urgency, in normal face-to-face collaboration they often do not need to, as we will discuss in the next chapter.
- The host also acts as the moderator in each of the view-sharing systems. This makes sense in a meeting presentation, where the host may be expected to dominate the discussion; it is likely that the designers of the view-sharing systems expected this to be their primary use. However, in a more collaborative setting, such as a design review, it is more likely that users would participate more equally. In such a setting, requiring the host to moderate requests for control could be an annoyance.

Shifting notification messages to a different modality could address most of the issues that have been raised. While auditory feedback could be used to manage the turn-taking process, distributed collaboration systems usually assume the presence of an audio channel to allow users to speak to one another as they work. This means that auditory cues could again be disruptive, or missed by users engaged in conversation. Another approach is to require users to verbally mediate requests for control, as in LiveMeeting. While this works well for small groups, it does not scale well because only one person can be heard at a time,

making it difficult, for example, to request control while another person is speaking. Instead, we consider haptics.

2.2 Haptics

Haptics research explores ways to communicate information through the sense of touch. The sensitivity and acuity of this sense is well known; by picking up an object and running our fingers along it, we can determine a wide range of properties such as its texture, compliance, warmth, contours, and heft [41]. In daily life, we receive a great deal of tactile feedback, along with information from our other senses. However, current user interfaces are highly visual in nature, with audio cues used primarily to reinforce what is shown on a visual display. As effective as this is, the visual and auditory senses cannot convey the same information as the haptic sense; in cases like manipulating an object, this information can be crucial for understanding or using a system. In other circumstances where the visual and auditory senses are already engaged in activity, haptic feedback has the potential to be an additional conduit through which information is communicated. In this section, we describe three areas of haptics research: the psychophysical properties of haptics, the use of haptics to generate physical forces in teleoperated and virtual environments, and haptic communication.

2.2.1 Psychophysical properties of haptics

To use haptics to convey any kind of information, we must understand the limits of our haptic abilities. Klatzky and Lederman measured the ability of subjects to identify haptic sensations at a “haptic glance” [41]. Blindfolded subjects who were permitted to explore everyday objects freely were able to identify them with 93% accuracy. The objects included items such as corduroy, chalk, and a bread pan. When subjects were only permitted to contact the object with their fingers (moving them as little as possible) and had exposure time limited to 3 seconds or less, subjects still were able to identify objects with above-chance accuracy. These researchers have also documented how we perceive texture when feeling it through a rigid probe [42, 43]. Inspired by the Tadoma method of communication, where a deaf listener can understand speech by putting his or her hands on the speaker’s face, Tan, Durlach, Reed, and Rabinowitz measured the rate at which vibrotactile information could be transmitted using a custom-built device called the TACTUATOR. They found the optimal rate

(maximizing the number of stimuli felt, while minimizing errors in identifying them) to be 2 – 3 stimuli / second [61].

2.2.2 Haptics in Teleoperated and Virtual Environments

The earliest applications of haptics have been in the areas of teleoperation and virtual environments. Teleoperation involves manipulating a robotic arm remotely, and is useful in situations where robots can be placed where humans cannot. Even with a visual display of the work area, the arm is difficult to control, since the visual display cannot convey properties like heft, compliance or resistance. Haptic feedback mimics these properties and affords operators a greater degree of control (see [56] for an example). Haptic feedback is also valuable in laparoscopic surgery simulation, where the goal is to minimize tissue damage [63]. Again, the feedback simulates the sensation that would be felt if the surgeon contacted the surface with tools directly.

Studies have examined whether haptic feedback can assist in target acquisition tasks routinely performed on desktop computers. Rosenberg and Brave [57] presented preliminary results suggesting that either passive or active haptic feedback in a force-feedback joystick would improve target acquisition times in a Fitts' Law task. The passive feedback was designed such that more force would be required to move beyond the target, and the active feedback was designed such that an attractive force field surrounded the target. Dennerlein, Martin, and Hasser [22] found that haptic feedback on a force-feedback mouse significantly improved performance on steering and combined steering / targeting tasks through a tunnel, which are similar to tasks like selecting a nested menu item. The haptic feedback was designed so that the tunnel walls would repel the cursor towards the center of the tunnel, with the force magnitude inversely proportional to the distance from the nearest wall.

Several researchers have examined whether haptic feedback can create a greater sense of presence or togetherness in a virtual collaborative environment. In one study by Basdogan, Ho, Srinivasan, and Slater [12], pairs of subjects were instructed to move a ring along a wire, trying to minimize contact between the ring itself and the wire. The ring only moved when both subjects exerted a certain amount of force, and the movement was based on the directions in which they exerted force. Subjects had significantly better performance in the haptic condition, compared to a condition without haptic feedback. They also reported a significantly higher sense of togetherness. A study by Oakley, Brewster, and Gray [51] asked

pairs of subjects to create UML diagrams using a groupware editor. Haptic effects were added so that the subjects could locate one another and move each other around; these effects were found to increase the sense of “presence,” which was defined as being physically present and engaged in a natural environment. Sallnas, Rassmus-Grohn, and Sjoström [59] studied presence more carefully, investigating whether haptic feedback could improve virtual presence, defined as the sense of being in the virtual environment, or social presence, the sense of being together with the remote collaborator. In their study, subjects worked together to stack blocks in a certain order, with haptic feedback simulating mass, friction, and damping. They found that the sense of virtual presence was significantly higher when haptic feedback was present, but there was no significant difference in the sense of social presence.

2.2.3 Haptic Communication

In the work reviewed so far, haptic feedback has been used to recreate real-world physical forces in virtual or teleoperated environments. However, recent research has asked whether haptic feedback could be used to send messages, where a haptic stimulus would be associated with a given meaning. This research has its basis in work on creating auditory signals. Gaver [30] proposed creating “auditory icons,” real-world sounds with an intuitive mapping to an action. In contrast, Brewster, Wright, and Edwards [18] proposed creating synthetic sounds called “earcons,” whose meanings would have to be learned.

Haptic messages would be especially beneficial in situations where the visual system (and possibly the auditory system) is highly engaged in a task. One example is in driving. Vehicle cockpits have become more complex; features like navigation systems are no longer exclusive to luxury vehicles. New features in audio and climate control systems have also contributed to this complexity. However, these systems often require the driver’s visual attention to be used, thus creating a dangerous distraction. For this reason, automobile manufacturers are introducing haptics into cockpits, with BMW’s iDrive [1] the first to market. The iDrive consists of a force-feedback rotary knob that is used to access vehicular functions displayed on a screen. Different detents are felt for different types of menu items (Swindells, C., personal communication, September 14, 2004). Cellular telephone manufacturers are also experimenting with haptics. Just as a customized ring-tone can identify a caller, they hypothesize that vibrotactile “touch tones” may perform the same function with less intrusiveness.

The Sensory Perception and INteraction (SPIN) Lab at the University of British Columbia is at the forefront of haptic communication research, investigating how to create haptic icons, haptic sensations that have an associated meaning. Previous work has included a graphical editor for creating haptic icons on a rotary knob [28], and a technique for designing haptic icons that are perceptually distinct from one another [46]. Our research used and extended this technique to create families of haptic icons. The Glasgow Multimodal Interaction Group is also active in creating “tactons,” which serve a similar purpose [17].

2.3 Summary

While many turn-taking protocols have been proposed, relatively few studies have evaluated their effectiveness. This may be in part due to the recent emphasis on studying multi-user distributed groupware. However, with the availability of commercial view-sharing software using different variations of the *give* and *take* protocols and relying on human mediation, renewed interest in this area is warranted. An analysis of four current systems reveals shortcomings in their implementations, rooted in their dependence on visual cues for sending messages. A possible way of addressing these shortcomings is through haptic feedback. Although research in haptic feedback has traditionally involved reproducing physical forces felt in the real-world, more recent research has turned to conveying abstract information through haptic icons.

Designing an Urgency-Based Turn-Taking Protocol

In the previous chapter, we discussed some of the shortcomings of current view-sharing systems and the turn-taking protocols they use. In particular, we noted their reliance on visual cues to convey requests for control and we proposed using haptic feedback to convey this information instead.

In this chapter, we discuss our first steps towards implementing a novel *urgency-based* turn-taking protocol that uses haptic icons to communicate information about the current turn-taking state. We define a haptic icon as a haptic stimulus to which a meaning has been assigned. The chapter is divided into three sections. In the first section, we provide motivation for and describe the protocol we developed. Next, we describe the haptic device we selected, a commercially produced mouse with a vibrotactile display embedded in it. We discuss the advantages and disadvantages of this approach. Finally, we describe the process used to prototype three families of haptic icons to support our protocol.

3.1 An Urgency-Based Turn-Taking Protocol

In discussing current commercial view-sharing systems, we noted that the systems do not allow users to indicate the urgency with which they want control. In this section, we continue to motivate the need for an *urgency-based* turn-taking protocol by drawing on conversational analysis and its insights into how we take turns speaking. We note that while we have different ways of obtaining the floor in spoken conversation, the turn-taking models for collaboration thus far do not support the same degree of flexibility. We then describe the

turn-taking protocol we developed that allows users to request control with different levels of urgency.

3.1.1 Turn-Taking in Conversation

In face-to-face conversation, a variety of techniques are used to obtain and maintain the floor. A model for turn-taking in conversation by Sacks, Schegloff, and Jefferson [58] observed that as a speaking turn comes to a close, three outcomes are possible: the speaker selects the next person to speak, a different speaker self-selects to speak next, or the speaker takes another turn. Additional research by Duncan [23] and Duncan and Niederehe [24] showed that several mechanisms are used in yielding a turn, maintaining a turn, and obtaining a turn:

- Turn-yielding may arise through speech content, syntax, intonation, paralanguage (vocal effects such as pitch, loudness, or stressing), or body-motion.
- A listener can encourage a speaker to continue through different back-channel behaviors, including head movements, short verbalizations (“mm-hmm”), or other short statements that indicate that the speaker has been understood.
- A speaker can suppress an attempt by a listener to take a speaking turn through a gesticulation, such as a raised hand.
- A listener can request a speaking turn through a shift in head direction away from the speaker, audible inhalation, initiation of a gesticulation, or paralinguistic overloudness in back-channel communication.

The manner in which these mechanisms are employed affects how the recipient perceives their meaning. When we are engaged in conversation and wish to speak, there are three strategies we might employ. We might wait until the speaker gives an appropriate yielding signal, take a breath, and start a speaking turn. However, if we perceive that the speaker has seriously misunderstood something that was said, we may wish to quickly clarify the situation, either through a combination of methods to request control (e.g., raising a finger and taking a sharp breath), or through a vigorous gesture, such as raising both hands quickly. Lastly, we may even find it necessary to interrupt the speaker. The ability to request control using these three strategies and the ability to convey the urgency with which we want control are typical features of everyday conversation.

3.1.2 Protocol Description

The substantial role communication plays in collaboration leads us to expect that flexibility in the means for requesting control will be useful. In particular, we believe that allowing users to express the urgency with which they wish to obtain control will be beneficial. However, none of the turn-taking protocols we have observed provide this flexibility. To address this, we devised a protocol where a user has three means of requesting control, roughly corresponding to the three strategies we described in the previous paragraph. We also shifted the burden of mediating requests for control from a host to a rule-based algorithm.

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, 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 showing the possible states in the protocol is shown in Figure 3-1, with an explanation of each of the transitions in Table 3-1.

3.2 Icon Delivery and Control Input Device: A Haptic Mouse

We selected Logitech *iFeel* mice to deliver haptic feedback. These are standard optical mice with an embedded vibrotactile display, using technology licensed from Immersion Corp. Haptic feedback is generated through a plastic gear train driving an eccentrically mounted rotating mass [10]. An obvious drawback to the mouse approach is that haptic feedback can only be felt when a user's hand is on the mouse, but we made the simplifying assumption that our evaluation would be mouse-based.

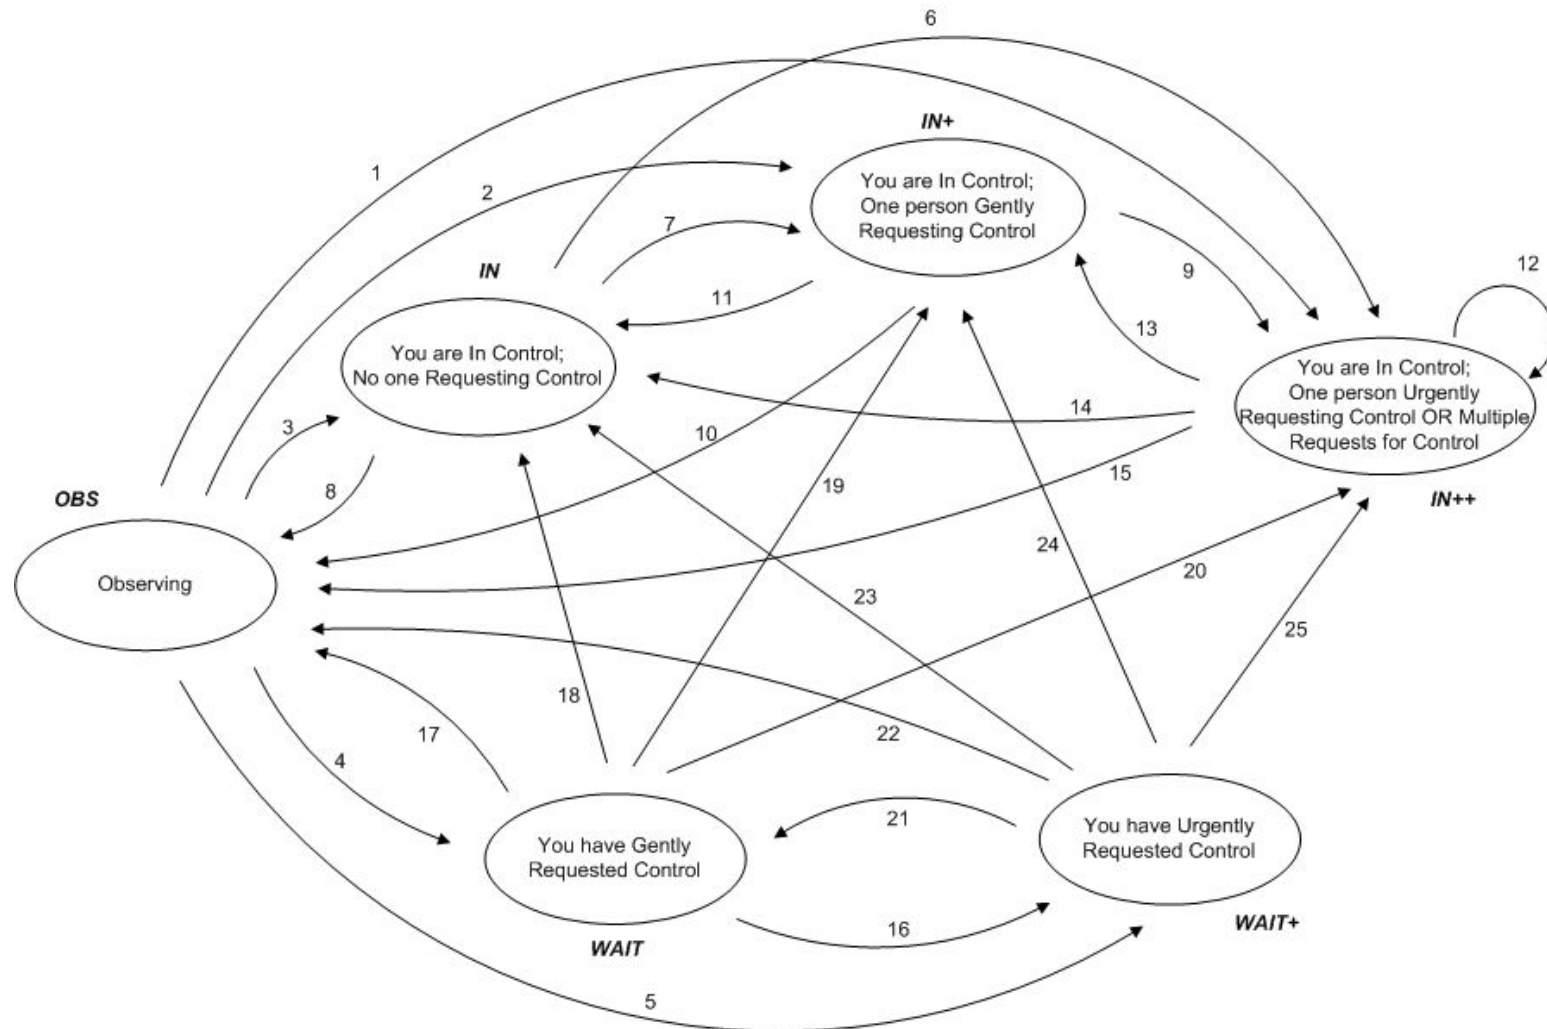


Figure 3-1 - State transition diagram for the *urgency-based* turn-taking protocol. Numbers correspond to transitions in Table 3.1. Labels for each state are shown in italicized text.

Initial State	Trans	Description
OBS	1	You have taken control from the person in control; there is one person urgently requesting control or there are multiple people requesting control.
	2	You have taken control from the person in control; there is one person gently requesting control.
	3	You have taken control from the person in control with no one waiting, OR you have obtained control immediately following a take, gentle request, or urgent request because no one was in control.
	4	You have gently requested control; someone else is currently in control
	5	You have urgently requested control; someone else is currently in control
IN	6	Someone has urgently requested control.
	7	Someone has gently requested control.
	8	You have released control OR someone took control from you.
IN+	9	Someone has requested control.
	10	You have released control OR someone took control from you.
	11	The person gently requesting control cancelled the request.
IN++	12	Another person has gently or urgently requested control OR a person has cancelled their request for control; someone is still urgently requesting control or multiple people are still requesting control.
	13	Someone has cancelled his or her request for control, but there is still someone gently requesting control.
	14	Someone has cancelled his or her urgent request for control; no one is requesting control.
	15	You have released control OR someone took control from you.
WAIT	16	You have urgently requested control; someone else is currently in control.
	17	You have cancelled your request for control.
	18	You have taken control from the person in control OR the person in control released it; no one else is waiting.
	19	You have taken control from the person in control OR the person in control released it, you were the first to gently request it, and no one urgently requested control; someone else is gently requesting control.
	20	You have taken control from the person in control OR the person in control released control, you were the first to gently request it, and no one urgently requested control; multiple people are gently requesting control.
WAIT+	21	You have downgraded your request for control to a gentle request.
	22	You have cancelled your request for control.
	23	You have taken control from the person in control OR the person in control released it; no one else is waiting.
	24	You have taken control from the person in control OR the person in control released control and you were the only person to urgently request it; someone else is gently requesting control.
	25	You have taken control from the person in control OR the person in control released control and you were the first to urgently request it; multiple people are requesting control or someone else is urgently requesting it.

Table 3-1 - Description of the state transitions in Figure 3-1.



Figure 3-2 - Logitech *iFeel* mouse with thumb buttons.

As part of our approach, we wanted to examine whether all actions related to turn-taking could be incorporated into the *iFeel* so that extra visual items would not be needed. This meant that we had to have a means of obtaining and releasing control other than by selecting a GUI widget. We achieved this by adding two thumb buttons to the mouse; Figure 3-2 shows a picture of the modified *iFeel*. The buttons were powered and polled through the parallel port of the computer to which the mouse was attached.

Originally, we had considered using a force-sensing resistor so that users would literally squeeze the mouse, and that the urgency of the request would be linked to the force applied. However, we decided that users might have difficulty determining how much force was required to gently request, urgently request, or take control. Buttons, with their binary action, would not have this problem. Although our mice can no longer be strictly called off-the-shelf, we note that commercially available mice have similar buttons, such as the Logitech MX 1000 [2]. In those mice, one button acts as the forward button, and one button acts as a back button during Internet browsing. We used a similar metaphor when designing the button presses for obtaining and releasing control: the front button was used for increasing the level of urgency, and the back button was used for canceling or releasing control. The button presses for obtaining and releasing control are shown in Table 3-2.¹

The frequency, amplitude, and “rhythm” of feedback provided by the *iFeel* can be manipulated. Had a different type of device been used, such as a force-feedback knob, other parameters could have been manipulated, such as waveform [46]. Stimuli with frequencies ranging from 0.01 Hz to 500 Hz can be created. A software API allows developers to specify intensity values between 0 and 10 000 to influence the amplitude of feedback at a given frequency. Using this API, developers can specify the initial intensity of the vibration (known

¹ In Study 3, where we evaluated our protocol, we decided not to allow urgent requests for control to be downgraded to gentle requests. This was done to simplify our protocol.

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

Table 3-2 - Haptic input for obtaining and releasing control.

as the attack), the final intensity of the vibration (known as the fade), and intensity in between (known as the magnitude). Different rhythms can be created by creating multiple stimuli and combining them.

Due to its design, the amplitude of feedback produced by the *iFeel* is dependent both on the intensity values specified and the frequency. Below 20 Hz and above 250 Hz the range of feedback feels constrained. As a further confound, human perception of the salience or intensity of haptic feedback depends on both amplitude and frequency. While the designers of the API for the *iFeel* could have accounted for this so that stimuli with the same magnitude at different frequencies have the same perceptual intensity, they chose not to. As a result, a 2000-magnitude vibration at 50 Hz feels noticeably stronger at 100 Hz.

Despite these limitations, we wanted to challenge the assumption that haptic technology is not ready for mainstream use by using off-the-shelf technology. This approach also enabled us to use Immersion Studio, a GUI application for generating haptic stimuli, to rapidly prototype different stimuli. Stimuli generated are written to a file and later recreated by accessing the file through a software API. Several programming languages are supported, including Visual Basic, C++, and Java. A screenshot of this application is displayed in Figure 3-3.

The implementation of our turn-taking protocol was tied to the haptic device we chose. For example, a force-feedback knob delivers different kinds and a wider range of sensations than a vibrotactile display, albeit across a smaller frequency range. However, in our intended application, users would not be able to hold the knob continuously. Another option could have been to place the vibrotactile display on an arm-band. This would allow users to receive haptic feedback regardless of their current actions, and raises interesting questions of when feedback should be delivered: for example, should a user continue to receive feedback when he or she temporarily steps out of the collaboration to take a phone

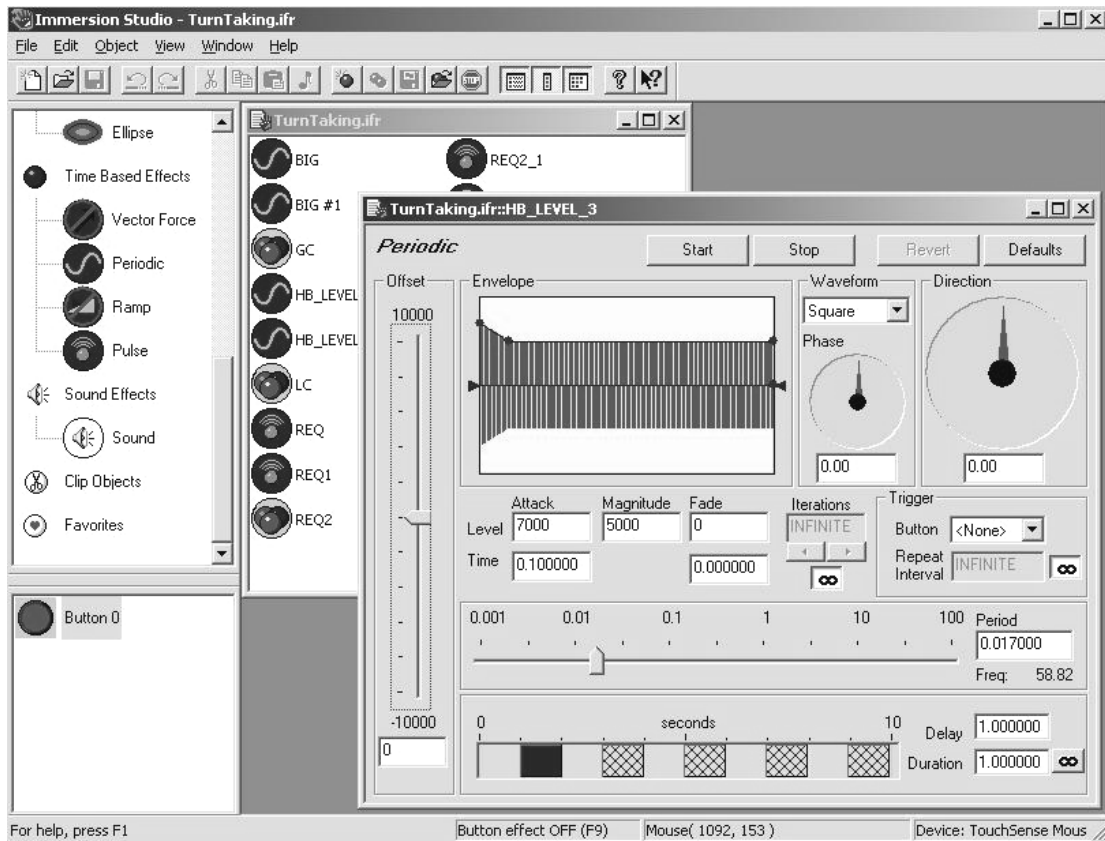


Figure 3-3 - Immersion Studio screenshot. The foreground window shows the settings for a periodic stimulus. Some settings, such as Offset, Waveform, and Direction, do not apply for a vibrotactile display, but would for a force-feedback knob.

call? Thus, our choice of haptic device could even create additional possibilities for the turn-taking protocol.

3.3 Prototyping Haptic Icons to Support the Protocol

As stated earlier, in our *urgency-based* protocol a user is always *in control*, *waiting for control*, or simply *observing*. There are six possible states, as shown earlier in Figure 3-1: one for observing, two while waiting for control, and three while in control. The two *waiting for control* states correspond to a user waiting to obtain control following a gentle or urgent request. The three *in control* states represent a user currently in control with no one requesting control, one collaborator gently requesting control, and one collaborator urgently requesting control (or multiple collaborators requesting control) respectively.

We decided not to provide any haptic feedback in the observing state, reasoning that users would find the mapping between no interaction and receiving no haptic feedback

intuitive. In the other five states, we provide periodic haptic feedback so that a user can ascertain their status simply by placing a hand on the haptic mouse. We also provide transitory haptic feedback to users when they gain and lose control. In total, we needed a set of seven haptic icons spanning three families: three representing the *in control* states, two representing the *waiting for control* states, and two representing *changes in control*.

Rather than simply creating seven stimuli and arbitrarily assigning them to meanings, we required stimuli within a family to be perceptually similar, yet distinguishable from one another, and different families to have distinctly different sensations. Instead of creating completely abstract associations between stimuli and meanings, we began by using some common metaphors. Our inspiration for the *waiting for control* icons was a person impatiently waiting in line, tapping or drumming his or her fingers on the counter. Thus, when a user makes a gentle request for control, he or she feels a single, periodic pulse, confirming that he or she is waiting for control. When a user makes an urgent request for control, he or she feels two narrowly spaced pulses, again repeated periodically. For the *change in control* icons, we created haptic equivalents of the two-tone sound played when a PCMCIA card is inserted to and removed from a Windows laptop. The gained control icon consisted of a short, moderate vibration, followed immediately by a longer, strong vibration. The lost control icon was the exact opposite. Our motivation for the *in control* icons was a heartbeat metaphor; as a person becomes more anxious, their heart beats harder and faster. Thus, when no one has requested control, the user in control receives a subtle vibration, but as gentle and urgent requests are made, the intensity of the feedback increases.

While preliminary prototyping yielded a possible set of icons, we needed to ensure that distinctions between the icon families were clear and that icons within a family were also mutually distinguishable. In the next chapter, we describe the methodology used to match haptic stimuli with each of the icons in our protocol.

Study 1: Selecting Haptic Icons for the Urgency-Based Turn-Taking Protocol

The design of Study 1 was based on work by MacLean and Enriquez in their investigation of the perceptual design of haptic icons [46]. In their study, they sought to identify the parameters humans use to categorize haptic stimuli delivered through a knob. They found that subjects categorized first on the frequency of the stimuli, next on the waveform or shape of the haptic stimuli, and finally the magnitude.

The primary goals of our study were to ensure that the *change in control* icons were mutually distinguishable, yet related to one another, and to select a set of *in control* icons, again distinguishable yet related, and distinctly different from the *change in control* icons. We did not include the two *waiting for control* icons, as we were quite certain the pulse-based stimuli would be perceived as very different from the vibration-based stimuli we evaluated. Besides this, we also wanted to see how subjects would categorize stimuli using a vibrotactile display, and whether the parameters would be similar to MacLean and Enriquez's results. Finally, we wanted to measure how noticeable and pleasant the various stimuli felt.

In designing the haptic icons, we wanted certain icons to be more noticeable or intrusive than others, so that they would draw a user's attention quickly. At the same time, we knew that prolonged exposure to intrusive icons would annoy users. For example, the *waiting for control* icons remind a user of an action they made themselves, so they should not be intrusive. Similarly, when a user is in control and no one is requesting control, the feedback provided should be quite subtle. However, we wanted to ensure that the user in control noticed requests for control, particularly urgent requests. We also felt that the *change in control* icons should be quite noticeable. With respect to the pleasantness of icons, we felt

that icons designed to be subtle should also be pleasant, and that it would be acceptable if more noticeable icons were somewhat less pleasant.

4.1 Multidimensional Scaling

Multidimensional scaling (MDS) is a method for identifying relationships in data. It plots data points in an m -dimensional, typically Euclidean space such that m is small and points that are near one another are considered similar to one another. Two- or three-dimensional solutions are common. This enables investigators to identify interesting features such as tight clusters of points or outliers quickly, something that would be much more difficult to do by looking at the raw data. This method can be used to confirm hypotheses about relationships, or it can be used to discover the relationships that exist in a set of data.

MDS algorithms take as input an $n \times n$ dissimilarity matrix that contains information on an n item set, namely the “distance” or “difference” between each item and the rest of the items in a set. An example would be a table on a road map that shows the driving distances between a set of cities. Depending on the algorithm used, the differences can be expressed in ratio, interval, or ordinal units. With this information, an n -dimensional graph could easily be constructed, but would be of limited use for visual examination unless n is quite small. MDS algorithms iterate to produce an m -dimensional matrix, where $m < n$ and m is defined by the user, where the distance information is preserved as much as possible. This data can then be graphed in an m -dimensional space. A goodness-of-fit test called *stress* indicates the degree to which the new data corresponds to the input set. Lower stress values indicate better fit and are desirable. The closer m is to n , the lower the stress but also the lower the benefits of this approach. When the stress values for several dimensions are graphed, an “elbow” in the line can typically be observed. Researchers typically use the dimension at which the elbow occurs to perform their analysis, as the increase in accuracy with higher dimensions is outweighed by the increase in difficulty of performing the analysis. Formal descriptions of MDS can be found in Young and Hamer [65], and Green, Carmone, and Smith [32].

MDS often uses subjective ratings of a set of items as input. The typical approach for generating the initial $n \times n$ matrix has subjects rate the similarity of pairs of items on a Likert scale. All possible pair-wise comparisons of items ($n(n - 1) / 2$) are rated by subjects and input into an $n \times n$ similarity matrix, where similar items have a high score. This matrix is then converted into a dissimilarity matrix, such that similar items have a low score (just as the

distance between neighboring cities is lower), and processed by the appropriate MDS algorithm.

Ward proposed an alternate method of gathering similarity data [64]. Ward's experiments dealt with discovering the salient parts of photographs. Rather than having subjects perform pair-wise comparisons, subjects were presented with all the photographs at once, and were instructed to sort them into up to 20 categories, using any criteria they wanted. This was repeated an additional four times, except that in these trials, the subjects had to sort the photographs into a fixed number of categories. By the end, subjects had sorted the photographs into four out of 3, 6, 9, 12, and 15 categories. The category skipped was the one that most closely matched the number of categories the subject used in the initial sort.

Once the categorizations were complete, similarity scores were calculated by summing the number of categories used each time a pair of stimuli appeared in the same category. Thus, if a pair of stimuli appeared in the same category in the 9 and 12 category sorts, it would have a similarity score of 21. Similarly, if a second pair of stimuli appeared in the same category in the 3 and 9 category sorts, it would have a score of 12. The first pair would be considered more similar than the second, since its score is higher.

There are several advantages to Ward's methodology. By allowing the subjects to place items into categories, the number of comparisons required is greatly reduced, saving time and likely improving the consistency of subjects' ratings over time. The categorization technique also enables subjects to compare an item to the other items in the category, again improving the likelihood of consistency. We used this technique in our study.

4.2 Method

Subjects were asked to categorize 26 haptic stimuli. This included the two *change in control* icons, and a set of 24 possible candidates for the three *in control* icons. The *gained control* icon consisted of a 100 Hz, 3000-magnitude vibration followed by 200 Hz, 8000-magnitude vibration; the *lost control* icon was the mirror opposite. The 24 *in control* candidates varied on three parameters: frequency, magnitude, and rhythm. Frequencies of 21, 59, and 100 Hz were used, as stimuli below 20 Hz did not have a sufficient range of magnitudes, and stimuli over 100 Hz produced confounding auditory noise. Four different magnitude levels were used: 500, 2000, 5000, and 8000, resulting in twelve combinations of stimuli. To this, we introduced a temporal variable, whereby the stimuli were played either in

a single 1000 ms burst, or in two 700 ms bursts separated by a 100 ms delay. The 26 stimuli are listed in Table 4-1.

We modified a Visual Basic application written by Enriquez and used in [46]. The application presented a user interface (shown in Figure 4-1) where the 26 stimuli were represented by small tiles and grouped at the bottom of the window. Subjects could play back

Stimulus	Frequency	Attack	Magnitude	Bursts
v2	59 Hz	0	500	1
v3		0	500	2
v4		4 000	2 000	1
v5		4 000	2 000	2
v6		7 000	5 000	1
v7		7 000	5 000	2
v8		10 000	8 000	1
v9		10 000	8 000	2
v11	21 Hz	0	500	1
v12		0	500	2
v13		4 000	2 000	1
v14		4 000	2 000	2
v15		7 000	5 000	1
v16		7 000	5 000	2
v17		10 000	8 000	1
v18		10 000	8 000	2
v19	100 Hz	0	500	1
v20		0	500	2
v21		4 000	2 000	1
v22		4 000	2 000	2
v23		7 000	5 000	1
v24		7 000	5 000	2
v25		10 000	8 000	1
v26		10 000	8 000	2

Table 4-1 - Haptic Stimuli evaluated in Study 1.

Single burst stimuli were played for 1000 ms. Each burst in the two-burst stimuli was played for 700 ms, separated by a 100 ms delay. Attack values were used to strengthen the initial sensation of all except for the weakest stimuli to make them more noticeable.

Stimulus v1 was the *gained control* stimulus: it consisted of a 100 Hz stimulus lasting for 250 ms, followed by a 50 ms delay, followed by a 200 Hz stimulus lasting for 250 ms. The first stimulus had an attack value of 5000 and a magnitude of 3000. The second stimulus had an attack value of 10 000 and a magnitude of 8 000. Stimulus v10 was the *lost control* stimulus, and was the exact opposite of the *gained control* stimulus

the stimuli by pressing on the tile. Subjects then sorted the stimuli into categories five times, based on Ward's method described above. On the initial sort, subjects could sort the stimuli into a minimum of 2 and a maximum of 15 categories.

Stimuli could be moved between categories as necessary. To mitigate learning effects, the tiles were labeled with random numbers from 1 – 26 and positioned randomly in

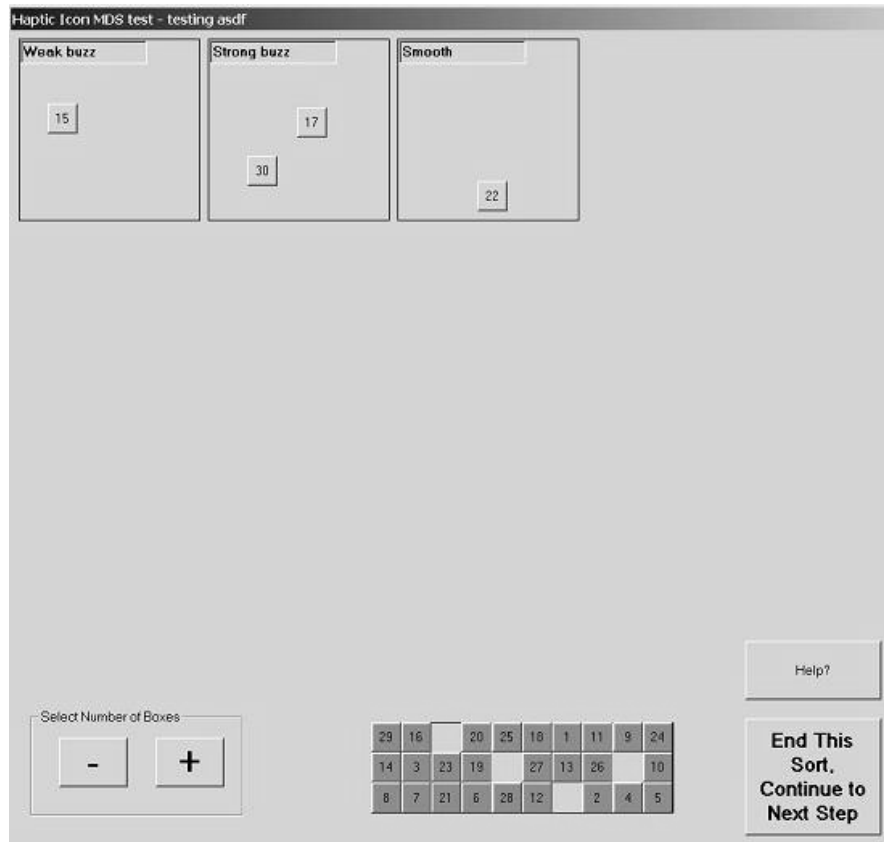


Figure 4-1 - Visual Basic application for sorting haptic stimuli. Here, three categories have been created and four stimuli placed in them. The subject has labeled the categories “Weak buzz,” “Strong buzz,” and “Smooth.”

the grid; the labels and positions were changed each of the five trials. Subjects were also asked to label each category with a descriptive name. After the study, subjects were presented with a Java application that allowed them to review each of the stimuli and rate them in terms of how noticeable and how pleasant they felt using a 5-point Likert scale.

The study was conducted in the experiment room of the Imager Graphics, HCI and Visualization Lab at the University of British Columbia. The *iFeel* mouse rested on a thick mouse pad, providing vibrational damping to improve the quality of the haptic feedback and

to minimize confounding audio noise. The study software ran on a Pentium IV 2.67 GHz computer with 512 MB of RAM. A 17" NEC display at a resolution of 1280x1024 was used. To mask audible noise from the *iFeel*, subjects wore Bose noise-canceling headphones, and listened to recorded white noise.

4.3 Analysis

Subject data were analyzed using SPSS 11.5 [6] using two variations of the ALSCAL MDS algorithm [60]: the Euclidean distance model and the Individual differences Euclidean distance model (also known as INDSCAL). The Euclidean distance model averages matrices and performs an analysis on a single matrix, whereas the INDSCAL model considers the importance of each dimension to each subject separately [65]. The differences in our solutions were relatively minor, but results from the INDSCAL algorithm are shown here.

To convert the subject data into the format required by the MDS algorithms, similarity scores were first calculated again based on Ward's methodology. The maximum score² was $3+6+9+12+15 = 45$. The scores were then converted to dissimilarity scores using the formula: $\text{Dissimilarity Score} = 1000 - (1000/45 * \text{Similarity Score})$.

Thus, if a pair of stimuli appeared in the same category in all sorts, it would have a similarity score of 45 and a dissimilarity score of 0. Conversely, if a pair of stimuli never appeared together, it would have a similarity score of 0 and a dissimilarity score of 1000.

4.4 Results

10 subjects (6 male, 4 female) were recruited to participate in the study, ranging in age from 21 to 31 years old. All had normal tactile sensitivity; 6 had no or little prior exposure to haptics and 4 were expert users of haptic devices. Three users were left-handed and 7 right-handed, but all used their right hand to control the mouse. The subjects were

² In fact, the maximum score depends on the first sort by the subject. If the subject sorts the stimuli into 4, 7, 10, or 13 categories, the maximum score is 46. If the subject sorts the stimuli into 2, 5, 8, 11, or 14 categories, the maximum score is 44. This affects the dissimilarity score calculated, but we found that differences in the resulting MDS graphs were minimal. Ward's method used the maximum possible score (in this case, 46). He also assumed that the minimum similarity score would be 1, not 0, since each item is at least part of the set of items being presented. To our knowledge, the process used by Ward has never been formally justified.

undergraduate and graduate students at the University of British Columbia and were paid \$10 for one hour's participation.

We first generated an MDS graph based all subjects' data for an overview of their categorization. Then, we generated graphs for subsets of the subjects, partitioning the data by gender, haptics experience, handedness, and outlier removal. In each case, we found that the elbow in the stress values occurred at the 3D MDS solution, yielding the best tradeoff between accuracy and interpretability. After this, we examined subjects' Likert scale responses regarding the noticeability and pleasantness of the stimuli. Based on our analyses, we selected the stimuli to be used in the *in control* states and confirmed the effectiveness of our *change in control* icons.

4.4.1 Analysis of MDS Graphs

Figures 4-2 to 4-5 show several views of the 3D MDS graph from all subjects. While it is difficult to interpret printouts of the graphs, SPSS has a "spin mode" that allows users to rotate the graphs at interactive rates. We relied heavily on this feature when analyzing the different graphs. As previously stated, we partitioned our data into several groups, and generated an MDS graph for each group. The groups were as follows:

- G1. Overall: all 10 subjects.
- G2. Male: the 6 male subjects.
- G3. Female: the 4 female subjects.
- G4. Left-handed: the 3 left-handed subjects (all used their right hand to control the mouse).
- G5. Right-handed: the 7 right-handed subjects.
- G6. Novices: the 6 subjects with little or no prior exposure to haptic devices – at most, using a vibrating cellular telephone or occasional use of a game-pad with a vibrotactile display.
- G7. Experts: the 4 subjects with substantial haptics experience, such as extensive use of vibrotactile game-pads or force-feedback devices like the PHANTOM [5].
- G8. Weird removed: the INDSCAL algorithm calculates a *weirdness* index that shows how each subject's weighting of each dimension in the solution differs from the average weighting. Higher weirdness values indicate a greater difference. We

removed the 3 subjects with the highest weirdness values and ran the analysis on the remaining 7 to see how the resulting graph would differ from using all the subjects. Since the weirdness index does not specify which ratings for pairs of stimuli cause a subject to be considered weird, we excluded all data for the 3 subjects.

The partitioning of subjects into groups is summarized in Table 4-2. It should be noted that all female subjects were novice haptic users, and all the expert haptics users were male. This was acceptable for our purposes, as our main intent was not to look for gender or experience differences in our subjects, but should be considered when interpreting the results.

We examined the MDS graph for each group to see how subjects clustered the 26 stimuli; the graphs are shown in Appendix G. In particular, we tried to identify the parameters subjects used to cluster stimuli, whether based on frequency, magnitude, number of bursts, or other criteria. We also report common features across groups that we noticed when exploring the data. Our observations have been summarized in Table 4-3 and are defined as follows:

- O1. Number of clusters: How many clusters of stimuli were identified. A cluster was loosely defined as two or more stimuli near one another. For example, in the analysis shown in Figures 4-2 to 4-5, the stimuli labeled v6, v8, v13, v15, and v17 were considered to be in a cluster.
- O2. Number of isolated stimuli: How many stimuli were observed that were not part of an obvious cluster.
- O3. Quality: How tightly packed each cluster was in a graph, relative to the clustering of the Overall graph (we use Average to denote the packing in the Overall graph).
- O4. Single / Double: When single-burst stimuli appeared in different clusters than double-burst stimuli, as opposed to single- and double-burst stimuli appearing in the same cluster.
- O5. Gained / Lost Control: When the gained control and lost control stimuli were clustered together, yet distinct from one another.




Figure 4-2 - MDS Graph for all 10 subjects, perspective projection.




Figure 4-3 - MDS Graph for all 10 subjects, looking down Dimension 1.




Figure 4-4 - MDS Graph for all 10 subjects, looking down Dimension 2.




Figure 4-5 - MDS Graph for all 10 subjects, looking down Dimension 3.

Subject	Male	Female	Left-handed	Right-handed	Novice	Expert	Weird Removed
1		•	•		•		
2	•			•		•	•
3	•		•		•		
4	•			•	•		•
5		•		•	•		
6	•			•		•	
7	•			•		•	
8		•		•	•		•
9		•	•		•		
10	•			•		•	

Table 4-2 - Partitioning of Study 1 subjects into groups for MDS analysis.

- O6. Weak: When the 500-magnitude stimuli were clustered together, regardless of frequency.
- O7. Strong 21/59 Hz: When the 5000- and 8000-magnitude 21 Hz and 59 Hz stimuli were clustered together.
- O8. Strong 100 Hz: When the 2000-, 5000-, and 8000-magnitude 100 Hz stimuli were clustered together.
- O9. Frequency Split: When stimuli of similar magnitude but different frequencies were clustered together and sub-clusters based on the frequency of the stimuli could be identified.
- O10. Frequency / Magnitude: When 5000- and 8000-magnitude stimuli of a particular frequency were clustered with *all* intensities of a different frequency of stimuli.
- O11. Weak / Isolated: When the isolated stimulus / stimuli have a magnitude of 2000.

The first three observations provide a sense of the MDS graph produced in each analysis. The “ideal” number of clusters is not strictly defined. If an MDS graph only contains a few clusters, it is likely that a single dimension dominates the others. However, if there are too many clusters, it can be difficult to tell what dimensions are being used to categorize the stimuli. In our results, there was little variation in the number of clusters across different groups of subjects (5-8, with most either 6 or 7). The presence of isolated stimuli in the data suggests that subjects did not agree on how to categorize certain stimuli. A closer examination of the outliers revealed that all are 500-magnitude stimuli; it is possible that their low intensity makes their frequencies more difficult to ascertain. The quality of the clustering also reflects the degree of consensus among subjects. A tighter clustering indicates that subjects agree that a set of stimuli is related. As might be expected, the group of haptics experts and the group with weird subjects removed had the tightest clustering.

All groups clearly distinguished between single- and double-burst stimuli when categorizing them except for the female subjects. One possible explanation was that all the female subjects were novice haptics users, but this was disproved by the novice users graph (based on 4 females and 2 males), which showed a distinction between single- and double-burst stimuli. Similarly, in the left-handed subjects graph, where 2 out of 3 subjects were female, the same distinction appeared. It is possible that the female subject who was removed

		Group							
		G1: Overall (N = 10)	G2: Male (N = 6)	G3: Female (N = 4)	G4: Left- handed (N = 3)	G5: Right- handed (N = 7)	G6: Novices (N = 6)	G7: Experts (N = 4)	G8: Weird Removed (N = 7)
Observation	O1: Number of clusters	7	6	7	5	8	7	6	6
	O2: Number of isolated stimuli	0	2	0	0	0	1	0	0
	O3: Quality	Average	Looser	Looser	Tighter	Average	Looser	Tighter	Tighter
	O4: Single / Double	•	•		•	•	•	•	•
	O5: Gained / Lost Control	•	•	•	•	•	•	•	•
	O6: Weak Cluster	•	•	•		•	•		
	O7: Strong 21 / 59	•	•	•		•	•		•
	O8: Strong 100	•	•	•	•	•	•	•	•
	O9: Frequency Split	•				•	•		
	O10: Frequency / Magnitude		•		•			•	
	O11: Weak / Isolated		•				•		

Table 4-3 - Observations from analysis of MDS graphs (N = 10). Columns show different sub-groupings of the same set of subjects.

in the Weird Removed analysis caused the inconsistency, or that the single-burst values could not be placed accurately in the 3D graph solution.

The analyses confirmed that the *change in control* icons are related to one another, as groups consistently placed the icons in their own tight cluster.

While we expected either frequency or magnitude to be used to differentiate between stimuli, our results were inconclusive. On one hand, many groups clustered the 500-magnitude stimuli together regardless of frequency; groups also clustered the 5000- and 8000-magnitude 21 Hz and 59 Hz stimuli. For some groups, within a cluster of stimuli with similar magnitudes, sub-clusters categorized by frequency were present, suggesting magnitude might dominate over frequency. On the other hand, the 100 Hz stimuli were consistently placed together; in a few cases, all four magnitudes were clustered together. We also observed a few instances of a curious and inexplicable grouping of stimuli, namely 5000- and 8000-magnitude 59 Hz stimuli with all magnitudes of the 21 Hz stimuli, such that the 59 Hz stimuli were closest to the 500- and 2000- magnitude 21 Hz stimuli.

In their experiment, MacLean and Enriquez found that subjects categorized stimuli delivered through a knob first by waveform, then by frequency, and finally by magnitude. In our study, subjects primarily used the number of bursts to distinguish between stimuli, a parameter that was not present in their study. After this, frequency and magnitude were used equally to categorize. The weak, 500-magnitude stimuli were often clustered together, perhaps due to their lack of salience. On the other hand, all except for the weakest of the 100 Hz stimuli were clustered together consistently instead of being clustered with the stronger 21 and 59 Hz stimuli. This suggests the overall salience of the 100 Hz stimuli was greater than the others due to its higher frequency. Had the magnitude levels been equalized so that stimuli felt equally intense across frequencies, our results likely would have been different.

4.4.2 Likert Scale Responses

Subjects rated how noticeable the stimuli felt on a five-point Likert scale, where a 1 meant “barely noticeable” and a 5 meant “very noticeable.” The results for each stimulus are listed in Appendix G. The trend in responses was as we expected: for each frequency, the larger the magnitude level specified, the higher the noticeability ratings. However, the number of bursts did not appear to have an effect on the ratings. An interesting observation was that subjects rated the 500-magnitude, 100 Hz stimuli more noticeable than the

equivalent 21 Hz and 59 Hz stimuli (the former was rated by most as a 2, while the latter two were rated as a 1). This is likely related to the limitations of the *iFeel* and to the way humans perceive intensity, as discussed.

Subjects also rated how pleasant the stimuli felt on a five-point Likert scale, where a 1 meant “very unpleasant” and a 5 meant “very pleasant.” As we suspected, there was an inverse relationship between stimulus magnitude and pleasantness rating: across all frequencies, the lower the magnitude, the higher the pleasantness rating. While there was little variation in ratings for each frequency, the 59 Hz and 100 Hz stimuli received slightly higher pleasantness ratings than the 21 Hz stimuli overall.

The Likert scale responses were used to ensure we picked appropriate stimuli to use. The *change in control* stimuli were designed primarily to be noticeable; since they are transient icons, their pleasantness was less of a concern. Our requirements for the *in control* icons were different: we wanted the basic *in control* icon (where no one else is requesting control) to be non-intrusive and very pleasant. We felt the icon indicating someone had gently requested control should be more noticeable, but still pleasant. Finally, we required the icon indicating multiple requests or an urgent request for control to be quite noticeable and perhaps somewhat annoying. These requirements were taken into consideration when selecting the *in control* icons.

4.4.3 Selecting the In Control Icons

Recall there are three possible states when a user is in control: the user may be (1) in control with no outstanding requests for control, (2) in control with a gentle request from another user, or (3) in control with an urgent request for control from another user or multiple requests. Based on the MDS analysis, we chose the v2 stimulus (500-magnitude, 59 Hz, single-burst) to represent the first state, the v6 stimulus (5000-magnitude, 59 Hz, single-burst) to represent the second state, and the v24 stimulus (5000-magnitude, 100 Hz, double-burst) to represent the third state. Rather than selecting stimuli within a single cluster, we opted to be conservative and choose stimuli from different clusters. If either frequency or magnitude had been a dominant factor in categorizing stimuli, we would have had more confidence in selecting stimuli from a single cluster.

In keeping with our conservative approach, we also analyzed the MDS graphs of each subject to examine the suitability of the stimuli; one subject placed the v24 stimulus

close to the change in control stimuli and a different subject placed the v2 and v24 stimuli together. In the first case, we felt that it was appropriate for the stimuli to be somewhat related, as the v24 stimulus should encourage a user to release control. However, the second case was puzzling; it is possible that one of the stimuli was misplaced by the subject, or that one of the stimuli was not placed well by the MDS algorithm.

The Likert scale ratings for each of the stimuli satisfied our requirements. On the 5-point noticeability scale, 8 out of 10 subjects rated the v2 stimulus as not very noticeable (1 or 2 on the Likert scale), and 8 subjects rated the v6 and v24 stimuli as quite noticeable (4 or 5). On the 5-point pleasantness scale, 7 subjects rated the v2 stimulus as quite pleasant (4 or 5), 9 subjects rated the v6 stimulus as somewhat pleasant (3 or 4), and 8 subjects rated the v24 stimulus as somewhat unpleasant (2 or 3). The responses for all stimuli are shown in Appendix H.

4.4.4 Confirming the Change in Control icons

In the MDS analysis, the *change in control* icons were consistently placed in a cluster together, supporting our desire for them to be related. In the Likert-scale responses for noticeability, 9 subjects rated the *gained control* stimulus as quite noticeable, giving it a 4 or 5. Six subjects gave the *lost control* stimulus a 4 or 5. In both cases, the remaining subjects rated the stimulus a 3. Since the stimuli were designed to be intrusive, this was a positive result. In terms of the pleasantness ratings, the *gained control* stimulus received a nearly even distribution of responses, while the *lost control* stimulus received more neutral or slightly favorable responses.

Unfortunately, we had no measure by which to judge whether the icons were different enough to be distinguishable from one another, since a MDS graph only shows the relative differences between stimuli. Indeed, when piloting the second study, subjects reported having to rely on the noise from the *iFeel* to distinguish the icons, forcing us to modify them. In retrospect, had we prototyped several variants of these icons just as we did for the *in control* icons, this problem might have been avoided.

4.5 Haptic Icons

Based on our initial prototyping and subsequent evaluation using MDS and subjective responses, we selected the stimuli shown in Table 4-4 to be used as our haptic

Family	State	Haptic Sensation
Change of Control	User has gained control of the shared application	0.25 s, 3000-magnitude, 100 Hz vibration, followed by a 0.05 s pause, followed by a 0.25 s, 8000-magnitude, 200 Hz vibration
	User has lost control of the shared application	0.25 s, 8000-magnitude, 200 Hz vibration, followed by a 0.05 s pause, followed by a 0.25 s, 3000-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 strongly requested control or multiple people have 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

Table 4-4 - Haptic icons selected after Study 1.

icons. The stimuli for the *waiting for control* icons were unchanged from the prototyping stage. As well, the stimuli for the *change in control* icons were also used as prototyped, since our analysis did not highlight any difficulties. Finally, we chose 3 stimuli from the 24 stimuli we evaluated for the *in control* icons.

4.6 Summary

In this chapter, we described how we optimized the set of haptic icons chosen to support our protocol using a technique based on Multidimensional Scaling. In particular, we were interested in selecting an appropriate set of *in control* icons and validating the *change in control* icons we had prototyped. We conducted a user study where subjects sorted a set of 26 stimuli into different numbers of categories, and rated the stimuli on their noticeability and pleasantness.

Based on the Study 1 results, we selected a set of haptic icons to support our turn-taking protocol. We next had to ensure that subjects could learn to identify the icons without extensive training. As well, in our collaborative system, subjects would have to be able to identify the icons while actively working on a primary task. In the next chapter we describe Study 2, where we address these issues.

Study 2: Learning and Using Haptic Icons in the Presence of Workload

In Study 1, we selected three families of haptic icons to represent the different states in our urgency-based turn-taking protocol. We were reasonably certain that each family would generally be perceived as distinct from the others, and that icons within a family would be perceived as distinct from one another. However, we neither knew how easily users would be able to learn the meanings associated with the stimuli, nor whether they could accurately and rapidly recall the meanings while engaged in other tasks. If extensive training was required to learn the icons, users likely would be reluctant to expend the effort. Furthermore, if users struggled to recall the meanings of the icons, the usefulness of this approach would be minimal.

In this chapter, we describe our second study, which we designed to address these questions. Although our intended use of the haptic icons is in a collaborative environment, we chose to evaluate single-user behavior and performance in Study 2. We begin by discussing the experiment, which consisted of a learning phase and an evaluation phase. Then, we list the measures we used to collect data and the research questions we addressed in the study. Following this, we present the results from the study and discuss their implications.

5.1 Experiment Procedure

The study was divided into a learning phase and an evaluation phase, both completed by subjects in a single 1.5 hour session. The purpose of the learning phase was to measure how quickly subjects could learn the 7 haptic icons planned for our *urgency-based* protocol to 90% accuracy. The evaluation phase was designed to measure subjects' ability to notice

changes in the haptic icons delivered and to identify them, with this taking place under different amounts of cognitive workload.

The study setup was nearly identical to Study 1. The study software, consisting of a multithreaded Java application, ran on Pentium IV 2.67 GHz computers with 512 MB of RAM, running Windows XP Professional. The displays used were 17" LCD monitors at a resolution of 1280 x 1024. Subjects again wore Bose QuietComfort2 noise-canceling headphones and listened to white noise to mask noise from the *iFeel*. Sessions were automated; to avoid subtle strategic bias from variations in instruction delivery [29], subjects read instructions on-screen and in a booklet provided at the beginning of the session.

The icons used in the study are shown in Table 5-1. They are nearly identical to the Study 1 stimuli, with minor changes to the *change in control* family. The changes were necessary because pilot subjects reported using the sound from the *iFeel* to distinguish between the *gained control* and *lost control* icons. As well, the meanings associated with the stimuli were changed, as we felt that learning our intended set of meanings would require an elaborate explanation. The labels we used in the study are shown in the last column of Table 5-1; they correspond to different states a person may experience during the day, and preserve

Family	ID	State	Haptic Sensation	Study 2 Label
Change in Control	CH+	User has gained control of the shared application	0.4 s, 1000-magnitude, 100 Hz vibration, followed by a 0.2 s, 8000-magnitude, 100 Hz vibration	Awake
	CH-	User has lost control of the shared application	0.2 s, 8000-magnitude, 100 Hz vibration, followed by a 0.4 s, 1000-magnitude, 100 Hz vibration	Asleep
In Control	IN	User is in control of the shared application	1 s, 500-magnitude, 60 Hz vibration; 1 s delay between iterations	Low Stress
	IN+	User is in control, but someone has gently requested control	1 s, 5000-magnitude, 60 Hz vibration; 1 s delay between iterations	Medium Stress
	IN++	User is in control, but someone has strongly 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	High Stress
Waiting for Control	WAIT	User has gently requested control	Single pulse; 1 s delay between iterations	Bored
	WAIT+	User has strongly requested control	Two pulses, separated by a 0.15 s pause; 1 s delay between iterations	Really Bored

Table 5-1 – Haptic icon set used in Study 2.

the original family relationships.

5.1.1 Learning Phase

During the learning phase of the study, subjects were instructed to learn the meanings associated with the 7 haptic stimuli “as quickly as possible”. To proceed to the evaluation phase, subjects had to score over 90% on a test. Subjects were first presented with an application that allowed them to play back the 7 icons as many times as they wanted in any order; they chose without penalty when to proceed to the test. A screenshot of the application is shown in Figure 5-1. The icons were arranged by family to facilitate hierarchical learning; subjects clicked on the button beside the icon to play its associated stimulus.

In the learning test, subjects felt a haptic icon once and identified it by selecting the correspondingly labeled radio button (Figure 5-2). Each icon was presented three times for a total of 21 trials, randomized with the constraint that the same icon was never presented twice in a row. To prevent positional memorization, the labeled radio buttons were randomly re-

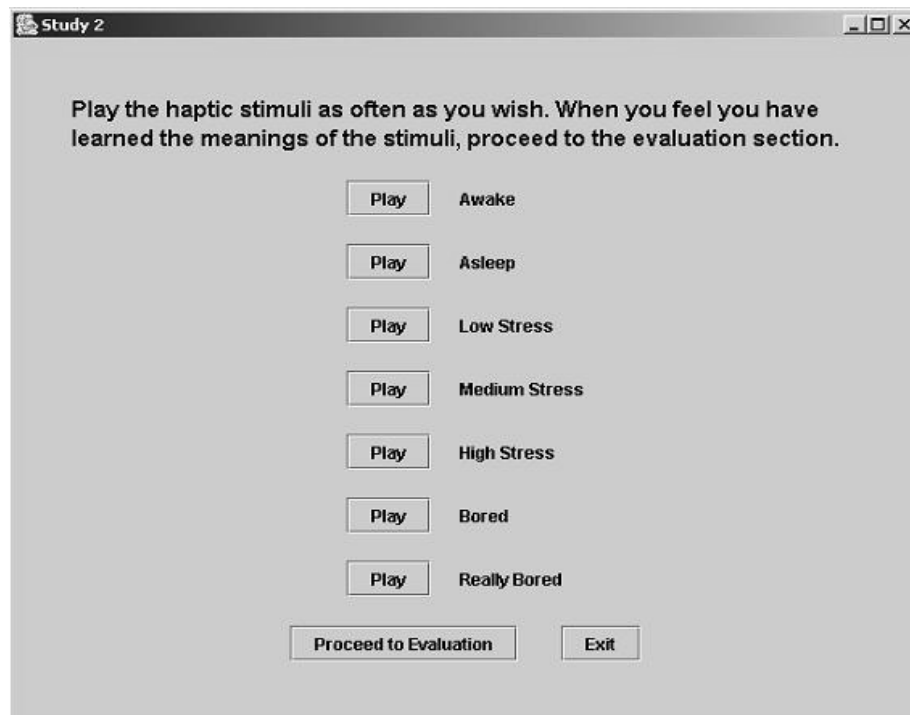


Figure 5-1 - Screen for exploring haptic icons in Study 2.

Study 2

Select which of the haptic stimuli you just felt, then press the "Next" button

Trial 1 of 21

☐ High Stress

☐ Asleep

☐ Low Stress

☐ Medium Stress

☐ Really Bored

☐ Bored

☐ Awake

Next

Figure 5-2 - Study 2 learning test.

ordered on each trial. As well, subjects were only told whether they had passed or failed the test, without any other specific performance feedback. When subjects correctly identified 19 or more icons, they proceeded to the evaluation phase; otherwise, they returned to the initial screen for more practice, before repeating the test.

5.1.2 Evaluation Phase

During the evaluation phase, subjects' ability to recall the meanings they learned in the learning phase was tested under three increasingly difficult conditions: haptic, haptic+visual, and haptic+visual+auditory, where "visual" and "auditory" represent the addition of workload tasks. Since our collaborative system is intended for use on a visual task, we did not include a haptic+auditory only condition in this study. The order of the conditions was counterbalanced across subjects. On average, trials for the three conditions were completed in 11, 12, and 13 minutes, respectively.

In the haptic condition, icons were presented in pairs. The transition from the first icon to the second occurred after a randomly chosen delay of 10, 15, or 20 seconds; non-periodic icons (CH-, CH+, as shown in Table 5-1) were repeated every 2 seconds. Subjects

were instructed to press the space bar as soon as they noticed the change. Although not specifically instructed to do so, all subjects used their non-mouse hand to press the space bar. Following the key press, a modal dialog box appeared that listed the 7 icons, again grouped by family. Subjects identified the second icon in the pair by selecting a radio button and pressing an OK button; they then proceeded to the next pair. If the subject had not pressed the space bar 10 seconds after the transition, this was counted as a “missed transition” and the dialog box was displayed, forcing the user to identify the second icon. As well, if the subject had not pressed the OK button on the dialog box 10 seconds after it was displayed (regardless of whether the subject missed the transition), this was counted as a “missed identification” and the haptic icon stopped playing. However, the user still had to select a radio button and press OK to proceed, based on their best guess. If the subject pressed the space bar before the transition occurred, this was counted as a “false alarm” and the subject was notified of his / her error.

Subjects responded to a total of 35 pairs, consisting of 5 transitions to each of the 7 icons; more transitions would have made the duration of the study unreasonable. We chose to use only the transitions that are possible in our turn-taking protocol, a subset of the 42 possible, to help us predict performance in Study 3. Transitions were presented in random order.

In the haptic+visual condition, subjects had to perform a visual task of solving a picture puzzle while performing the icon identification described for the haptic condition. An image was randomly selected from a set of 65 images, subdivided into a grid of 12 pieces, and the pieces were randomly rearranged. Subjects were instructed to rearrange the pieces to restore the original image, which was displayed beside the scrambled puzzle. A screenshot of this application is shown in Figure 5-3. A puzzle piece could be swapped with any other piece by dragging the piece on top of the other. When the subject had successfully solved a puzzle, a new puzzle was presented. The images were taken from the author’s personal photo collection and cropped to be roughly the same size. The same image was never repeated during a session.

In the haptic+visual+auditory condition, subjects had to listen for a keyword to be spoken while performing the tasks described in the haptic+visual condition. The keyword “blue” was spoken 30 times at random intervals interspersed with approximately 120 enunciations of 14 other colors in this condition, thus requiring subjects to attend to the audio



Figure 5-3 - Visual distracter task in Study 2 (image is one example out of 65 possible images). Subjects had to rearrange puzzle pieces on the left to match the image on the right.

stream. When subjects heard the keyword, they had to press the “b” key on the keyboard before the next color was spoken. Subjects had a minimum of 5 seconds to respond. All other presses were counted as misidentifications. Again, subjects pressed the key with their non-mouse hand without being explicitly directed to do so.

In each condition, subjects first practiced on seven pairs of icon transitions to familiarize themselves with the user interface for that condition and thereby mitigate learning effects. A random set of transitions was used. Subjects were also given an opportunity before each condition to review the 7 icons, using the same UI as in the learning phase of the experiment. This was done because pilot subjects reported becoming unsure over time as to whether they had associated the stimuli with their meanings correctly. This is probably because subjects never received reinforcement in icon identification in either the learning or the evaluation phase.

In summary, the evaluation phase was a 3 conditions x 7 icons x 5 transitions design, where all factors were within-subjects. The order of the 3 conditions was counterbalanced, and icon transitions were delivered randomly within each condition. Thus, subjects each completed a total of 105 trials.

5.2 Performance Metrics

We measured several aspects of subjects' performance, including:

- Time spent **learning** the associations between stimuli and their meanings.
- Time required to **detect** icon transition for each trial.
- Time required to further **identify** the second icon in the pair, once the transition had been detected.
- The number of **false alarms**, **missed transitions** and **missed identifications**.
- The number of **correctly identified** icons.
- The number of **visual puzzles solved** in the haptic+visual and haptic+visual+auditory conditions.
- The number of **audio keywords correctly** and **incorrectly** identified in the haptic+visual+auditory condition.

5.3 Hypotheses

Our hypotheses were as follows:

- *Detection Time Hypothesis:* Detection time for haptic icon transitions will increase with added workload.
- *Identification Time Hypothesis:* Identification time for the second icon in a pair will increase with added workload.
- *Correct Identification Hypothesis:* Number of correctly identified icons will decrease with added workload.
- *“Mistake” Hypothesis:* Number of false alarms, missed transitions, and missed identifications will increase with added workload.

As our hypotheses show, we expected performance to degrade as workload increased. While we didn't establish specific thresholds, we knew that if performance degraded substantially with increased workload, the utility of the haptic icons would be compromised. We expected detection times to be affected the most by workload, but hoped that icons designed to be intrusive (such as IN++, CH+ and CH-) would be affected less than icons

designed to be subtle (such as IN, WAIT, and WAIT+). As well, while we expected identification times to increase, we hoped that the changes would be minimal; stable identification times would suggest that subjects had internalized the icons, just as they might learn to recognize a physical item by its texture. With respect to the correct identification and mistake hypotheses, we again hoped that large changes would not occur, as misidentification or mistakes would likely be highly disruptive in our collaborative environment (such as a user suddenly releasing control when no one was requesting).

5.4 Results

Six males and 6 females participated in the study. Subjects ranged from 17 – 28 years old and were relatively naïve to haptic feedback; 5 subjects reported having no experience with haptic devices, while 7 occasionally used vibrating game controllers. Due to an oversight when screening subjects, one subject participated who had also participated in Study 1. His data was compared to the other subjects' data; it was subsequently used because he did not appear to be an outlier. Subjects were paid \$10 for a 1.5 hour session. To encourage brisk execution, subjects were informed that the four subjects with the best overall performance would receive an additional \$10. To avoid biasing any one task, instructions explicitly directed subjects to pay equal attention to the haptic, visual, and auditory tasks in order to maximize their "score".

A series of repeated-measures ANOVAs with an alpha level of 0.05 was run. When the data failed Mauchly's Test of Sphericity, the Huynh-Feldt correction was applied, reducing the degrees of freedom in several *F*-tests. In keeping with the exploratory nature of this work, we conducted post-hoc pair-wise comparisons liberally, but for protection used a Bonferroni adjustment, also at a 0.05 alpha level.

5.4.1 Learning Time

The learning time was measured as the amount of time subjects spent exploring the haptic icons using the GUI shown in Figure 5-1. The total time spent exploring the icons and taking the learning test was not used, as this would unfairly penalize subjects who attempted the learning test multiple times. Subjects spent between 56 and 446 seconds playing back the haptic icons (mean 177 seconds, standard deviation 114 seconds).




Figure 5-4 - Mean detection times for each condition (ms).

5.4.2 Detection Time Hypothesis

Detection time was calculated from the time the second icon in a pair began playing to the time when the subject pressed the space bar. If subjects missed the transition, the detection time was set at 10 seconds. Our statistical analysis yielded a main effect of condition, a main effect of icon, and an interaction effect between icon and condition. Figure 5-4 provides an overview of the detection time results.

As we hypothesized, the condition had a significant impact on the detection time ($F_{1,297, 14,270} = 20.359, p < 0.001$, partial $\eta^2 = 0.649$). The detection times for each condition are shown in Table 5-2. Mean detection time in the haptic+visual condition was nearly double that of the haptic condition, and the haptic+visual+auditory mean detection time was 22% longer than the time in the haptic+visual condition. Both pair-wise comparisons were significant.

Condition	Mean	95% Confidence Interval	
		Lower Bound	Upper Bound
haptic	1815	1186	2444
haptic+visual	3507	2421	4594
haptic+visual+auditory	4269	2998	5540

Table 5-2 - Mean detection times for each condition (ms).

The significant interaction between condition and icon indicates that the detection times of some icons were more sensitive to condition than others ($F_{7.940, 87.342} = 4.472$, $p < 0.001$, partial $\eta^2 = 0.289$). We compared detection times for each icon across different condition pairs; the differences are summarized in Table 5-3. As shown in the last column of that table, detection times for all except IN++ in the highest-workload condition (haptic+visual+auditory) were significantly greater than in the lowest-workload condition (haptic). With respect to the comparisons between the two other condition pairs, the results were not as strong. Only three icons requiring a longer detection time in the haptic+visual condition as compared to the haptic condition, and one icon in the case of the haptic+visual+auditory compared to the haptic+visual.

Looking more closely at specific icons, the IN++ icon was designed to be the most intrusive of our icons, and it is therefore not surprising that there was no difference in detection times across conditions. By contrast, the *waiting for control* icons were designed to be the least intrusive, as they confirm a user's actions, rather than conveying the intentions of others. It is therefore not surprising that in two out of the three condition pairs there were differences in detection times. In other words, we found that increasing workload impacted the detection of a nonintrusive icon, but did not impact that of an intrusive one. However, the results for the *change in control* icons were counter to our expectations. Both were intended to be intrusive but the detection time analysis revealed that they behaved more like the nonintrusive icons.

The presence of an interaction effect means that main effects should be treated with caution, as they may be due to the interaction effect. From Figure 5-4, it seemed likely that the interaction effect was caused by the WAIT and WAIT+ stimuli, whose detection times

Icon	Condition Pairs		
	h vs. h+v	h+v vs. h+v+a	h vs. h+v+a
CH+	0.884	0.010*	0.015*
CH-	0.022*	0.435	0.039*
IN	0.128	0.218	0.020*
IN+	0.338	0.168	0.028*
IN++	0.797	0.743	0.452
WAIT	< 0.001*	1.000	< 0.001*
WAIT+	0.002*	1.000	0.001*

Table 5-3 - p -values for differences in detection times across conditions. Items with an asterisk (*) are significant.

were substantially higher in the haptic+visual and haptic+visual+auditory condition than the haptic condition. To test this observation, we re-ran our analysis excluding the WAIT and WAIT+ stimuli and found that there was no interaction effect. However, the main effects of condition and icon remained, suggesting that the main effects observed in our original analysis were not simply a result of the interaction effect.

5.4.3 Identification Time Hypothesis

Identification time was calculated from the appearance of the modal dialog box listing the 7 icons (whether the subject had detected the change or missed it) to the subject pressing the OK button. Figure 5-5 shows a graph of the mean identification times in each condition for each icon. Table 5-4 shows the mean identification times for each condition. The data revealed a marginally significant main effect of condition suggesting that the condition impacted identification times as we hypothesized ($F_{2, 22} = 3.175, p = 0.061$, partial $\eta^2 = 0.224$). However, post-hoc comparisons did not confirm that any condition supported significantly faster identification than the others.

There was also a significant main effect of icon, indicating that some icons took longer to identify than others ($F_{6, 66} = 20.993, p < 0.001$, partial $\eta^2 = 0.656$). Comparisons revealed that identification of the *change in control* icons, and in particular of CH-, took

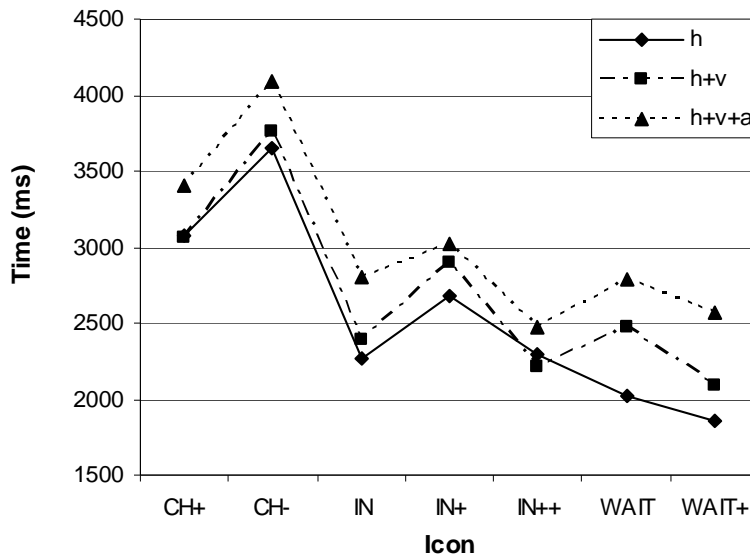


Figure 5-5 - Mean identification times for each condition (ms).

Condition	Mean	95% Confidence Interval	
		Lower Bound	Upper Bound
haptic	2548	2260	2836
haptic+visual	2698	2293	3103
haptic+visual+auditory	3022	2769	3276

Table 5-4 - Mean identification times for each condition (ms).

significantly longer than the others, suggesting that subjects found CH- the most difficult to identify. This was confirmed by our data; CH- was mistaken for CH+ or IN+ four times more often than those icons were mistaken for CH-.

We unexpectedly found a significant main effect of trial ($F_{4, 44} = 3.325, p = 0.018$, partial $\eta^2 = 0.232$). Recall that five transitions were made to each of the icons. Comparisons showed that identification times for the fifth transition were 13% faster than for the first transition. This indicated that despite our practice transitions in each condition, subjects were still learning and improving as the study progressed.

5.4.4 Correct Identification Hypothesis

Contrary to our expectations, condition did not significantly impact the rate of correct identification. On average, subjects identified icons correctly 95% of the time in all three conditions. A significant effect of haptic icon was found ($F_{2,395, 26,342} = 3.384, p = 0.042$, partial $\eta^2 = 0.235$) but none of the pair-wise comparisons were significant. To probe this result, we examined the subject data to see where mistakes occurred. Half of the mistakes involved the *change in control* stimuli: mistaking CH+ for CH- and vice versa; and mistaking CH- for IN+ and vice versa. The IN stimulus was also sometimes mistaken for the IN+ stimulus. While not significant, this trend along with our other observations about the *change in control* icons indicates that subjects struggled with them.

5.4.5 “Mistake” Hypothesis

In each condition, we measured the number of times subjects pressed the space bar before the haptic icon transition, the number of times subjects failed to press the space bar within 10 seconds of a haptic icon transition, and the number of times subjects failed to identify the haptic icon within 10 seconds of the selection dialog appearing. We found a

significant effect of condition for the first two measures, but not for the third, where only one instance of a missed identification occurred.

There are several possible reasons for false alarms to occur: a subject may have been certain that a transition occurred when it hadn't, a subject may have been uncertain as to whether a transition had occurred and subsequently decided to check, or a subject may have accidentally pressed the space bar instead of performing another action. For example, one subject reported pressing the space bar instead of the 'b' key repeatedly in the haptic+visual+auditory condition to identify an audio keyword. Nonetheless, the occurrence of false alarms can point to the effect of workload on subjects. A significant effect of condition was found ($F_{2, 22} = 12.815$, $p < 0.001$, partial $\eta^2 = 0.538$). Table 5-5 shows the number of false alarms in each condition; the number of false alarms in the haptic+visual+auditory condition was significantly greater than the haptic condition ($p = 0.004$) and the haptic+visual condition ($p = 0.021$).

Condition also had a significant impact on the number of missed transitions as we hypothesized ($F_{2, 22} = 13.822$, $p < 0.001$, partial $\eta^2 = 0.557$). Table 5-6 summarizes the percentage of missed transitions in each condition. All of the pairwise comparisons were significant.

The results also revealed a significant interaction between condition and icon, indicating that transitions to some icons were missed more in some conditions than in other conditions ($F_{12, 132} = 3.402$, $p < 0.001$, partial $\eta^2 = 0.236$). Post-hoc comparisons revealed a difference between the haptic and the haptic+visual+auditory condition for the IN icon. There were also differences between the haptic and the haptic+visual conditions, as well as

Condition	Number of False Alarms	95% Confidence Interval	
		Lower Bound	Upper Bound
haptic	1.500	0.215	2.785
haptic+visual	4.417	1.456	7.378
haptic+visual+auditory	8.917	4.894	12.939

Table 5-5 - Number of false alarms in each condition.

Condition	% Missed	95% Confidence Interval	
		Lower Bound	Upper Bound
haptic	1.7	0.2	3.1
haptic+visual	10.5	4.4	16.6
haptic+visual+auditory	18.8	9.8	27.8

Table 5-6 - Percentage of missed transitions for each condition.

between the haptic and the haptic+visual+auditory conditions for both of the *waiting for control* icons. These results show that transitions to the three subtlest icons were often overlooked as workload increased.

As before, we noted that the *waiting for control* icons could be responsible for the interaction effect, and re-ran our analysis without them. No interaction effect was found, but again a main effect of condition was present.

5.4.6 Distracter Task Performance

Subject performance on the distracter tasks suggests that they did not simply focus on identifying haptic icons. Subjects placed between 224 and 535 puzzle pieces during the two conditions with the visual puzzle task, with an average of 387 pieces. There was no significant difference in the number of pieces placed in each condition. Given that the two conditions took a combined total of approximately 25 minutes to complete, this means that on average, subjects placed a puzzle piece every 3 – 7 seconds, with an average rate of one piece every 4 seconds. Large individual differences are to be expected, since the task involves spatial reasoning abilities, but the average rate and its consistency across conditions strongly suggest that they were highly engaged. Performance on the audio distracter task was also acceptable: subjects correctly identified between 13 and 30 out of 30 keywords, with an average of 27 identifications.

5.5 Discussion

Learning Times and Distracter Task Performance

The short learning times exceeded our expectations, particularly since subjects were not given any hints or strategies to use to learn the icons, and since the learning test did not inform subjects which icons they had misidentified. While the labels we gave the icons were not completely random, the associations between the labels and the haptic stimuli were still quite abstract.

Our results show that subjects were engaged in the distracter tasks. The challenge of completing picture puzzles seemed quite appealing; several subjects informally remarked that the experiment was fun. Thus, we believe that our results are a good indication of

performance on a collaborative task, where users would be engaged in a visual primary task and haptic feedback would provide turn-taking information.

Detection Times

The increase in icon detection times across conditions as overall workload increased was expected. Of more interest was the size of the change from condition to condition. The mean detection time in the haptic+visual+auditory condition was double that of the haptic condition, but at approximately 4.3 seconds, quite acceptable for our purposes.

It was particularly important for the *change of control* icons and the IN++ icon to be detected and identified quickly regardless of condition. This proved true for IN++, but CH+ was detected quickly and not identified quickly, and CH- was neither detected nor identified quickly. In post-study interviews, subjects reported having the most difficulty identifying the CH- icon, especially as compared to the CH+ icon and the IN+ icon. This was also clear in our analyses. We attribute the difficulty to the modifications we made to the *change of control* icons immediately before the study in response to our pilot subjects' reports that they used their sounds to identify them. The changes inadvertently introduced the side effect of making them less distinguishable. Had we re-piloted the study, it is likely we would have discovered this.

We were interested to find that the mean detection time in the haptic+visual+auditory condition was significantly greater than in the haptic+visual condition. The auditory task was specifically designed to be straightforward so as not to unduly overload the user. However, we observed that even an easy auditory task made a significant difference in the detection time and in the number of missed transitions and false alarms. Given that subjects would be conversing in our collaborative system, this might seem cause for concern. However, we designed the conditions in this study to be a conservative evaluation of our collaborative system. In this study, the haptic, visual, and auditory tasks were all unrelated; in our system, the visual and auditory channels would be used in concert to accomplish the collaborative task, and the haptic channel would mediate the turn-taking. We expect that the cognitive load associated with this combined use of the visual and auditory channels would be lower than using the channels independently.

Identification Times and Accuracy

Although identification times increased marginally across conditions, we were pleased to find a very high degree of accuracy in haptic icon identification, regardless of condition. It is possible the accuracy would have been even higher with a different set of *change in control* stimuli, as they accounted for half of the errors. At the same time, it is possible our results would have differed had we used a different approach to gather the identification data, such as verbal reports. A modal dialog box allowed us to measure the identification time more precisely than a verbal report, and it allowed us to force subjects to identify each icon. One consequence of this choice was that subjects did not have to attend to the distracter tasks when the dialog box was open. This could possibly assist them in their identification of the haptic icons, as they could focus on a single task. However, it is not certain that subjects would be able to identify the icons while working on distracter tasks in parallel; they might simply pause long enough to identify the icon before resuming their tasks.

5.6 Summary

In this study, we evaluated the ability of subjects to learn a set of seven haptic icons and recall their meanings under different levels of workload. The set included icons designed to be nonintrusive, icons designed to be intrusive, icons to notify a user of others' intentions, and icons designed to confirm a user's own actions. Our results were encouraging.

Despite shortcomings with the *change in control* icons introduced during the pilot of the experiment, subjects were able to recall the meanings of seven icons to 90% accuracy on a test after approximately three minutes. Without any distracter tasks, average detection time for a icon transition was 1.8 seconds; as the workload increased, detection time increased significantly, but was still acceptable at 4.3 seconds in the haptic+visual+auditory condition. Icons that were designed to be nonintrusive were affected more than icons designed to be relatively intrusive.

Surprisingly, accuracy remained constant regardless of workload, and identification times were not affected to the same extent as the detection times. Subjects also showed a reasonable ability to perform the haptic identification, visual puzzle and audio keyword tasks simultaneously, arguably a more difficult task than our intended collaborative environment

poses. With this knowledge, we proceeded to our final study, where groups of subjects used our protocol in a collaborative task.

Study 3: Evaluating the Urgency-Based Turn-Taking Protocol

Using the results from Study 1, we selected a set of haptic icons to support our urgency-based turn-taking protocol. In Study 2, we evaluated subjects' ability to learn the icons, as well as their ability to recall the icons' meanings under different levels of workload. We designed Study 3, described in this chapter, based on our positive results; it is an exploratory, observational user evaluation of our protocol. Groups of 4 users completed furniture-layout tasks using three different combinations of haptic and visual modalities.

6.1 Research Questions

Our goal in Study 3 was to address the following research questions:

- I. Can subjects learn the meanings associated with the haptic stimuli in a reasonable amount of time?
- II. How will collaborative style be impacted by the different conditions?
- III. How will equitability of control sharing be impacted by the different conditions?
- IV. Which modality (visual, haptic, or combined) will subjects prefer for interaction information and control?
- V. How will task performance be impacted by the different conditions, if at all?

6.2 Conditions

Our goal was to compare protocol mediation by the traditional visual modality (costly in attention and screen space) with the potentially less intrusive haptic channel. Thus, all three conditions in our study used our new turn-taking protocol and were designed as follows:

1. **Visual:** The visual condition shows a User Window and Button Bar. The User Window displays who is in control, who has gently and urgently requested control, and a list of the group members. The Button Bar allows users to request and release control. Both objects, shown in Figure 6-1, always float beside the shared application window.
2. **Haptic:** The haptic condition uses the haptic input described in Chapter 3 and the haptic icons shown in Table 6-1. Addressing the difficulties with the *change in control* icons identified in Study 2, we re-introduced a delay between the bursts and modified their length to make them more distinguishable. The User Window from the visual 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.
3. **Haptic + Visual:** This condition combines haptic input and feedback from the haptic condition with the User Window and Button Bar from the visual condition. Subjects can use either the Button Bar or haptic inputs to request and release control.



Figure 6-1 - User Window and Button Bar.

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 <i>Feel:</i> System “powering up” – weak buzz followed by a strong buzz
	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 <i>Feel:</i> System “powering down” – strong buzz followed by a weak buzz
In Control	User is in control of the shared application	1 s, 500-magnitude, 60 Hz vibration; 1 s delay between iterations <i>Feel:</i> A light heartbeat
	User is in control, but someone has gently requested control	1 s, 5000-magnitude, 60 Hz vibration; 1 s delay between iterations <i>Feel:</i> A stronger heartbeat
	User is in control, but someone has strongly 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 <i>Feel:</i> A very strong heartbeat
Waiting for Control	User has gently requested control	Single pulse; 1 s delay between iterations <i>Feel:</i> Tapping a single finger on a table
	User has strongly requested control	Two pulses, separated by a 0.15 s pause; 1 s delay between iterations <i>Feel:</i> Drumming two fingers on a table

Table 6-1 - Haptic icons used in Study 3.

To minimize confounds, the same information is available in each condition, in particular persistent information about who is in control or requesting control at a given urgency. In the visual condition this information is displayed continuously in the User Window, rather than through the transient tool-tips or dialog boxes used in current solutions. In the haptic condition we continuously transmit requests for control, and the user can invoke the User Window to identify who has made the requests.

We note that these conditions are not purely ‘visual’ or ‘haptic’, since both (haptic) motor activities and some visuals are necessarily involved in all cases; the conditions distinguish the primary source of information and control.

We could have made many other comparisons, for example, compared our protocol to the more common *give* and *take* protocols. However, this would have required an elaborate and lengthy study. We felt it was more important to first investigate the protocol itself, and compare its haptic and visual instantiations.

6.3 Study Setup

We modified an open-source view-sharing system called Virtual Network Computing (VNC) [7] to implement our protocol. VNC consists of client and sever applications; using the client, a user can control the desktop of a remote computer running the server. Unlike the web conferencing systems described in Chapter 2, VNC is primarily used for remote desktop administration. It does allow multiple users to view the same desktop, but does not have features to help users decide who is in control. By default, VNC uses a *free-floor* protocol to mediate control: the server simply handles keyboard and mouse inputs in the order in which they arrive, which can cause unpredictable results.

We modified both the client and server to use our turn-taking protocol. First, we modified the Remote Frame-Buffer (RFB) protocol used by VNC to communicate information between the clients and the server. We added support for messages such as requests for control and changes in control. The client was altered to support haptic input through the modified *iFeel* mice and traditional GUI input through the Button Bar, as well as to provide haptic feedback and display information in the User Window based on messages received from the server. The changes were implemented such that any of these elements could be enabled or disabled at will, allowing us to easily reconfigure the client for each of the three conditions. The server was modified to process requests for control, to only accept mouse and keyboard events from the client in control, and to keep the clients informed of changes in the turn-taking state. The server had no knowledge of what input or output methods were being used on each client.

The study was conducted in the Sensory Perception and Interaction (SPIN) Lab at the University of British Columbia. To simulate a distributed setting, subjects were seated at workstations as shown in Figure 6-2 such that they could not easily see each other. Subjects wore Sennheiser HD280 headphones and Sony ECM-T115 lapel microphones so they could communicate with one other easily. The computers used by the subjects included Pentium III and Pentium IV computers, with clock speeds ranging from 733 MHz to 2 GHz, and between 256 and 512 MB of RAM. Despite the variation in hardware, application performance was similar across computers because the software used in the study was not computationally intensive. Each computer had a 17" LCD display with 1280x1024 screen resolution.

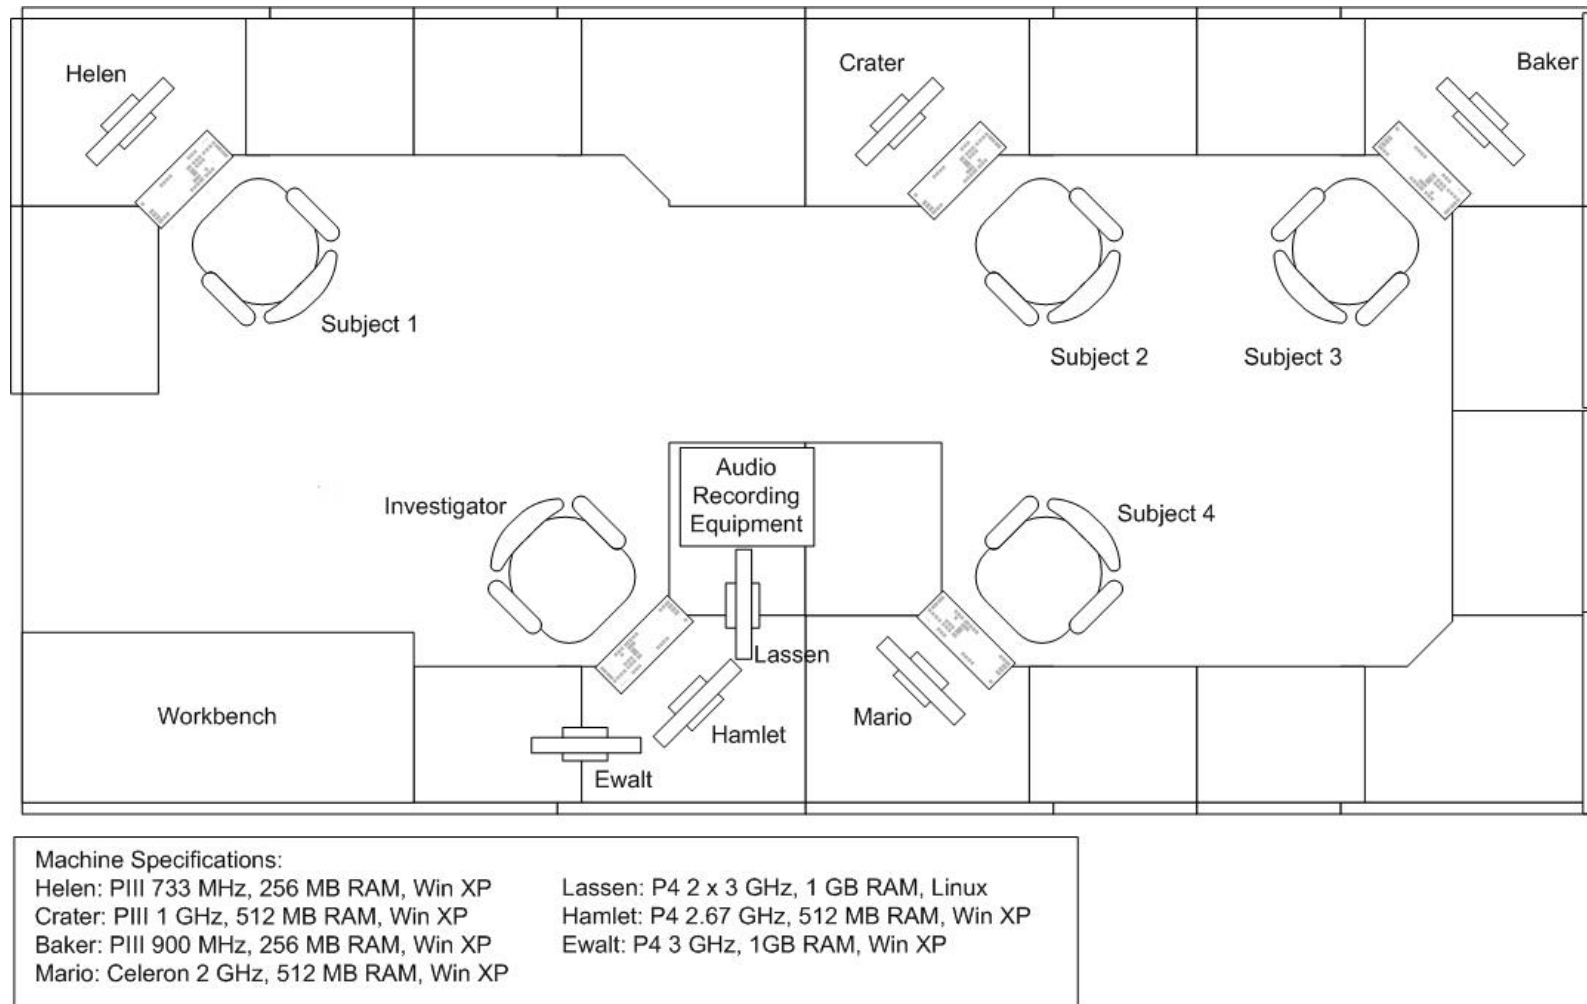


Figure 6-2 - Layout of experiment room in Study 3.

Unlike the web conferencing systems described in Chapter 2, where one user's computer acts as the host for the other collaborators, we used a separate computer called Hamlet to host the shared application. Each subject's computer ran a VNC client that connected via a 100 Mbps LAN to a VNC server running on Hamlet, a Pentium IV 2.67 GHz computer with 512 MB of RAM. We chose this approach to simplify our modifications to the VNC server; it was designed to share the entire desktop of the computer on which it was running rather than a specific application. This meant we could not easily display elements like the User Window and Button Bar on the server without all of the clients seeing copies of them as well. To allow the investigator to monitor the turn-taking state as groups worked on each condition, the computer called Ewalt was used as an "observer", essentially an inactive fifth client machine. Ewalt connected to Hamlet using the same VNC client as the subjects' computers, displaying Hamlet's desktop, the User Window, and the Button Bar from the visual condition. All of these computers ran Windows XP.

The audio setup used in the study is shown in Figure 6-3. Microphone pickup from the subjects and the investigator was fed into a Mackie 12-channel mixer. For subjects 1 – 3, Eurorack mixers were used to boost the signal from the microphones to the mixer. The output from the Mackie mixer was sent to a tape deck for recording and to a headphone amplifier so that subjects could hear each other speaking. The amplifier was required because the Mackie mixer only had 3 outputs, and we required one for each of the five pairs of headphones used.

We automated data collection in several ways. The VNC client and server were instrumented to record information in log files. Recording the efforts of each group as they worked (including their turn-taking state) required a somewhat complicated setup. We turned to an open-source, Linux screen recording program called *vncrec* [37]; it uses a modified VNC client to receive screen updates from a VNC server and record them to disk. This program ran on a computer called Lassen, running Red Hat Linux (kernel 2.4.20-31.9smp). We then ran a VNC server on Ewalt and configured *vncrec* to access it. Therefore, Lassen recorded the screen of our "observer" machine Ewalt. Another Linux open-source program, *transcode* [13], was later used to create MPEG4 movies of the video data. Group conversations were recorded on audio tape, as described.

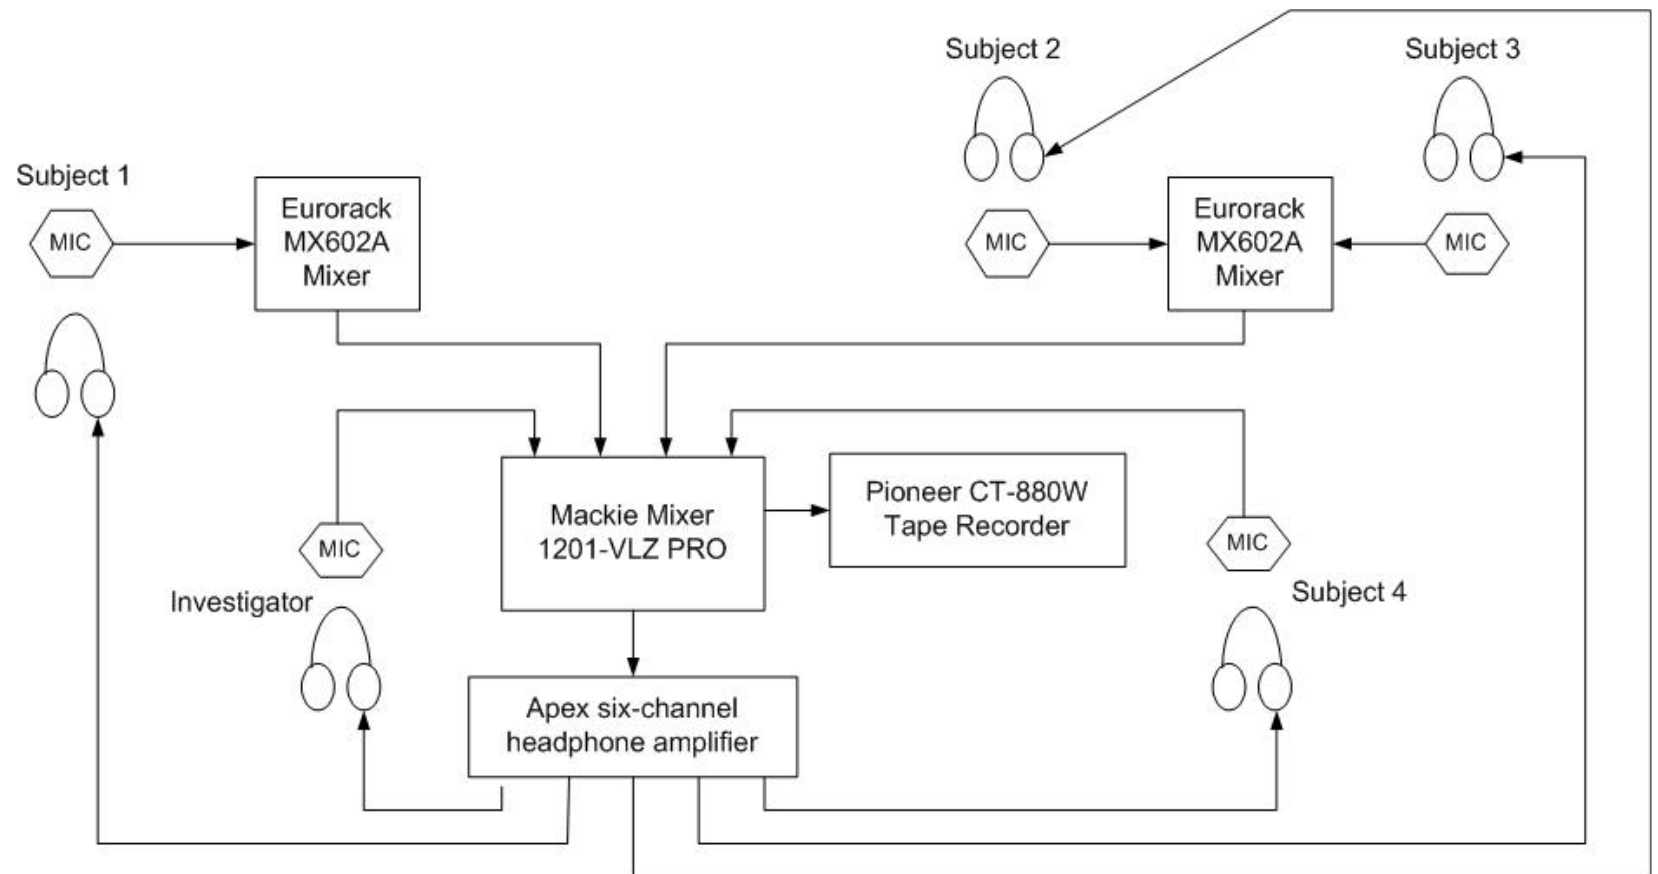


Figure 6-3 - Audio setup of experiment room in Study 3.

6.4 Task

Our task was designed to closely approximate real-world group collaboration. Several task characteristics were deemed important:

- Groups should share a common body of knowledge, but each individual should possess specific, specialized knowledge.
- Groups should work towards a well-defined set of goals, but there should be constraints on how the goals can be achieved.
- Group members should have conflicting interests, but collaboration should not be adversarial.

We developed a furniture-layout task that satisfied these characteristics. We created three isomorphic tasks for our three conditions, all centering around furniture layout in a typical graduate-level lab in computer science. To deepen the collaborative aspect, the three tasks shared an identical set of eight firm constraints that *had* to be observed as specific task goals were met, and an identical set of 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 was that three-foot wide walkways had to exist to each piece of furniture in the room; a soft constraint was that noisy areas should be isolated from workstations. The tasks are described further in Appendix F, and the initial layouts for each of the tasks are shown in Appendix I.

Groups were given 20 minutes to formulate a solution for each task. All members knew the complete set of goals they were to achieve; however, each member was responsible for two hard and two soft constraints, which were provided in written form. To mitigate subject development of expertise in one particular aspect of the task, subjects 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. Tasks were completed using Microsoft Visio, a diagramming tool.

6.5 Study Procedure

Subjects individually completed a training phase before working together on the furniture-layout tasks. First, subjects learned how to use Microsoft Visio. We simplified its interface to make it easier to learn, hiding all functions not needed in the study. Subjects 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 training exercises, ensuring that they understood its use at this level. Next, subjects were trained to identify the haptic icons we used, first reading descriptions of the stimuli and their protocol-based meanings. A Java training application similar to the one used in Study 2 was provided: subjects were shown a screen with a button for each of the 7 icons, ordered by family, which they could play as many times as they wanted, in any order. Hints were provided as descriptions on that same screen to help subjects learn the associations (see “Feel” descriptions, Table 6-1). When subjects felt they had learned the icons and their meanings, they proceeded to a learning test.

In the learning test, identical in format to the test used in Study 2, each icon was presented three times in random order, for a total of 21 trials. Subjects had to identify 19 or more icons correctly to pass; otherwise, they returned to the initial screen for more practice before repeating the test. Subjects identified icons by clicking on the correspondingly labeled radio button; to prevent subjects from learning positional rather than meaning associations with the stimuli, the labeled radio buttons were randomly re-ordered on each trial. No feedback was given during the test other than whether they passed or failed.

After the training phase, the group completed the three study conditions. Each condition was preceded by a five-minute 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 minutes working on a furniture-layout task. After each condition, subjects individually completed a questionnaire and were given five-minute rest breaks. At the end of the study, subjects individually completed an overall questionnaire. They were then interviewed and debriefed as a group.

The study required one three-hour session, for which subjects were each paid \$25. As an incentive, groups were told that their task solutions would be evaluated, and the top ¼ of the groups would each receive a \$40 bonus.

6.6 Study Design and Subjects

A within-subjects design was used so that subjects could compare the modalities in the three conditions. We adopted a “2x2+1” design. The visual and haptic conditions were counterbalanced, as were two out of the three tasks. The remaining task was always paired with the haptic+visual condition, and it was always the last condition presented. By placing this condition last, we were able to record and ask which modality subjects relied on, having had equal exposure to the other two approaches. Our design required four groups, an appropriate number for an exploratory study.

Subjects were recruited with the following constraints: each group had to have at least one male and one female in it, and all members of the group had to be acquainted with one another. These constraints were imposed to simulate real-world group composition. Subjects were not screened with respect to Visio experience because the interface to Visio was reduced to the extent that only novice behavior was permitted, reducing any advantage advanced users might have. No person with familiarity of haptic stimuli similar to those used in our study was allowed to participate.

6.7 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 specific goals, while respecting the constraints.

We also collected data from questionnaires and post-study interview data. The questionnaires consisted of Likert-scale and open-ended questions, and questions where subjects ranked the modalities in order of preference.

6.8 Results

Four groups of 4 subjects participated in the study, with 16 subjects in total (8 male, 8 female). All subjects were students at the University of British Columbia. Subjects ranged in age from 18 to 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: 6 reported none, 6 used game controllers with vibrotactile displays, 3 had used other haptic devices, and one did not respond to the question. Each group is described below:

- *The Engineers (3 males, 1 female)* had taken over a year of undergraduate engineering courses together, participated jointly in extracurricular activities, and kept in touch even though one subject had changed faculties.
- *The Long-Time Friends (2 males, 2 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 (1 male, 3 females)* were completing their Education degrees, as part of a cohort of approximately 40 people who took all their classes together for one year.
- *The Graduate Students (2 males, 2 females)* consisted of two male-female pairs from different research labs in Computer Science and Electrical and Computer Engineering. While each pair knew each other, the pairs did not; we permitted this due to recruitment difficulty.

We now summarize the study results according to our research questions. Although we did not anticipate any statistically significant results given the small number of groups, for completeness and curiosity we did run ANOVAs on some of the dependent measures across the conditions and, where significant, we report those results. Complete results from the post-condition questionnaires can be found in Appendix J.

6.8.1 Learning and Using Haptic Stimuli

Unlike Study 2, the learning component of the study was designed to ensure that subjects could achieve a threshold level of performance within a reasonable amount of time, rather than test how quickly subjects could learn the haptic icons. Therefore, subjects were encouraged to learn ‘carefully’ rather than ‘quickly’, and it is possible that learning times reported here could have been even lower.

As in Study 2, learning time was calculated as the time spent exploring the haptic stimuli in the Java application. The learning test itself was not included, as results would be skewed by subjects who adopted an aggressive strategy, trying the test quickly and needing multiple attempts to pass it. Indeed, we observed that 3 of the 5 subjects with the longest

learning times required only one attempt to pass the learning test. Nine subjects only required one attempt to pass the test, 5 subjects required two attempts, 1 subject required 3 attempts, and 1 subject required 5 attempts. Interestingly, the subject who required the most attempts had a learning time close to the average time. Learning times ranged from 51 – 270 seconds, with a mean time of 135 seconds and a standard deviation of 64 seconds. By comparison, the mean time in Study 2 was 177 seconds with a standard deviation of 114 seconds. The decrease in learning time was likely related to the provision of hints to help subjects learn the icons.

These results show that associations between a moderate-sized set of well-designed haptic stimuli and compatible meanings can be quickly learned to a high degree of accuracy.

6.8.2 Collaborative Style

We investigated several aspects of group collaboration. First, we were interested in the approaches that groups would use to solve the floor-layout tasks. Second, we examined the impact of condition on the frequency of control transfer. Third, we investigated the distribution of different verbal methods for gaining control. Fourth, we explored the distribution of the different methods for gaining control, and the influence of method on wait times.

Groups used a variety of strategies to solve the floor-layout tasks. Although we did not search for links between the strategies employed and task performance or collaborative style, we include this information for completeness. The Engineers and Teachers typically added and repositioned furniture piece-by-piece within the room, finding appropriate locations as they worked. The Long-Time Friends and the Graduate Students preferred to move large groups of furniture outside the room to create open space, then rearranged the furniture as it was moved back in. At the beginning, all groups except for the Engineers shared their constraints with one another before starting to work; by the final task, the groups had become sufficiently familiar with the constraints such that this was not needed.

Across all groups and conditions, the number of control changes ranged from 8 to 29. As shown in Table 6-2, there were nearly twice as many changes in the haptic and haptic+visual conditions compared to the visual condition; a one-way 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.

	Visual	Haptic	Haptic+Visual
Engineers	13	22*	17
Long-Time Friends	20	18*	25
Teachers	8*	20	18
Graduate Students	9*	21	29
Average	12.50	20.25	22.25
Std. Deviation	5.45	1.71	5.74

Table 6-2 - Number of changes in control for each condition; * denotes condition seen first (N=16).

We expected verbal communication to play a role in mediating turn-taking among group members, just as it does in face-to-face collaboration. However, we were curious as to the extent to which it would be used in the different conditions. We counted the number of explicit, implicit, and nonverbal requests for control across groups in each condition. We defined an explicit request for control to be statements such as, “I want control.” Statements like, “I have an idea” were categorized as implicit requests for control. A nonverbal request was defined as using the turn-taking protocol to indicate a desire for control without uttering a word. Explicit and implicit requests for control were typically used in conjunction with the turn-taking protocol. On some occasions, subjects made multiple verbal requests in an attempt to obtain control and share their ideas.

The distribution of explicit, implicit, and nonverbal requests is shown in Figure 6-4. Across all conditions, nonverbal requests were used most frequently; a 3x3 (condition x method) repeated-measures ANOVA showed a significant effect of method ($F_{2,6} = 10.615, p = 0.011$, partial $\eta^2 = 0.780$) with post-hoc comparisons showing that nonverbal requests were used more often than implicit requests ($p = 0.009$). While we hypothesized that there might be a significant effect of condition as well, noticing there were more requests in the haptic and haptic+visual conditions than the visual condition, the statistical analysis was only marginally significant ($F_{2,6} = 3.717, p = 0.089$, partial $\eta^2 = 0.553$). Although no significant interaction was found, it is interesting to note that nonverbal requests were used more in the haptic+visual condition than either the haptic or visual conditions (58% versus 51% and 45%). This might indicate that subjects relied more on the turn-taking protocol as they became familiar with it.

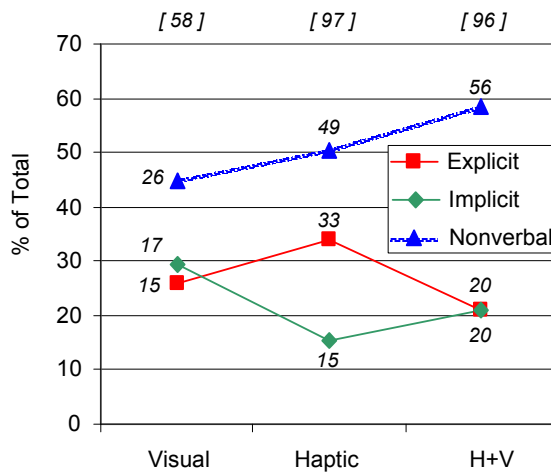


Figure 6-4 - Distribution of verbal methods of requesting control. Numbers in italics show actual counts and numbers in square brackets show total counts (N = 16).

Recall that subjects could attempt to gain control by *taking*, *urgently requesting*, *gently requesting*, or by directly *obtaining* control. The obtain control category represents a control acquisition through any of the other 3 methods when no one is already in control. It is distinguished because we cannot determine from the data when subjects knew someone was in control or not; for example, a *take* in the former situation may represent an aggressive collaborative style, whereas in the latter situation it is simply one of 3 methods to assume control when no one else has it.

Figure 6-5 shows the distribution of non-verbal methods for gaining control used for each condition. It does not include obtain controls, which represent 12, 39, and 46 attempted acquisitions in the visual, haptic, and haptic+visual conditions respectively. The figure shows a clear preference for *gentle requests* over *takes* and *urgent requests* in all conditions: gentle requests accounted for 60-73% of the acquisition requests made when control was held by another.

By contrast, the balance between the two stronger acquisition methods, *takes* and *urgent requests*, suggests their use differed by condition. In the visual condition, there was a clear preference for *take* over *urgent request* (37% compared to 2%), whereas in the haptic and haptic+visual conditions they were used more equally (16% compared to 16%, and 16% compared to 12%, respectively). A 3x3 (method x condition) repeated-measures ANOVA showed a significant effect of method ($F_{2,6} = 17.401, p = 0.003$, partial $\eta^2 = 0.853$) with post-




Figure 6-5 - Distribution of non-verbal methods for requesting control. Numbers in italics show actual counts and numbers in square brackets show total counts. These do not include directly *obtained controls*; the values are 12, 39, and 46 for the visual, haptic, and haptic+visual conditions respectively (N=16).

hoc pair-wise comparisons showing that gentle requests outnumbered urgent requests ($p = 0.054$) and takes ($p = 0.093$).

The data suggests increased access of our protocol's key feature, the graded request, in conditions including haptics: urgent requests represented 11% of all requests across all conditions, distributed as 2/16/12% in the visual, haptic and haptic+visual conditions. Finally, across all conditions, control holders were more responsive to urgent requests, releasing control in an average of 19.4 seconds as opposed to 29.3 seconds for gentle requests (Table 6-3 and Table 6-4). Several additional details aid interpretation of the numbers:

- Having only one urgent request in the visual condition may have been because request urgency was not emphasized in that condition; the requestor's name simply appeared in a different list in the User Window.
- One group, the Graduate Students, never used urgent requests under any condition. Additionally, some subjects reported post-study that they did not feel the need to make urgent requests because a gentle request was fast enough.
- In most groups, a subject typically retained 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 the obtain control counts for the haptic and haptic+visual conditions.

Conditions including haptics thus seemed to facilitate increased turnover in control frequency and a more even distribution of urgent requests and takes as compared to the visual condition, indicating a generally more dynamic collaborative style. Group practices also played an important role.

6.8.3 Equitability of Sharing Control

Although control turnover was more frequent in the presence of haptics, we wanted to know whether haptics promoted equitability of control amongst team members. By examining how much time subjects 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 subject spent in control of Visio in any one turn, before releasing or losing control, was noticeably larger in the visual condition than the other conditions (Table 6-5). A one-way repeated-measures ANOVA 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 statistically significant.

This suggests that subjects in control of Visio were more responsive to requests for control in the haptic and haptic+visual conditions than in the visual condition. This hypothesis is supported by the average wait durations after a gentle request (Table 6-3) and an urgent request (Table 6-4) control, both of which are shorter when haptic feedback is present. For gentle requests, a one-way ANOVA showed a significant main 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 again not statistically significant. Three of the four groups did not make urgent requests at all in the visual condition, making that measure more difficult to compare across conditions.

Examination of the overall percentage of time each subject was in control under each condition revealed an interesting finding. For each group, a spread was calculated by subtracting the percentage of the subject in control the least from the percentage of the subject in control the most (Table 6-6). We observed that the spreads were larger in the visual condition as compared to haptic+visual and especially the haptic condition. Statistical analysis with a one-way ANOVA yielded a main effect of condition ($F_{2,6} = 37.405, p < 0.001$

	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
Col Sum / # Groups	58.1	23.9	6.0
Std. Deviation	37.3	14.6	2.3

Table 6-3 - Gentle Requestor's Perspective. Average time from a gentle request until gaining control (sec), by group and condition. Unweighted table average 29.3. (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	-	-	-
Col Sum/# Groups	67.0	13.7	9.2
Std. Deviation	-	17.0	11.1

Table 6-4 - Urgent Requestor's Perspective. Average time from an urgent request until gaining control (sec), by group and condition. '-': no urgent requests made. Unweighted table average = 19.4. (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
Col Sum / # Groups	104.3	50.8	56.3
Std. Deviation	39.5	12.9	16.6

Table 6-5 - Control-Holder's Perspective. Average lengths of periods in control before releasing or losing control (sec), by group and condition. Unweighted table average = 70.5. (N=16).

	Visual	Haptic	Haptic+Visual
Engineers	44	23	35
Long-Time Friends	39	25	24
Teachers	39	13	22
Graduate Students	47	25	35
Col Sum/# Groups	42	22	29
Std. Deviation	4	6	7

Table 6-6 - Equitability of Control Time. Spread of percentage of time in control, between group members most and least in control (%). Unweighted table average = 0.31. (N=16).

partial $\eta^2 = 0.926$), with post-hoc pair-wise comparisons showing that the spread was significantly larger in the visual condition than the haptic ($p = 0.011$) and haptic+visual ($p = 0.014$) conditions.

Together these results suggest that the sharing of control is more equitable in the presence of haptics.

6.8.4 Subject Preferences

At the end of the experiment, subjects ranked the three conditions in order of preference for obtaining control (input to the system), for displaying the turn-taking state (output/feedback from the system), and for overall preference. Table 6-7 shows that the haptic+visual condition was overwhelmingly preferred overall (12 subjects) as compared to the two single-modality conditions, which each received an equal number of overall first order rankings (2 subjects each). It is thus unsurprising that haptic+visual was also favored for both obtaining control and displaying state (11 subjects). However, the haptic condition was the next preferred for obtaining control, whereas for displaying state, the visual condition was next preferred.

When asked to justify their overall rankings, subjects who ranked the haptic+visual condition first noted that the haptic feedback notified them of changes in state, while the User Window indicated who was in control or requesting control. One subject noted the lack of haptic feedback when no one was in control, and sometimes displayed the User Window to confirm that fact. This may be a limitation of our design that could be considered in the next iteration. Another difficulty that arose occasionally in the haptic condition was that the person gaining control was not always who the person releasing control expected it would be, based on the audio dialogue. To address this haptically, it might be necessary to add subject identification information to the stimuli; further study would be required to determine if this

	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

Table 6-7 - Condition Preference. Number of subjects who ranked a given condition first / second / third (N = 16).

is perceptually workable.

Subjects clearly expressed a desire for the User Window to be present. However, the data did not show that they actually used it. After each condition, subjects were asked to indicate how much they agreed or disagreed with the statement, “I constantly monitored the User Window when someone asked me for control.” In the haptic condition, 11 subjects either disagreed or strongly disagreed with the statement, compared to 6 subjects in the haptic+visual condition and 3 subjects in the visual condition. Furthermore, the number of times the User Window was opened was measured in the haptic condition; on average, each subject opened it fewer than three times.

In summary, subjects preferred the haptic+visual condition overall. However, the justification for preferring features specific to the visual condition is ambiguous and we anticipate that it might diminish with familiarity of the haptic features.

6.8.5 Task Performance

For completeness, we checked task performance across conditions, even though significant variation was not expected due to the short task duration. Task solutions were evaluated based on how well they satisfied the goals for the task, while following the specified hard and soft constraints. We used a points-based system: points were awarded for satisfying a specific goal, and points were deducted for failing to satisfy a constraint. The penalty for violating hard constraints was more severe than violating soft constraints. Prior to running the study, a reference solution for each task was created to estimate its maximum possible score.

Each group’s task solution was scored, and the resulting score was divided by the reference solution score (Table 6-8). As would be expected, a significant learning effect was found ($F_{2,6} = 27.167, p = 0.001$, partial $\eta^2 = 0.901$), with post-hoc comparisons showing that scores in the second and third condition to be significantly better than the first ($p = 0.04$ for both). Since the order of presentation of the visual and haptic conditions was counterbalanced, we compared their scores. On average, groups had slightly better performance in the haptic condition than the visual condition, but the difference is not statistically significant. The better average performance in the haptic+visual condition is almost certainly due to learning effects.

	Visual	Haptic	Haptic+Visual
Engineers	0.69	0.28*	0.84
Long-Time Friends	0.60	0.44*	0.65
Teachers	0.22*	0.70	0.73
Graduate Students	0.41*	0.79	0.83
Col Sum/# Groups	0.48	0.55	0.76
Std. Deviation	0.21	0.23	0.09

Table 6-8 - Normalized task scores. ‘’ denotes condition seen first (N=4).**

We also checked to ensure that there was no effect of task in the first two conditions, which were counterbalanced 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 results suggest that the different modalities did not impact task performance. Longer exposure to the conditions, however, may yield measurable differences.

6.9 Discussion

Impact on Quality of Collaboration and Equitable Sharing

We found that our background-level, haptically-supplied information increased both overall turnover of control and usage of graded requests, and it also seemed to promote equitability in the total amount of time each subject was in charge. To explain these positive effects, we hypothesize that users may have found this method of control exchange either less cumbersome or more informative – since no condition *compelled* users to give up control more than others, any increase in either overall frequency or use of a specific method was presumably by preference.

Thus we see consistent evidence across the available indicators that this kind of carefully designed haptic feedback can facilitate a more active, equitable distributed collaborative style, in ways that might apply to co-located collaboration as well. It will be of interest to see whether in more extensive tests, this extends beyond notions of fairness to better performance.

Stakes

Urgent requests represented 11% of all requests made when another subject had control (2% in the visual condition, 16% in the haptic condition, 12% in the haptic+visual condition). We observe that while gentle requests will presumably dominate any effective interaction, the low stakes typical in an experiment suggest that this result may be conservative: despite our incentives, our subjects were conscientious and polite as well as new to the concepts involved. In a real-world (non-experiment) setting, when group members have strongly vested interests in the project outcome, the availability of choices other than rudeness or silence might allow high-stakes collaboration to be intense without becoming adversarial. This benefit could apply to both co-located and distributed situations.

Familiarity and Learning

We required our subjects to quickly learn several new concepts, including the notion of distributed collaboration, a turn-taking protocol, and identifying haptic stimuli. Newness may thus be a factor in some observations. For example, subjects clearly preferred having the User Window displayed all the time, but did not use it as much as their preference might suggest. With time, will they outgrow this desire for the familiar? Will they gain both the ability to process the haptic icons more automatically, and the confidence to rely on this channel?

Utilization of Information in Haptic Icons

Our Study 3 learning data demonstrates our subjects' facility in learning our haptic icons, but does not tell us how easily haptic content was accessed under workload. Near the end of the study we began to follow up on a question of subjects using the haptic icons merely as binary triggers. Some claimed they used the icons only as notification; others said they were able to identify specific meanings. However, they may have been unaware of their identification ability, and of the extent to which they were utilizing meanings. Subjects in Studies 2 and 3 occasionally expressed surprise when they passed the learning test, and our Study 2 results demonstrated that subjects consistently identified haptic icons at 95% accuracy across different levels of workload. There is substantial evidence for nonconscious perception and utilization of visual stimuli [49, 54], and it is likely that this mechanism exists

for haptic perception. However, longer exposure and different tests will be required to definitively answer this question.

Value of an Urgency-Based Protocol

We were not able to explicitly compare our augmentations to the presumed ideal of co-location, nor to a protocol without graded control requests. However, when given the chance to easily make graded requests, subjects did so; further, control-holders were overall more responsive to urgent requests than to gentle requests. Together, these results demonstrate the potential impact of the urgency-communicating aspect of our protocol on both control holder and requestor behavior, and suggest that our subjects found it useful. This justifies proceeding to the next step, which is to establish whether this protocol can bring the quality of collaboration and task performance to co-located levels or even beyond.

6.10 Summary

We conducted an exploratory, observational study to evaluate our urgency-based turn-taking protocol in three conditions: one that relied primarily on haptic cues, one that relied on visual cues, and one that used both modalities. Our goals were to see how quickly subjects could learn to identify the haptic icons, how the different conditions would affect group collaboration and task performance, and which modality subjects would prefer. Four groups of subjects from diverse academic backgrounds participated in the study.

Subjects were able to learn the haptic icons in an average of 135 seconds, nearly a minute faster than subjects in Study 2. In conditions with haptic feedback, we observed that more changes in control occurred between subjects and that subjects were more responsive to requests for control. When another subject was in control, subjects preferred to obtain control by gently requesting it and waiting. Although not heavily used, we noted that urgent requests were used more frequently in the haptic and haptic+visual conditions than in the visual. As well, takes were used less frequently in the haptic and haptic+visual conditions than the visual condition.

We noticed that turn-taking was more equitable in conditions with haptic feedback. This did not come at the expense of task performance, as performance was steady across conditions. Subjects overwhelmingly preferred the haptic+visual condition, finding haptic feedback useful as a mechanism for notifying them of changes in state, and the visual

information useful for discerning the identify of a person requesting control. However, based on their usage of the User Window, it is possible that with more exposure, subjects' perceived need for visual feedback would decrease.

Conclusions and Future Work

In this thesis, we described the design, implementation, and evaluation of an urgency-based turn-taking protocol, where haptic icons delivered through a vibrotactile display inform a user of their current collaborative state. We summarize our major findings, discuss their implications, and discuss future work.

7.1 Using MDS to Categorize Haptic Icons

In Study 1, we used an MDS algorithm to select 3 stimuli from a set of 24 to use as our *in control* icons, and attempted to verify that the *change in control* icons we had prototyped were distinct from these stimuli and from each other. As well, we hoped to identify the dimensions used by subjects to distinguish between vibrotactile stimuli.

Unlike MacLean and Enriquez, who found that subjects categorized stimuli delivered through a force-feedback knob based on waveform, then frequency, then magnitude [46], our results were somewhat mixed. Most subjects categorized primarily on the number of bursts delivered, placing single- and double-burst stimuli into separate clusters. After this, frequency and magnitude played roughly equal roles. A possible reason for this is that we were unable to deliver four perceptually equal magnitudes across the three frequencies. Using a different vibrotactile display (such as a voice coil) would allow us to rectify this problem, and may lead to more conclusive results.

Given our results, we conservatively selected 3 stimuli to use as the *in control* icons that were as mutually distinct from one another as possible and from the *change in control* icons. This strategy appeared to be successful, as Studies 2 and 3 showed that subjects did not have difficulty learning or distinguishing these icons. However, while our analysis showed

that the *change in control* icons were perceptually similar, it could not predict that they would be difficult to distinguish. As a result, the icons had to be adjusted twice to improve their distinctiveness. Further work is needed to quantify the distinctiveness of stimuli and establish minimum thresholds to ensure they are distinguishable. MDS may not be suitable for this role.

7.2 Learning Haptic Icons

In Study 2 and Study 3, subjects were able to learn the set of 7 haptic icons in a short period of time: 177 seconds on average in Study 2 and 135 seconds on average in Study 3. The difference is likely related to the provision of hints in Study 3 to help subjects learn the icons. In both cases, subjects had to correctly identify 90% of the icons presented in a test to complete the learning process. These results show that it is possible to design haptic icons that can be learned with modest effort and in a short period of time, something necessary for them to be adopted more widely as an interaction method.

Admittedly, the haptic stimuli were not randomly associated with meanings, but were intentionally designed to be as intuitive as possible by drawing on common metaphors. Nonetheless, the haptic icons effectively represent different gradients of three concepts. Further work could explore how many concepts can be learned, and how many gradients for each concept. In addition, the ability of subjects to retain the meanings they have learned could be measured, to see whether subjects can commit the haptic icons to long-term memory, just as they learn to recognize textures and surfaces in the physical world.

7.3 Identifying Haptic Icons while Engaged in Other Tasks

In Study 2, the main goal was to evaluate how increased levels of workload would affect subjects' performance. Subjects identified haptic icons in three conditions: one with no distracter tasks, one with a visual distracter task, and one with both a visual and an auditory distracter task. As might be expected, the time required for subjects to detect changes in haptic icons increased significantly with workload, from 1.8 seconds in the haptic condition to 4.3 seconds in the haptic+visual+auditory condition. Identifying the icons took an additional 2.5 – 3.0 seconds; the effect of condition was marginally significant. Combining the detection and identification times, a subject was still able to notice and identify a change

in haptic icon within 7.3 seconds, which we deemed acceptable for our purposes. Subjects identified icons with 95% accuracy across conditions, a result that pleasantly surprised us.

To explore further the effect of workload on these measures, it would be useful to use a distracter task where the workload imposed can be systematically increased. This would allow us to discern how performance degrades - whether it gradually decreases or whether it drops off sharply after a certain point. For example, the visual puzzle task could be modified so that subjects must complete a puzzle within a certain period of time, with the amount of time decreasing as each level is passed. Another possible task would be for subjects to play a block-placement game like Tetris, with difficulty increasing each level. A Tetris-like task that requires continuous attention, paired with verbal identification of haptic icons, would encourage subjects to identify the icons in parallel with completing the task.

7.4 Haptic Feedback for Mediating Turn-Taking

In Study 3, we conducted an exploratory observational study to examine how groups of subjects would use our turn-taking protocol when collaborating on furniture-layout tasks. Groups used three different implementations of the protocol: one that relied primarily on haptic interaction, one that was visual, and one that combined both modalities. Although the lack of statistical power from the small number of groups meant that statistical tests often returned marginally significant or non-significant results, we still noticed some interesting trends. Our results showed that more turnovers in control occurred and subjects were faster to respond to requests for control in conditions with haptic feedback. Thus, it seems that the haptics provided a convenient and effective channel for conveying information.

We also found that sharing of control was more equitable in conditions with haptic feedback. The improved equitability did not come at the expense of task performance, as task performance was fairly stable across conditions. When asked which modality they preferred for obtaining control, conveying state, and overall, in each case subjects overwhelmingly chose haptic+visual as their first choice. Subjects who chose haptic+visual overall liked having the haptic feedback to notify them of the current state, but also wanted to see the User Window so that they could tell who was in control or waiting for control.

7.5 Value of the Urgency-Based Turn-Taking Protocol

When we designed the urgency-based turn-taking protocol, we assumed that gentle requests would be used far more often than urgent requests or takes. Our Study 3 results showed this to be the case. As well, our data suggests that in conditions with haptic feedback, subjects relied more on urgent requests and less on takes to obtain control. We believe this is because subjects felt that the haptic feedback indicating an urgent request would effectively communicate the sense of urgency to the subject in control. Our results also showed that subjects were more responsive to urgent requests for control than gentle requests.

In laboratory studies, it is nearly impossible to recreate the stakes a task may carry in the real world. Real-world collaboration may include individuals with hidden or conflicting agendas, or individuals who dislike each other but must work together. At times, the collaboration may even be adversarial. In contrast, subjects in our study were polite, conscientious, and willing to share information with one another. As a result, the perceived need to urgently request or take control may have been lower in our evaluation than in a real-world setting. Since our protocol is built into a robust view-sharing system and the modified *iFeel* does not impose special hardware requirements, it would be feasible to conduct a field study with subjects who normally use another view-sharing system to collaborate. The systems could be compared to ascertain how the different protocol and the provision of haptic feedback affect collaboration.

7.6 Future Work

Besides the possibilities described in the previous sections, there are many avenues for future research. Owing to the exploratory nature of this work, the cost of paying subjects, and the amount of time required to analyze each group's data, we designed the study so that only four groups were required, and only compared different implementations of our turn-taking protocol. It would be useful to replicate the study with more groups, using a fully-counterbalanced design in place of our "2x2+1" design. As well, a comparison of our protocol to the different variants of the *give* protocol used in current view-sharing systems would be informative.

Subjects only used our system for a brief period of time. A longitudinal study would allow us to see whether the effects that haptic feedback had on group collaboration would continue, decline, or strengthen. As subjects grow comfortable with the system, we could also see whether their preferences change, such as whether they continue to feel the User Window is necessary. Prolonged exposure would also allow us to learn whether the haptic icons are too intrusive or too subtle; it is possible that the icon magnitudes would have to be calibrated for each user so that they are sufficiently noticeable without being annoying.

Group size is quite likely to influence collaboration using our turn-taking protocol. While our protocol could be used with a group as small as two, its benefit over any other protocol would be minimal. However as group size increases, we feel our protocol would be increasingly useful since more individuals will be vying for control and verbal mediation will become increasingly difficult. To test this assertion, we could compare two (or more) sizes of groups, as in previous studies [48]. Of course, with too large a group, collaboration using any means would be very difficult.

Further studies could also examine the effect of group composition on collaboration using our turn-taking protocol. Groups composed of close friends may use our protocol differently than groups composed of strangers, as might groups who have a strong natural leader (or leaders) versus groups that do not. It would also be interesting to see whether our protocol leads to increased participation from introverted individuals who are hesitant to assert themselves verbally.

We have shown that our haptically-supported protocol improves distributed collaboration in a furniture-layout task. It is reasonable to presume this will generalize to other common tasks, such as design reviews or document editing. Our protocol could also be used to coordinate resources, such as the positions of air and ground crews battling a forest fire. In co-located or distributed meetings, haptic feedback could remind a speaker of others' desire to speak. If a participant moderates the meeting, haptic feedback may allow him or her to concentrate more on the meeting, and less on monitoring the other participants' wishes. Similarly, during presentations an audience could indicate their interest, boredom, or confusion, and the presenter could receive the information through haptic icons, rather than having to monitor status displays. Our protocol could also be useful in specialized applications such as air-traffic control, where a controller routes aircraft in a given zone and passes control to other controllers as the aircraft leave it. The urgency with which an aircraft

must be handled could be set either by the controller or by an intelligent system monitoring flight paths.

Research on haptic communication is still in its infancy. This thesis has shown that haptic icons can be learned quickly and used effectively in an environment where a user's primary focus is on a visual task. Subjects in our studies not only understood messages delivered through the haptic sense, but also responded to those haptic messages more quickly than the same messages delivered through the traditional visual modality. Our *urgency-based* turn-taking protocol also shows promise, as subjects took advantage of the different ways of obtaining control and responded appropriately to gentle and urgent requests for control.

Bibliography

1. BMW AG. *iDrive*. Accessed: September 14, 2004. <http://www.bmw.com/generic/com/en/fascination/technology/technologyguide/index.html>.
2. Logitech. *Logitech MX 1000*. Accessed: September 14, 2004. <http://www.logitech.com/index.cfm/products/details/US/EN,CRID=3,CONTENTID=9043>.
3. Macromedia, Inc. *Macromedia Breeze*. Accessed: September 3, 2004. <http://www.macromedia.com/software/breeze/>.
4. Placeware Inc. *Microsoft LiveMeeting*. Accessed: September 3, 2004. <http://main.placeware.com/>.
5. SensAble Technologies, Inc. *PHANTOM devices*. Accessed: September 14, 2004. http://www.sensable.com/products/phantom_ghost/phantom.asp.
6. SPSS, Inc. *SPSS*. Accessed: September 14, 2004. <http://www.spss.com>.
7. RealVNC Ltd. *Virtual Network Computing*. Accessed: September 12, 2004. <http://www.realvnc.com/>.
8. WebEx Communications Inc. *WebEx*. Accessed: April 6, 2004. <http://www.webex.com/>.
9. Microsoft Corporation. *Windows NetMeeting*. Accessed: April 6, 2004. <http://www.microsoft.com/windows/netmeeting/>.
10. Amato, I. (2001). Helping Doctors Feel Better. *Technology Review*, 104 (3), 64-69.
11. Baecker, R.M., Grudin, J., Buxton, W.A.S., and Greenberg, S. (1995). *Readings in Human-Computer Interaction: Towards the Year 2000*. San Francisco, California: Morgan Kaufmann Publishers.
12. Basdogan, C., Ho, C., Srinivasan, M., and Slater, M. (2000). An experimental study on the role of touch in shared virtual environments. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7 (4), 443-460.
13. Bitterberg, T. *transcode*. Accessed: September 13, 2004. <http://zebra.fh-weingarten.de/~transcode/>.
14. Boyd, J. (1993). Floor control policies in multi-user applications. In *INTERACT '93 and CHI '93 Conference Companion on Human Factors in Computing Systems* (pp. 107-108). Amsterdam, The Netherlands: ACM Press.
15. Boyle, E., Anderson, A., and Newlands, A. (1994). The Effects of Visibility on Dialogue and Performance in a Cooperative Problem Solving Task. *Language & Speech*, 37 (1), 1-20.

16. Boyle, M., and Greenberg, S. (2002). GroupLab Collabrary: A Toolkit for Multimedia Groupware. In *ACM CSCW 2002 Workshop on Network Services for Groupware*.
17. Brewster, S., and Brown, L. (2004). Non-visual information display using tactons. In *Extended Abstracts of the 2004 Conference on Human Factors and Computing Systems* (pp. 787-788). Vienna, Austria: ACM Press.
18. Brewster, S., Wright, P., and Edwards, A. (1993). An evaluation of earcons for use in auditory human-computer interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 222-227). Amsterdam, The Netherlands: ACM Press.
19. Chan, A., MacLean, K., and McGrenere, J. (2004). Designing Haptic Signals to Support Peripheral Awareness in the Presence of Workload. *Technical Report TR-2004-15*, Department of Computer Science, University of British Columbia.
20. Chan, A., McGrenere, J., and MacLean, K. (2004). Haptic Support for Urgency-Based Turn-Taking. *Technical Report TR-2004-14*, Department of Computer Science, University of British Columbia.
21. Crowley, T., Milazzo, P., Baker E., Forsdick, H., and Tomlinson, R. (1990). MMConf: An Infrastructure for Building Shared Multimedia Applications. In *Proceedings of the 1990 ACM Conference on Computer-Supported Cooperative Work* (pp. 329-342). Los Angeles, California: ACM Press.
22. Dennerlein, J., Martin, D., and Hassler, C. (2000). Force-feedback improves performance for steering and combined steering-targeting tasks. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. The Hague, Netherlands: ACM Press.
23. Duncan, S. (1972). Some Signals and Rules for Taking Speaking Turns in Conversations. *Journal of Personality and Social Psychology*, 23 (2), 283-292.
24. Duncan, S., and Niederehe, G. (1974). On Signalling That It's Your Turn to Speak. *Journal of Experimental Social Psychology*, 10, 234-247.
25. Ellis, C., Gibbs, S., and Rein, G. (1991). Groupware: some issues and experiences. *Communications of the ACM*, 34 (1), 39-58.
26. Engelbart, D. *Toward High Performance Knowledge Workers*. Accessed: April 6, 2004. <http://www.bootstrap.org/augdocs/augment-81010.htm>.
27. Engelbart, D., and English, W. (1968). A Research Center for Augmenting Human Intellect. In *AFIPS Conference Proceedings of the 1968 Fall Joint Computer Conference* (pp. 395-410). San Francisco, California.
28. Enriquez, M., and 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, California.
29. Enriquez, M., and MacLean, K. (2004). Impact of Haptic Warning Signal Reliability in a Time-and-Safety-Critical Task. In *12th Annual Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, IEEE-VR2004*. Chicago, USA.
30. Gaver, W. (1993). Synthesizing auditory icons. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 228-235). Amsterdam, The Netherlands: ACM Press.
31. Gibbs, S. (1989). LIZA: An Extensible Groupware Toolkit. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 29-35). ACM Press.

32. Green, P., Carmone Jr., F., Smith, S. (1989). *Multidimensional Scaling: Concepts and Applications*. Needham Heights, Massachusetts: Allyn and Bacon.
33. Greenberg, S. (1990). Sharing views and interactions with single-user applications. In *ACM/IEEE Conference on Office Information Systems* (pp. 227-237). Cambridge, Massachusetts.
34. Greenberg, S. (1991). Personalizable Groupware: Accomodating Individual Roles and Group Differences. In *Proceedings of the European Conference of Computer-Supported Cooperative Work (ECSCW '91)* (pp. 17-32). Amsterdam, Netherlands.
35. Grudin, J. (1994). Groupware and social dynamics: eight challenges for developers. *Communications of the ACM*, 37 (1), 92-105.
36. Gutwin, C., Roseman, M., and Greenberg, S. (1996). Workspace Awareness Support With Radar Views. In *Conference companion on Human factors in computing systems: common ground* (pp. 210-211). Vancouver, British Columbia: ACM Press.
37. Hayashi, Y. *vncrec*. Accessed: September 13, 2004.
<http://www.sodan.org/~penny/vncrec/>.
38. Inkpen, K., McGrenere, J., Booth, K., and Klawe, M (1997). Turn-Taking Protocols for Mouse-Driven Collaborative Environments. In *Graphics Interface '97* (pp. 138-145). Kelowna, Canada.
39. Johansen, R. (1988). *Groupware: Computer Support for Business Teams*. The Free Press.
40. Johansen, R. (ed.). (1989). *User approaches to computer-supported teams*. Hillsdale, New Jersey: Lawrence-Erlbaum Associates.
41. Klatzky, R., and Lederman, S. (1995). Identifying objects from a haptic glance. *Perception & Psychophysics*, 57 (8), 1111-1123.
42. Klatzky, R., and Lederman, S. (2003). Touch. In A.F. Healy, and Proctor, R.W. eds. *Handbook of Psychology*, (pp. 147-176). New York: John Wiley & Sons.
43. Klatzky, R., Lederman, S., Hamilton, C., Grindley, M., and Swendsen, R. (2003). Feeling textures through a probe: Effects of probe and surface geometry and exploratory factors. *Perception & Psychophysics*, 65 (4), 613-631.
44. Knister, M., and Prakash, A. (1990). DistEdit: a Distributed Toolkit for Supporting Multiple Group Editors. In *Proceedings of the 1990 ACM Conference on Computer-Supported Cooperative Work* (pp. 343-355). Los Angeles, California: ACM Press.
45. Lauwers, J., and Lantz, K. (1990). Collaboration awareness in support of collaboration transparency: requirements for the next generation of shared window systems. In *Proceedings of the SIGCHI conference on Human factors in computing systems: Empowering people* (pp. 303-311). Seattle, Washington.
46. MacLean, K., and Enriquez, M. (2003). Perceptual Design of Haptic Icons. In *Proceedings of Eurohaptics*. Dublin, Ireland.
47. McKinlay, A., Arnott, J., Procter, R., Masting, O., and 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).
48. McKinlay, A., Procter, R., Masting, O., Woodburn, R., and Arnott, J. (1994). Studies of turn-taking in computer-mediated communications. *Interacting with Computers*, 6 (2), 151-171.
49. Merikle, P., Smilek, D. and Eastwood, J.D. (2001). Perception Without Awareness: Perspectives from Cognitive Psychology. *Cognition*, 79, 115-134.

50. Myers, B., Chuang, Y., Tjandra, M., Chen, M., and Lee, C. *Floor control in a Highly Collaborative Co-Located Task*. Accessed: April 6, 2004. <http://www-2.cs.cmu.edu/~pebbles/papers/pebblesfloorcontrol.pdf>.
51. Oakley, I., Brewster, S., and Gray, P. (2001). Can You Feel the Force? An investigation of haptic collaboration in shared editors. In *Proceedings of Eurohaptics*. Birmingham, UK.
52. Patterson, J. (1990). Rendezvous: an architecture for synchronous multi-user applications. In *Proceedings of the 1990 ACM conference on Computer-supported cooperative work* (pp. 317-328). Los Angeles, California.
53. Pedersen, E., McCall, K., Moran, T., and Halasz, F. (1993). Tivoli: An electronic whiteboard for informal workgroup meetings. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 391-398). Amsterdam, The Netherlands: ACM Press.
54. Rensink, R.A. (2004). Visual Sensing Without Seeing. *Psychological Science*, 15 (1), 27-32.
55. Roseman, M., and Greenberg, S. (1996). Building Real Time Groupware with GroupKit, A Groupware Toolkit. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 3 (1), 66-106.
56. Rosenberg, L. (1993). The use of Virtual Fixtures to enhance telemanipulation with time delay. In *Proceedings of Advances in Robotics, Mechatronics, and Haptic Interfaces, ASME* (pp. 29-36).
57. Rosenberg, L., and Brave, S. (1996). Using force feedback to enhance human performance in graphical user interfaces. In *Conference companion on Human Factors in Computing Systems* (pp. 291-292). Vancouver, British Columbia: ACM Press.
58. Sacks, H., Schegloff, E., and Jefferson, G. (1974). A Simplest Systematics for the Organization of Turn-Taking for Conversation. *Language*, 50 (4), 696-735.
59. Sallnas, E., Rassmus-Grohn, K., and Sjostrom, C. (2000). Supporting presence in collaborative environments by haptic force feedback. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7 (4), 461-476.
60. Takane, Y., Young, F.W., and de Leeuw, J. (1977). Nonmetric individual differences multidimensional scaling: An alternating least squares method with optimal scaling features. *Psychometrika*, 42, 7-67.
61. Tan, H., Durlach, N., Reed, C., and Rabinowitz, W (1999). Information transmission with a multifinger tactual display. *Perception & Psychophysics*, 61 (6), 993-1008.
62. Tse, E., and Greenberg S. (2004). Rapidly Prototyping Single Display Groupware through the SDGToolkit. In *Proceedings of the Fifth Australasian User Interface Conference* (pp. 101-110).
63. Wagner, C.R., Stylopoulos, N., and Howe, R. D (2002). The role of force feedback in surgery: analysis of blunt dissection. In *10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*.
64. Ward, L. (1977). Multidimensional Scaling of the Molar Physical Environment. *Multivariate Behavioral Research*, 12, 23-42.
65. Young, F., and Hamer, R. (1987). *Multidimensional Scaling: History, Theory, and Applications*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, Inc.

Appendix A: **Study 1 Materials**

On the following pages material related to Study 1 is shown, including the consent form signed by subjects, the instructions subjects read, pre- and post-study questions subjects were asked, the results from the MDS analysis of subject data, and subjects' rating of the noticeability and pleasantness of stimuli.

A.1 Consent Form

Subjects were required to read and sign the following consent form before participating in Study 1.

PARTICIPANT'S COPY CONSENT FORM

Department of Computer Science
2366 Main Mall
Vancouver, B.C. Canada V6T 1Z4
tel: (604) 822-9289
fax: (604) 822-5485

Project Title: Physical and Multimodal User Interfaces – Usability and Psychophysics

Principal Investigator: K. MacLean, tel. 604-822-8169

Co-Investigator: Andrew Chan, tel. 604-318-0039 MSc. Candidate,
Department of Computer Science

The purpose of this study is to examine changes in human performance as well as interaction when using different types of haptic (touch sense) displays.

The task you will perform has been programmed on a computer. You will be asked to respond to each successive task by pressing buttons on a keyboard and interacting with a haptic display. You may be asked to wear headphones for the delivery of auditory input. Please tell the experimenter if you find the auditory stimulus level uncomfortable, and it will be adjusted. You will be asked to answer questions in a questionnaire and an interview as part of the experiment. The interview will be audio taped for later analysis.

You will receive specific instructions for the task before you begin (i.e., which buttons to press for a given response). If you are not sure about any instructions, do not hesitate to ask.

TIME COMMITMENT: 1 hour session

HONORARIUM: You will receive \$10 for participating in this experiment.

CONFIDENTIALITY: *Your identity will be confidential: you will not be identified by name in any study reports. Experimental results will be stored in a secure Computer Science account accessible only to the experimenters. The*

anonymous data from the experiment will be used in a Master's thesis and possibly in a scholarly publication.

You understand that the experimenter will ANSWER ANY QUESTIONS you have about the instructions or the procedures of this study. After participating, the experimenter will answer any questions you have about this study.

You understand that you have the RIGHT TO REFUSE to participate or to withdraw from the study at any time without penalty of any form.

You hereby CONSENT to participate in this study and acknowledge RECEIPT of a copy of the consent form:

NAME _____
(please print)

SIGNATURE _____ DATE _____

If you have any concerns regarding your treatment as a research subject you may contact the Research Subject Information Line in the UBC Office of Research Services at 604-822-8598.

A.2 Study 1 Instructions

Subjects read the following instructions on the computer screen during Study 1. The investigator did not verbally instruct subjects except to answer questions related to the study procedure.

Overview

In this study, you will be sorting different kinds of haptic (touch sense) stimuli based on how similar they feel. The stimuli will be delivered through the grey computer mouse attached to this workstation. Feel free to take breaks as needed, and if you experience any discomfort, please let the investigator know. At any time, you may withdraw from the study without penalty.

Detailed Instructions

Once you press the button below, a new screen will be shown. There will be a set of small, numbered tiles at the bottom, and a set of containers at the top left. The objective is to sort the tiles into the containers so that similar tiles are in the same container. Click with the left mouse button on any tile to play back the stimulus associated with it. Click with the right mouse button on a tile to move it; one click will pick up the tile, and a second click will put it down. You can play back a stimulus or move a tile as many times as you want.

Each container has a text field at the top. In this field, describe the kinds of tiles you are placing into that container. Make your description as detailed as you can.

The first time you perform this sort, you initially will be presented with two containers. Add more containers if you feel that the tiles fall into more than two categories. You can add and remove containers by pressing the + (plus) and - (minus) buttons in the lower left-hand corner. At least two containers and no more than fifteen can be used.

When you have finished sorting the tiles, press the "End This Sort" button in the lower right-hand corner. You will then be given a new set of tiles to sort, with a fixed number of containers in which to place them. You must place at least one tile in each container, and label the containers as before.

In total, you will sort five sets of haptic stimuli. Please take your time and try your best. We are interested in how well you sort the tiles, not how quickly you sort them. When you have completed the study, a message box will notify you.

Clicking the Help button at any time will display these instructions.

Thank you for participating in this study! Please put on the headphones and begin by pressing the button below.

A.3 Pre- and Post-Study Questions

Subjects were asked the following questions before and after the study respectively.

Pre-Study Questions

What is your name?

How old are you?

Which hand is your dominant hand?

Left Right

Which hand do you use most often to control a mouse?

Left Right

Do you have any previous experience with haptic devices (devices that communicate information through the sense of touch)?

Yes No

If yes, what kind of devices have you used, and for how long?

Post-Study Questions

Did you experience any fatigue during the experiment?


Do you have any thoughts / comments about the experiment?

A.4 MDS Graphs

For each of the groups discussed in Chapter 4, four different views of the 3D MDS solution are shown: a perspective projection that shows the three dimensions and a view looking down each of the dimensions. The coordinates for each stimulus in the solution are listed in a table following the plots. Interested readers should use a tool capable of plotting and animating the graphs, as it will greatly improve comprehension. SPSS 11.5 was used to analyze the graphs, but the latest version at the time of writing (version 12) no longer has this feature.

Overall

(N = 10)




Configuration derived in 3 dimensions

Stimulus Coordinates

Stimulus Number	Stimulus Name	Dimension		
		1	2	3
1	V1	-.7196	-.4066	-1.7551
2	V2	-.7782	.5300	1.5802
3	V3	-.7679	-.4243	1.6265
4	V4	-.2649	1.3324	.9019
5	V5	-.0899	-1.0119	1.3953
6	V6	.8518	1.3214	-.4428
7	V7	.8149	-1.3812	-.2132
8	V8	1.0477	1.1818	-.4839
9	V9	.8438	-1.3495	-.4207
10	V10	-.6376	-.4539	-1.7806
11	V11	.6557	.8011	1.4541
12	V12	.2581	-.7623	1.6353
13	V13	1.2516	1.1380	.1181
14	V14	1.1658	-1.2128	.1240
15	V15	1.3763	.9339	-.4051
16	V16	1.2488	-1.0776	-.5140
17	V17	1.4077	.8491	-.5118
18	V18	1.2547	-1.0606	-.5207
19	V19	-1.1520	1.0046	.7136
20	V20	-.8884	-.6835	1.3804
21	V21	-1.1577	1.1682	-.3353
22	V22	-1.3762	-.9007	-.1917
23	V23	-.9255	1.1322	-.8906
24	V24	-1.2746	-.8464	-.7165
25	V25	-.9317	1.0799	-.9655
26	V26	-1.2128	-.9013	-.7819

Male

(N = 6)



Configuration derived in 3 dimensions


Stimulus Coordinates

Stimulus Number	Stimulus Name	Dimension		
		1	2	3
1	V1	.2202	.3630	-2.0221
2	V2	-.6278	1.3473	.9415
3	V3	.6451	.8340	1.5502
4	V4	-1.3365	.8004	-.0697
5	V5	1.2112	.5924	.9610
6	V6	-1.3882	-.6760	.3325
7	V7	1.1877	-.6999	-.8664
8	V8	-1.2430	-.9286	.2937

9	V9	1.0785	-.7434	-1.0281
10	V10	.2329	.3260	-2.0203
11	V11	-.7900	-.7408	1.4562
12	V12	1.0515	-.7974	1.1357
13	V13	-1.0162	-1.1373	.8119
14	V14	.9924	-1.3015	-.3232
15	V15	-.8965	-1.3140	.7014
16	V16	.7183	-1.3533	-.9008
17	V17	-.8471	-1.3589	.6967
18	V18	.7261	-1.3437	-.9088
19	V19	-.9833	1.3017	-.3391
20	V20	.8984	.7578	1.3877
21	V21	-1.0267	1.0892	-.7919
22	V22	1.1164	1.1130	.5645
23	V23	-1.1382	.8356	-.9321
24	V24	1.1532	1.1356	.2937
25	V25	-1.1469	.8238	-.9363
26	V26	1.2087	1.0749	.0120

Female

(N = 4)




Configuration derived in 3 dimensions

Stimulus Coordinates

Stimulus Number	Stimulus Name	Dimension		
		1	2	3
1	V1	.3678	-1.8429	-.2072
2	V2	-.9249	1.1383	-.9703
3	V3	-.9526	1.2901	-.6437
4	V4	.0160	.6697	-1.5986
5	V5	.3743	1.7039	-.2064
6	V6	.9180	-.9912	-1.0163
7	V7	1.0316	.7789	1.0546
8	V8	1.0773	-.8110	-1.0805
9	V9	.7715	.8435	1.3688
10	V10	.3790	-1.8432	-.2880
11	V11	.3267	.8398	-1.5697
12	V12	-.0686	1.5107	-.9738
13	V13	1.0051	-.4926	-1.3693
14	V14	.8249	.9474	1.3315
15	V15	1.4241	-.7862	-.1646
16	V16	.9076	.5821	1.4706
17	V17	1.4699	-.6145	.3107
18	V18	.9360	.5753	1.4372
19	V19	-1.3875	.0756	-.8379
20	V20	-1.3305	.7401	-.6063
21	V21	-1.4628	-.7032	.1307
22	V22	-1.4502	-.2174	.5662
23	V23	-1.0857	-1.1166	.6178
24	V24	-1.0770	-.7934	.9932
25	V25	-1.0029	-.9302	1.0504
26	V26	-1.0870	-.5535	1.2013

Left-handed

(N = 3)



Configuration derived in 3 dimensions

Stimulus Coordinates

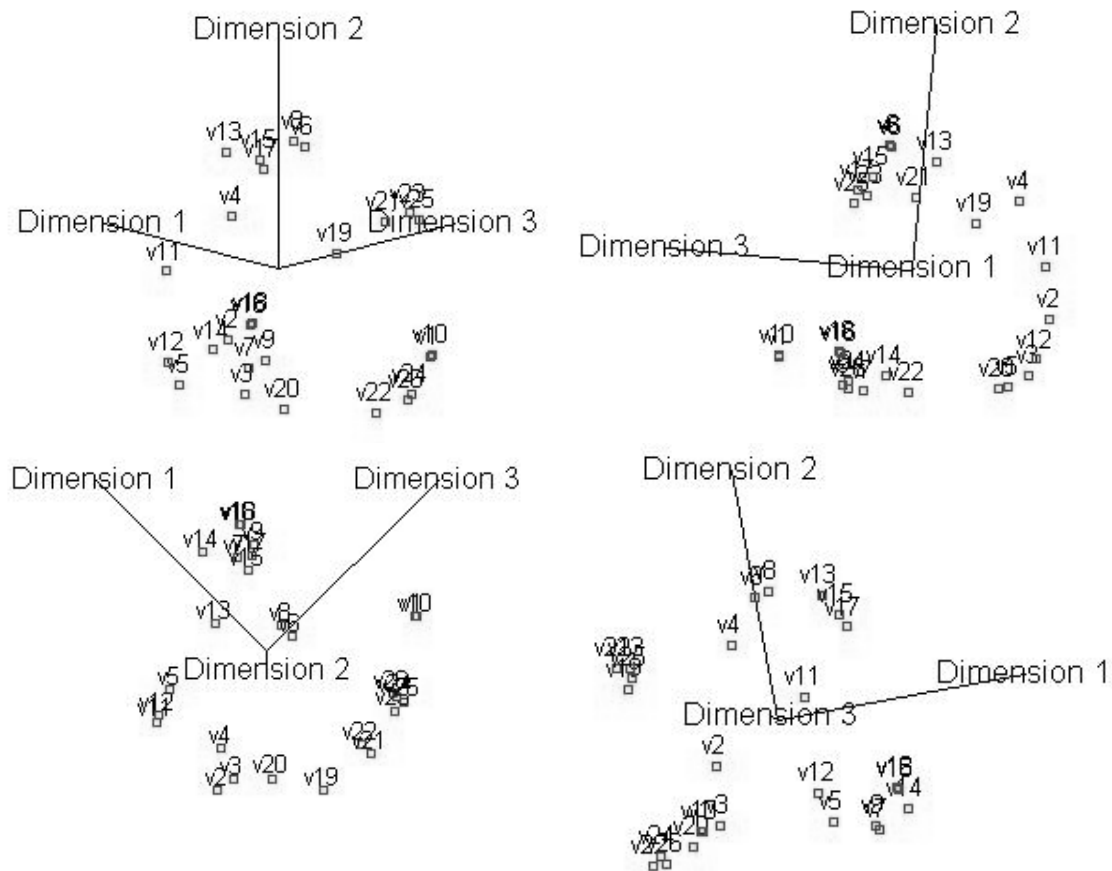
		Dimension		
Stimulus Number	Stimulus Name	1	2	3
1	V1	-.5040	1.3319	-1.3669
2	V2	-.8535	-.7829	-1.2517
3	V3	.8128	-1.5155	-.3900
4	V4	-1.1577	-.3866	.9986
5	V5	1.5715	-.0582	.4526
6	V6	-1.0919	-.0575	1.1936
7	V7	1.5302	.5597	.2600
8	V8	-1.1079	.0908	1.2327
9	V9	1.2311	1.0727	.0978
10	V10	-.3970	1.3882	-1.3400
11	V11	-.8882	.0953	1.4541
12	V12	1.6019	.1462	.5007

Stimulus Coordinates

Stimulus Number	Stimulus Name	Dimension		
		1	2	3
1	V1	.7841	.9692	-1.3296
2	V2	-.4954	.5482	1.6606
3	V3	.3670	.3548	1.7524
4	V4	-1.2265	.7157	.8603
5	V5	.9630	.1309	1.3981
6	V6	-1.3695	-.7095	-.6703
7	V7	1.1771	-1.0891	-.4065
8	V8	-1.1784	-.9996	-.7153
9	V9	1.2120	-1.0451	-.5080
10	V10	.8114	.8634	-1.3866
11	V11	-.8012	-.3923	1.5286
12	V12	.6255	-.3160	1.6458
13	V13	-1.3153	-1.0957	-.0630
14	V14	1.0207	-1.3616	.1781
15	V15	-1.1521	-1.1645	-.5665
16	V16	.9321	-1.3894	-.4747
17	V17	-1.0746	-1.2114	-.6493
18	V18	.9206	-1.3843	-.4871
19	V19	-.9656	1.1484	.7883
20	V20	.7813	.5070	1.4932
21	V21	-1.0603	1.2942	-.3114
22	V22	1.1108	1.2337	-.2047
23	V23	-1.0384	1.0742	-.8388
24	V24	.9587	1.1418	-.8449
25	V25	-.9669	1.0809	-.9316
26	V26	.9799	1.0959	-.9170

Novices

(N = 6)



Configuration derived in 3 dimensions


Stimulus Coordinates

Stimulus Number	Stimulus Name	Dimension		
		1	2	3
1	V1	-.7017	-.8036	1.5473
2	V2	-.4867	-.1086	-1.7377
3	V3	-.5549	-.7312	-1.5379
4	V4	-.1035	1.1190	-1.2296
5	V5	.6305	-.8985	-1.3038
6	V6	.2472	1.5630	.4235
7	V7	1.1251	-1.0933	.4555
8	V8	.4085	1.5802	.4400
9	V9	1.0956	-1.0473	.7369
10	V10	-.7188	-.7856	1.5576

11	V11	.5671	.4393	-1.6241
12	V12	.5309	-.5708	-1.6171
13	V13	.9426	1.4524	-.1503
14	V14	1.4548	-.9157	.2185
15	V15	1.0905	1.2159	.6114
16	V16	1.3934	-.7175	.7951
17	V17	1.1584	1.0703	.7695
18	V18	1.3862	-.6945	.8121
19	V19	-1.2499	.8417	-.7174
20	V20	-.8667	-.9048	-1.1916
21	V21	-1.3267	1.0726	.0626
22	V22	-1.3099	-1.0433	-.0882
23	V23	-1.1787	1.0431	.6569
24	V24	-1.2069	-.9737	.6550
25	V25	-1.1769	.9464	.8073
26	V26	-1.1492	-1.0553	.6485

Experts

(N = 4)




Stimulus Coordinates

Stimulus Number	Stimulus Name	Dimension		
		1	2	3
1	V1	-.4386	-.8231	-1.7677
2	V2	-.3448	-.8690	1.6080
3	V3	1.0341	-1.1496	.6411
4	V4	-1.0432	-.4261	1.2580
5	V5	1.2651	-.7417	-.6822
6	V6	-1.2829	1.0019	-.0424
7	V7	1.3476	.8356	.3454
8	V8	-1.0327	1.1720	-.4021
9	V9	1.3752	.8111	.1469
10	V10	-.3673	-.7695	-1.8045
11	V11	-1.2738	.9746	-.2513
12	V12	1.4381	.7144	.3878
13	V13	-1.1571	1.1341	-.3859
14	V14	1.1837	1.1629	.2616
15	V15	-1.0931	1.1769	-.5115
16	V16	1.0860	1.2910	.2531
17	V17	-1.0831	1.1727	-.5600
18	V18	1.0850	1.2927	.2584
19	V19	-.5171	-.9635	1.4675
20	V20	1.2524	-1.0573	.1421
21	V21	-.7313	-.9642	1.3046
22	V22	.4296	-1.1030	-1.3429
23	V23	-.8627	-.8577	1.2845
24	V24	.3043	-1.1035	-1.4171
25	V25	-.8631	-.8619	1.2787
26	V26	.2895	-1.0502	-1.4700

Weird removed

(N = 7)



Configuration derived in 3 dimensions

Stimulus Coordinates

Dimension

Stimulus Number	Stimulus Name	1	2	3
1	V1	1.1459	-.7297	-1.3671
2	V2	-.3770	-.6677	1.5953
3	V3	1.4653	-.5435	.5993
4	V4	-1.2912	-.8086	.4880
5	V5	.6705	1.2186	.8731
6	V6	-1.3301	-.5658	-.7880
7	V7	.1106	1.5927	.5636
8	V8	-1.2706	-.3964	-.9910
9	V9	.0833	1.6429	.3992
10	V10	1.1234	-.6792	-1.4140
11	V11	-1.4774	-.3790	-.6370

12	V12	.2060	1.5723	.6455
13	V13	-1.0466	-.2386	-1.3878
14	V14	-.0790	1.7385	-.1929
15	V15	-.7460	.0008	-1.6233
16	V16	-.2116	1.6872	-.4494
17	V17	-.7327	.0280	-1.6294
18	V18	-.2220	1.6761	-.4779
19	V19	-.3568	-.8647	1.4792
20	V20	1.5768	-.3143	.5316
21	V21	-.5413	-.9109	1.4100
22	V22	1.5652	-.7118	-.0317
23	V23	-.6898	-.9766	1.2523
24	V24	1.5376	-.7366	-.0773
25	V25	-.6770	-.9714	1.2671
26	V26	1.5645	-.6624	-.0374

A.5 Likert Scale Responses

Subjects were asked to rate each of the 26 stimuli presented in Study 1 on a 5-point Likert scale, based on how noticeable and how pleasant they felt. Their ratings are summarized below, along with the average rating. Due to a bug in the logging code, subject responses for v26 were not written to the files.

How Noticeable

(1 = barely noticeable, 5 = very noticeable)

Stimulus	Subject										Average
	1	2	3	4	5	6	7	8	9	10	
v1	5	4	5	5	5	4	4	3	5	4	4.40
v2	1	1	2	1	2	1	1	1	1	1	1.20
v3	1	1	2	1	2	1	1	1	1	1	1.20
v4	2	3	4	4	3	2	3	3	2	3	2.90
v5	3	2	3	2	3	2	3	2	4	3	2.70
v6	2	4	5	4	4	4	4	3	4	4	3.80
v7	3	4	5	4	4	4	4	4	4	4	4.00
v8	4	4	5	5	4	5	4	5	5	4	4.50
v9	4	4	5	5	4	5	5	5	5	5	4.70
v10	4	5	5	3	5	3	3	3	5	5	4.10
v11	1	1	2	4	2	1	3	1	2	1	1.80
v12	1	2	2	1	2	1	3	1	2	2	1.70
v13	2	2	4	2	3	4	4	2	4	2	2.90
v14	2	3	5	3	3	2	3	3	4	3	3.10
v15	4	4	5	3	4	5	4	4	5	4	4.20
v16	4	4	5	3	4	4	4	4	5	5	4.20
v17	5	4	5	4	4	5	5	5	5	5	4.70
v18	5	4	5	4	5	4	5	3	5	5	4.50
v19	2	2	3	4	3	1	2	2	2	2	2.30
v20	1	2	2	2	3	2	2	2	4	2	2.20
v21	2	2	4	3	4	4	3	2	4	2	3.00
v22	2	2	4	4	4	4	4	3	4	3	3.40
v23	3	4	5	4	4	5	4	3	4	3	3.90
v24	4	3	5	4	4	5	5	5	5	3	4.30
v25	3	4	5	5	4	4	4	5	5	5	4.40

How Pleasant

(1 = very unpleasant, 5 = very pleasant)

Stimulus	Subject										Average
	1	2	3	4	5	6	7	8	9	10	
v1	1	3	4	1	5	2	2	4	3	4	2.90
v2	4	4	3	1	4	5	4	5	4	3	3.70
v3	5	5	3	1	5	5	4	5	4	3	4.00
v4	3	4	3	4	3	5	4	3	3	3	3.50
v5	3	4	3	2	4	4	4	3	3	4	3.40
v6	4	4	3	4	3	4	3	3	2	4	3.40
v7	4	3	3	4	4	2	3	3	2	3	3.10
v8	1	3	3	5	4	4	3	2	2	2	2.90
v9	3	3	3	4	4	2	2	2	1	3	2.70
v10	3	3	3	2	5	4	3	4	3	4	3.40
v11	4	4	3	4	4	4	3	4	3	4	3.70
v12	5	4	3	3	5	5	4	3	3	2	3.70
v13	4	4	3	4	4	4	3	3	3	3	3.50
v14	2	3	3	5	4	4	3	2	2	3	3.10
v15	3	3	3	3	4	4	3	2	2	2	2.90
v16	3	3	2	3	4	2	3	2	2	3	2.70
v17	2	3	2	4	4	1	3	1	1	1	2.20
v18	2	3	2	4	4	3	3	2	1	2	2.60
v19	3	4	3	4	2	5	4	4	4	3	3.60
v20	2	4	3	3	2	4	4	4	4	4	3.40
v21	3	4	3	4	1	4	4	5	3	4	3.50
v22	3	4	3	5	2	3	4	4	3	4	3.50
v23	2	3	3	4	2	3	3	4	3	4	3.10
v24	2	3	3	4	3	4	3	2	2	3	2.90
v25	4	3	3	4	2	2	3	4	2	4	3.10

Appendix B: **Study 2 Materials**

On the following pages material related to Study 2 is shown, including the consent form signed by subjects, the instructions subjects read, and interview questions subjects were asked after the study.

B.1 Consent Form

Subjects were required to read and sign the following consent form before participating in Study 2.

PARTICIPANT'S COPY CONSENT FORM

Department of Computer Science
2366 Main Mall
Vancouver, B.C. Canada V6T 1Z4
tel: (604) 822-9289
tax: (604) 822-5485

Project Title: Physical and Multimodal User Interfaces – Usability and Psychophysics

Principal Investigator: K. MacLean, tel. 604-822-8169

Co-Investigator: Andrew Chan, tel. 604-318-0039 MSc. Candidate, Department of Computer Science

The purpose of this study is to examine changes in human performance as well as interaction when using different types of haptic (touch sense) displays.

The task you will perform has been programmed on a computer. You will be asked to respond to each successive task by pressing buttons on a keyboard and interacting with a haptic display. You must wear headphones for the delivery of auditory input. Please tell the experimenter if you find the auditory stimulus level uncomfortable, and it will be adjusted. You will be asked to answer questions in an interview as part of the experiment. The interview will be audio taped for later analysis.

You will receive specific instructions for the task before you begin (i.e., which buttons to press for a given response). If you are not sure about any instructions, do not hesitate to ask.

TIME COMMITMENT: 60 to 90 minutes

HONORARIUM: You will receive \$10 for participating in this experiment. In addition, the four participants with the best performance in this experiment will receive an additional \$10.

CONFIDENTIALITY: *Your identity will be confidential: you will not be identified by name in any study reports. Experimental results will be stored in a secure Computer Science*

account accessible only to the experimenters. The anonymous data from the experiment will be used in a Master's thesis and possibly in a scholarly publication.

You understand that the experimenter will ANSWER ANY QUESTIONS you have about the instructions or the procedures of this study. After participating, the experimenter will answer any questions you have about this study.

You understand that you have the RIGHT TO REFUSE to participate or to withdraw from the study at any time without penalty of any form.

You hereby CONSENT to participate in this study and acknowledge RECEIPT of a copy of the consent form.

If you are one of the four participants who have the best performance in this experiment, you will be contacted by email after the conclusion of the experiment. You will have three weeks from that point to collect the \$10 at a time of your choosing.

If you have any concerns regarding your treatment as a research subject you may contact the Research Subject Information Line in the UBC Office of Research Services at 604-822-8598.

You hereby CONSENT to participate in this study and acknowledge RECEIPT of a copy of the consent form:

NAME _____

(please print)

SIGNATURE _____

DATE _____

B.2 Study 2 Instructions

In Study 2, subjects read instructions on-screen and in a booklet. Lengthier instructions were put in the booklet to facilitate reading, and on-screen prompts directed subjects to read specific pages at certain points during the study. The investigator did not verbally instruct subjects except to answer questions related to the study procedure.

Overview (*shown on-screen at the beginning of the study*)

This study is divided into two phases. In the first phase, you will be introduced to a set of haptic (touch-sense) stimuli delivered through a haptic mouse. Each stimulus has been given a meaning; your task is to learn the meanings. This will be explained in detail. In the second phase, you will be tested on your ability to recall the meanings under different conditions.

You must wear headphones during this study. Among other things, background noise will be played through the headphones to block out any external noise. Please notify the investigator if the volume level is too loud and it will be adjusted.

The study will take between 60 and 90 minutes to complete. You will be given the opportunity to take rest breaks at several points during the study.

Specific instructions for each section of the study have been printed and are in a small booklet. Please read page 1 of the instructions now.

If you have any questions, ask the investigator. Otherwise, please put on the headphones and press the "Next" button to begin.

Phase One Instructions (*page 1 of the instruction booklet*)

This is the "exploration phase" in the study. Your task is to learn the meanings associated with seven haptic stimuli as quickly as possible.

On the next screen, seven haptic stimuli are listed in a grid. Each stimulus represents different emotional states an individual may experience during the day. Play each stimulus by pressing the "Play" button to the left of its label. You may play the stimuli as many times as you like.

When you feel that you have learned the meanings of the stimuli, you may proceed to a short evaluation. A number of stimuli will be presented in random order and you will be asked to identify them. When you are able to identify more than 90% of the stimuli, you will proceed to the next phase of the study. Otherwise, you will be given additional time to learn the meanings of the stimuli. Further instructions will be given before the evaluation begins.

The time taken to learn the meanings of the stimuli and successfully complete the evaluation will be recorded. Try to learn as quickly as possible, but make sure you learn the meanings well. You will need this information in phase two of the study.

Evaluation Instructions (*page 2 of the booklet; read once the subject has finished exploring the stimuli*)

We will now evaluate your knowledge of the seven stimuli.

On the next screen will be a listing of the seven stimuli. After a few seconds' delay, one of the seven stimuli will be played. Choose which stimulus you think it is, then press the "Next" button. A total of 21 stimuli will be presented. When you are able to identify 19 stimuli correctly, you will proceed to the second phase of the study. Otherwise, you will be returned to the previous screen so that you can spend more time learning the stimuli.

Phase Two Instructions (*read once the subject has passed the evaluation described above*)

In the second phase of the experiment, your ability to recall the meanings of the stimuli you learned will be tested under different conditions. Before each condition, you will be given a chance to review the haptic stimuli. You will also receive specific instructions and be given time to familiarize yourself with the interface.

Your performance will be measured in each condition. At the conclusion of the study, the four participants with the best performance will receive an extra \$10. The instructions for each condition will tell you how to maximize your performance.

Press "Next" to review the haptic stimuli.

Condition 1 Instructions (*page 3 of the booklet - the three conditions were counterbalanced; the instructions for the haptic condition are shown*)

In the first condition you will be identifying haptic stimuli, much as you did in the first phase. However, the stimuli will now be played continuously and in pairs. The first stimulus in the pair will be played for a period of time, directly followed by the second stimulus. As soon as you feel the change from the first stimulus to the second, press the space bar. A dialog box will open, and you will be asked to identify the second stimulus. This dialog box will also appear if you have not pressed the space bar within 10 seconds of the second stimulus playing. After you make your selection, a new pair of stimuli will begin playing.

The first stimulus in the new pair of stimuli may be same as the second stimulus from the previous pair. Do not be misled by this. As well, try not to press the space bar before the second stimulus starts, as this will slow you down.

Your performance in this condition will depend on how quickly you notice the second stimuli, and how accurately you identify them.

Condition 2 Instructions (*page 4 of the booklet - the three conditions were counterbalanced; the instructions for the haptic+visual condition are shown*)

In the second condition, you will be identifying haptic stimuli and solving picture puzzles.

The haptic stimuli will be presented in exactly the same way as in the first condition. They will now be played continuously and in pairs. The first stimulus in the pair will be played for a period of time, directly followed by the second stimulus. As soon as you feel the change from the first stimulus to the second, press the space bar. A dialog box will open, and you will be asked to identify the second stimulus. This dialog box will also appear if you have not pressed the space bar within 10 seconds of the second stimulus playing. After you make your selection, a new pair of stimuli will begin playing.

The first stimulus in the new pair of stimuli may be same as the second stimulus from the previous pair. Do not be misled by this. As well, try not to press the space bar before the second stimulus starts, as this will slow you down.

Each picture puzzle consists of an image that has been subdivided into pieces and scrambled. Rearrange the puzzle pieces to restore the original image. The original image is provided as a guide. Move a puzzle piece by clicking with the left mouse button and dragging it to a new location; the pieces will be swapped. When you finish a puzzle, a new one will be shown.

Your performance in this condition will depend on how well you perform these two tasks simultaneously. Try to identify the haptic stimuli as quickly and accurately as possible, while completing as many puzzles as you can.

Condition 3 Instructions (*page 5 of the booklet - the three conditions were counterbalanced; the instructions for the haptic+visual+audio condition are shown*)

In the third condition, you will be identifying haptic stimuli, solving picture puzzles, and listening for an audio keyword to be spoken.

The haptic stimuli will be presented in exactly the same way as in the first two conditions. They will be played continuously and in pairs. The first stimulus in the pair will be played for a period of time, directly followed by the second stimulus. As soon as you feel the change from the first stimulus to the second, press the space bar. A dialog box will open, and you will be asked to identify the second stimulus. This dialog box will also appear if you have not pressed the space bar within 10 seconds of the second stimulus playing. After you make your selection, a new pair of stimuli will begin playing.

The first stimulus in the new pair of stimuli may be same as the second stimulus from the previous pair. Do not be misled by this. As well, try not to press the space bar before the second stimulus starts, as this will slow you down.

The picture puzzles will also be presented as they were previously. Each picture puzzle consists of an image that has been subdivided into pieces and scrambled. Rearrange the

puzzle pieces to restore the original image. The original image is provided as a guide. Move a puzzle piece by clicking with the left mouse button and dragging it to a new location; the pieces will be swapped. When you finish a puzzle, a new one will be shown.

Different colours will be spoken at selected intervals. When you hear the word "blue" spoken, press the "b" key on the keyboard as quickly as possible.

Your performance in this condition will depend on how well you perform all three of the tasks simultaneously. Try to identify the haptic stimuli as quickly and accurately as possible, while completing as many puzzles as you can, and indicating when a keyword has been spoken.

B.3 Post-Study Interview Questions

After the study, subjects were asked the following questions.

In this study, you were exposed to seven haptic stimuli. These stimuli were named: Awake, Asleep, Low Stress, Medium Stress, High Stress, Bored, and Really Bored.

- Which of the stimuli did you find most noticeable?
- Which of the stimuli did you find least noticeable?

In the second part of the study, you felt the stimuli being played continuously for a fair length of time.

- Which of the stimuli did you find most pleasant to feel? Why?
- Which of the stimuli did you find least pleasant to feel? Why?

There were three conditions presented in the second part of the study: one where you only had to identify haptic stimuli; one where you identified haptic stimuli and solved picture puzzles; and one where you identified haptic stimuli, solved picture puzzles, and listened for a keyword.

- Which condition was the easiest to complete?
- Which condition was the next easiest? How much more difficult was it than the easiest one? What made it more difficult?
- Which condition was the most difficult? How much more difficult was it than the previous one? What made it more difficult?

Did you experience any fatigue during the experiment?

Do you have any thoughts / comments about the experiment?

Appendix C:

Study 3 Materials

On the following pages material related to Study 3 is shown, including the consent form signed by subjects, the instructions subjects read, questionnaires subjects filled out at various stages of the study, subjects' responses to the Likert scale questions on the questionnaires, and the tasks subjects completed with reference solutions and a scoring guide.

C.1 Consent Form

Subjects were required to read and sign the following consent form before participating in Study 3.

PARTICIPANT'S COPY CONSENT FORM

Department of Computer Science
2366 Main Mall
Vancouver, B.C. Canada V6T 1Z4
tel: (604) 822-9289
tax: (604) 822-5485

Project Title: Physical and Multimodal User Interfaces – Usability and Psychophysics

Principal Investigator: K. MacLean, tel. 604-822-8169

Co-Investigator: Andrew Chan, tel. 604-318-0039 MSc. Candidate, Department of Computer Science

The purpose of this study is to examine changes in human performance as well as interaction when using different types of haptic (touch sense) displays.

The task you will perform has been programmed on a computer. You will be asked to respond to each successive task by pressing buttons on a keyboard and interacting with a haptic display. You must wear headphones for the delivery of auditory input. Please tell the experimenter if you find the auditory stimulus level uncomfortable, and it will be adjusted. You will be asked to answer questions in an interview as part of the experiment. The interview will be audio taped for later analysis.

You will receive specific instructions for the task before you begin (i.e., which buttons to press for a given response). If you are not sure about any instructions, do not hesitate to ask.

TIME COMMITMENT: 2 ½ to 3 hours

HONORARIUM: You will receive \$25 for participating in this experiment. In addition, the top 1/3 groups with the best performance in this experiment will receive an additional \$40.

CONFIDENTIALITY: *Your identity will be confidential: you will not be identified by name in any study reports. Experimental results will be stored in a secure Computer Science account accessible only to the experimenters. The*

anonymous data from the experiment will be used in a Master's thesis and possibly in a scholarly publication.

If your group is eligible for the additional \$40, you will be contacted by email after the conclusion of the experiment. You will have three weeks from that point for any group member to collect the \$40 at a time of your choosing.

You understand that the experimenter will ANSWER ANY QUESTIONS you have about the instructions or the procedures of this study. After participating, the experimenter will answer any questions you have about this study.

You understand that you have the RIGHT TO REFUSE to participate or to withdraw from the study at any time without penalty of any form.

If you have any concerns regarding your treatment as a research subject you may contact the Research Subject Information Line in the UBC Office of Research Services at 604-822-8598.

You hereby CONSENT to participate in this study and acknowledge RECEIPT of a copy of the consent form:

NAME _____
(please print)

SIGNATURE _____

DATE _____

C.2 Study 3 Instructions

Subjects were given a binder of instructions and told when to read each section. The first two sections were identical for all subjects and used during the training portion of the study. When subjects were working together as a group, each subject received slightly different instructions to mimic real-world collaboration, where individuals share a set of knowledge, but also have their own priorities and objectives.

Using Microsoft Visio

(Subjects were given a demonstration of Microsoft Visio, then instructed to complete the exercises shown below. They were told that they did not have to read this tutorial unless they wished to clarify something that they did not understand)

Microsoft Visio is a powerful tool for creating diagrams. You will be working with a simplified version of Visio, and only need to understand a few basic concepts.

General Layout

The largest part of a Visio window is the drawing canvas where diagrams are created and modified. The canvas has a grid on it; each square is 1 foot x 1 foot. To the left of the drawing canvas is a shape stencil, which contains items of furniture you will need when working on the problems. There are three menus, the Plan, File, and Edit menus; the only menu you need is the Edit menu, which is described below. The Plan menu is not always visible.

Adding and Deleting Items

To add an item of furniture to the drawing canvas, click on the desired item in the shape stencil and drag it to the desired location on the canvas. An item of furniture can also be repositioned on the canvas by clicking on it and dragging it to the new location. To delete an item from the canvas, click on it and press the "Delete" key on the keyboard.

Rotating Items

When you click on an item, a green dot appears just outside of it, as in Figure 1. You can rotate an item by clicking on that dot and dragging the mouse in the direction in which you want the item rotated. When you move the cursor over the dot, a second, smaller dot appears, as shown in Figure 2. This dot defines the point about which the item is rotated. By default, this point is in the center of the item, so that the item rotates without changing location. However, you can change this point by clicking on the dot and moving it.

Note: It is easy to accidentally move the center of rotation, when you wanted to move the entire item. This occurs when you click on the center of the item (where the center-of-rotation dot is located) and drag it to a new location. If this occurs, the best thing to do is undo the action (choose Undo from the Edit menu).




Figure 1: Rotate Tool




Figure 2: Center of Rotation

Working with Groups

Visio also allows you to select groups of items and move or delete the entire group. This can be done by dragging a rectangle that completely encompasses all the items to be grouped together (see Figure 3). Items included in the group are highlighted, as shown in Figure 4. The group can be moved by dragging any item in the group. It can be rotated using the Rotate dot. To delete the group, press the "Delete" key. To undo the grouping, click anywhere outside the group.




Figure 3: Selecting a Group of Objects




Figure 4: The Selected Objects

The Edit Menu

In addition to these features, Visio has the standard Cut / Copy / Paste functions available under the Edit menu. The Undo / Redo features can also be found here.

Exercises

Exercise 1: Drag a Filing Cabinet from the stencil anywhere onto the drawing canvas. Move the 4' x 4' workstation to the right and place it underneath the printer. Delete the printer.

Exercise 2: Rotate the 4' x 4' workstation 90 degrees clockwise so that it looks like Figure 5. Then, move the pivot point from the center to the upper right-hand corner of the workstation. Observe what effect this has when you rotate the workstation.

Exercise 3: Select the 6' x 6' workstation and the bookshelf, and move them as a group to the right, so that they are immediately above the sofa.




Figure 5: Rotated workstation

Please show the investigator your solutions to the three exercises. If you have questions, please ask the investigator.

Using Haptics to Share Control

(Subjects read this section before proceeding to the training program where they could explore the different haptic icons)

For certain parts of the experiment, you will use the mouse to obtain and release control. The mouse you are using has two buttons located near your thumb. The following table shows the commands these buttons activate:



Command	Action
Gently Request Control	Press the front button once
Urgently Request Control	Press the front button twice (if you have already Gently Requested Control, you only need to press it once)
Take Control	Hold down the front button for two seconds, then release
Cancel a Request for Control	Press the rear button
Release Control	Press the rear button

As well, for certain parts of the experiment the *state* you are in (observing, in control, or waiting for control) will be communicated through your *haptic* (touch) sense. The mouse at your workstation will deliver a unique sensation for each of the states, so that you can tell what the current state is. There are also two signals to inform you when you have Gained Control and Lost Control of Visio. In total, there are eight signals, listed below:

State	Signal	Sensation Felt
Observing	You are neither In Control nor Requesting Control	None
Change of Control	You have just Gained Control	Gentle buzz followed by Strong buzz
	You have just Lost Control	Strong buzz followed by Gentle buzz
In Control	You are In Control, and no else has Requested Control from you	Gentle buzz (*)
	Someone has Gently Requested Control from you; you are In Control	Moderate buzz (*)
	Someone has Urgently Requested Control, or multiple people have Requested Control from you; you are In Control	Two strong buzzes (*)
Requesting Control	You have Gently Requested Control from the person In Control	One Tap (*)
	You have Urgently Requested Control from the person In Control	Two Taps (*)

All of the stimuli marked (*) are played periodically (every few seconds) while you are in that state.

Notice that your actions will not only influence the stimuli *you* feel, but also what the person *in control* feels. For example, when you gently request control, you feel a tapping sensation. The person in control will feel the gentle buzzing stimulus change to a moderate buzz.

To help you learn to identify these signals, a small training program has been written to demonstrate these signals. When you have finished reading these instructions, start the training program shown on the screen. Follow the instructions that are presented. When you are done, notify the investigator.

Warm-Up Exercise

(Before each condition, the group completed a warm-up exercise to familiarize themselves with the condition. Each subject received the four steps shown in a different order, and subjects completed a different ordering in each condition).

This brief exercise will allow you to practise obtaining control from one another while using Microsoft Visio. When you are instructed to begin, carefully follow the instructions below. If you Lose Control while carrying out a step, repeat the step again. Place a check mark beside each step when you have completed it.

Step 1: _____

Wait 5 seconds, then Request Control. When you Gain Control, move a 4x4 Workstation to the lower left-hand corner of the room and rotate it 180 degrees so that the chair portion (the red part) faces up. Release Control.

Step 2: _____

Wait 10 seconds, then Urgently Request Control. When you Gain Control, move a Printer and a Photocopier to the lower right-hand corner of the room. Release Control.

Step 3: _____

Wait 20 seconds.

Step 4: _____

Take Control. Move a 2x2 Chair and a 2' Circular Table to the center of the room. Release Control.

Task: Adding Workstations

(This task was presented with either the Haptic or Visual conditions in the study. Each subject only was shown two of the must-satisfy constraints and two of the try-to-satisfy constraints; the complete sets are shown here. The same constraints were used across tasks and conditions, but subjects were responsible for a different set each time)

Read the following description carefully, but do *not* begin working on it or discussing a solution with your friends until the investigator instructs you to begin. If the description is unclear, you may ask the investigator questions. Please try to ask them before you start working.


A start-up company is growing rapidly, and has hired your company to reorganize its office space. It has some *specific* constraints:

- 5 more workstations are needed immediately for new hires. Adding up to 10 would be better.
- The components of the snack bar (a fridge, a coffee desk, and a bookshelf filled with candy, all located at the upper-left corner of the office) should stay close together.
- Several staff greatly enjoy using the "Putting Green." While they want it to stay as large as possible, a smaller version would be tolerable. It must stay at least 6' x 6' in size.
- Having a social area is highly favoured, but its location is flexible.

In addition to those specific constraints, there are also *general* constraints that you *must* satisfy, and constraints that you should *try to* satisfy. The initial layout of the room may not satisfy these constraints. Each member of your group has been assigned two different *must-satisfy* constraints and two *try-to-satisfy* constraints. The constraints that you are responsible for are listed below.

Your *must-satisfy* constraints are as follows:

- There must be walkways to every item of furniture in the room. Each walkway must be at least 3' wide (as shown in the image on the right)
- Most items of furniture have at least one edge with a dark border (as shown on the right); the entire dark edge must be accessible by a walkway
- There must be 1 bookshelf for every 4 workstations in the room
- There must be 1 filing cabinet for every 3 workstations in the room
- Filing cabinets and bookshelves must not be placed in front of windows
- Furniture must not block doorways
- Furniture must not overlap other pieces of furniture or walls
- Bookshelves must be placed against walls or back-to-back



Your *try-to-satisfy* constraints are as follows:

- You should avoid having workstations that directly face one another and touch; staff find it distracting when someone is sitting across from them
- The entire length of whiteboards should be accessible by a walkway
- You should keep workstations a reasonable distance from high-noise areas (e.g. social areas, doorways, photocopiers) so that people can concentrate
- Windows should be accessible by a walkway
- There should be two routes from each workstation to each entrance, in case of a fire or similar emergency
- Small clusters of 2-4 workstations can be useful for staff working on projects together, but larger ones should be avoided as they tend to be noisy
- Bookshelves, workshelves, and filing cabinets should be placed against or close to walls, unless they are being used to partition a room
- You should try to re-use as much of the existing furniture as possible, besides any furniture the description explicitly says can be discarded

You will have 20 minutes to work with your coworkers to create an office layout that satisfies the different constraints listed. Remember that if your group does well on these tasks, you will receive a cash bonus! A points-based system will be used to evaluate each group:

- Groups will be given points for satisfying each *specific* constraint.
- They will be penalized heavily for every *must-satisfy* constraint that is not satisfied.
- Additional points will be awarded based on the extent to which *try-to-satisfy* constraints are met, and on the aesthetic appeal of the solution.

This means your group's success depends heavily on whether your individual *must-satisfy* and *try-to-satisfy* constraints are satisfied. Enforce your constraints as best as you can.

When you have finished reading these instructions, please let the investigator know. The investigator will tell you when you may begin working on the problem.

Task: Rearranging Workstations

(This task was presented with either the Haptic or Visual conditions in the study. Each subject only was shown two of the must-satisfy constraints and two of the try-to-satisfy constraints; the complete sets are shown here. The same constraints were used across tasks and conditions, but subjects were responsible for a different set each time)

Read the following description carefully, but do *not* begin working on it or discussing a solution with your friends until the investigator instructs you to begin. If the description is unclear, you may ask the investigator questions. Please try to ask them before you start working.


The QA department at a software company has packed workstations together to use space efficiently. However, the staff have complained that they don't have enough personal space, and that the noise level can be disruptive. They have hired your company to reorganize the room layout. Consultation with the staff yielded the following:

- The staff would like to have workstations grouped in side-by-side pairs, or even separate from one another
- The staff want to keep the social area and the complete access to windows
- There are five more workstations than the department needs, but it would be best to keep as many of the extra ones for summer interns to use

In addition to those specific constraints, there are also *general* constraints that you *must* satisfy, and constraints that you should *try to* satisfy. The initial layout of the room may not satisfy these constraints. Each member of your group has been assigned two different *must-satisfy* constraints and two *try-to-satisfy* constraints. The constraints that you are responsible for are listed below.

Your *must-satisfy* constraints are as follows:

- There must be walkways to every item of furniture in the room. Each walkway must be at least 3' wide (as shown in the image on the right)
- Most items of furniture have at least one edge with a dark border (as shown on the right); the entire edge must be accessible by a walkway
- There must be 1 bookshelf for every 4 workstations in the room
- There must be 1 filing cabinet for every 3 workstations in the room
- Filing cabinets and bookshelves must not be placed in front of windows
- Furniture must not block doorways



- Furniture must not overlap other pieces of furniture or walls
- Bookshelves must be placed against walls or back-to-back

Your *try-to-satisfy* constraints are as follows:

- You should avoid having workstations that directly face one another and touch; staff find it distracting when someone is sitting across from them
- The entire length of whiteboards should be accessible by a walkway
- You should keep workstations a reasonable distance from high-noise areas (e.g. social areas, doorways, photocopiers) so that people can concentrate
- Windows should be accessible by a walkway
- There should be two routes from each workstation to each entrance, in case of a fire or similar emergency
- Small clusters of 2-4 workstations can be useful to staff working on projects together, but larger ones should be avoided as they tend to be noisy
- Bookshelves, Workshelves, and Filing Cabinets should be placed against or close to walls, unless they are being used to partition a room
- You should try to re-use as much of the existing furniture as possible, besides any furniture the description explicitly says can be discarded

You will have 20 minutes to work with your coworkers to create an office layout that satisfies the different constraints listed. Remember that if your group does well on these tasks, you will receive a cash bonus! A points-based system will be used to evaluate each group:

- Groups will be given points for satisfying each *specific* constraint.
- They will be penalized heavily for every *must-satisfy* constraint that is not satisfied.
- Additional points will be awarded based on the extent to which *try-to-satisfy* constraints are met, and on the aesthetic appeal of the solution.

This means your group's success depends heavily on whether your individual *must-satisfy* and *try-to-satisfy* constraints are satisfied. Enforce your constraints as best as you can.

When you have finished reading these instructions, please let the investigator know. The investigator will tell you when you may begin working on the problem.

Task: Replacing Workstations

(This task was presented with the Haptic+Visual condition in the study. Each subject only was shown two of the must-satisfy constraints and two of the try-to-satisfy constraints; the complete sets are shown here. The same constraints were used across tasks and conditions, but subjects were responsible for a different set each time)

Read the following description carefully, but do *not* begin working on it or discussing a solution with your friends until the investigator instructs you to begin. If the description is unclear, you may ask the investigator questions. Please try to ask them before you start

working.


Thanks to new funding, a research lab is discarding their collection of workstations and purchasing new workstations and equipment. They have hired your company to figure out how to arrange the new furniture. The lab has the following requirements:

- 6'x6' L-shaped workstations are strongly preferred, or at least 5' or 6' long workstations, instead of the current 4'x4' workstations. 15 workstations are needed.
- A prototyping shop measuring 16' x 10' will be created. Power tools and machinery will be put in it, but furniture must not be placed inside it.
- The prototyping shop must be placed against a wall, and it must have two entrances.
- It would be best to isolate the prototyping shop from the workstations in the lab.

In addition to those specific constraints, there are also *general* constraints that you *must* satisfy, and constraints that you should *try to* satisfy. The initial layout of the room may not satisfy these constraints. Each member of your group has been assigned two different *must-satisfy* constraints and two *try-to-satisfy* constraints. The constraints that you are responsible for are listed below.

Your *must-satisfy* constraints are as follows:

- There must be walkways to every item of furniture in the room. Each walkway must be at least 3' wide (as shown in the image on the right)
- Most items of furniture have at least one edge with a dark border (as shown on the right); the entire edge must be accessible by a walkway
- There must be 1 bookshelf for every 4 workstations in the room
- There must be 1 filing cabinet for every 3 workstations in the room
- Filing cabinets and bookshelves must not be placed in front of windows
- Furniture must not block doorways
- Furniture must not overlap other pieces of furniture or walls
- Bookshelves must be placed against walls or back-to-back



Your *try-to-satisfy* constraints are as follows:

- You should avoid having workstations that directly face one another and touch; staff find it distracting when someone is sitting across from them
- The entire length of whiteboards should be accessible by a walkway
- You should keep workstations a reasonable distance from high-noise areas (e.g. social areas, doorways, photocopiers) so that people can concentrate
- Windows should be accessible by a walkway
- There should be two routes from each workstation to each entrance, in case of a fire or similar emergency

- Small clusters of 2-4 workstations can be useful to staff working on projects together, but larger ones should be avoided as they tend to be noisy
- Bookshelves, workshelves, and filing cabinets should be placed against or close to walls, unless they are being used to partition a room
- You should try to re-use as much of the existing furniture as possible, besides any furniture the description explicitly says can be discarded

You will have 20 minutes to work with your coworkers to create an office layout that satisfies the different constraints listed. Remember that if your group does well on these tasks, you will receive a cash bonus! A points-based system will be used to evaluate each group:

- Groups will be given points for satisfying each *specific* constraint.
- They will be penalized heavily for every *must-satisfy* constraint that is not satisfied.
- Additional points will be awarded based on the extent to which *try-to-satisfy* constraints are met, and on the aesthetic appeal of the solution.

This means your group's success depends heavily on whether your individual *must-satisfy* and *try-to-satisfy* constraints are satisfied. Enforce your constraints as best as you can.

When you have finished reading these instructions, please let the investigator know. The investigator will tell you when you may begin working on the problem.

C.3 Questionnaires

Subjects completed questionnaires consisting of Likert-scale questions and open-ended questions before and after the study, and after each of the study conditions.

Pre-Study Questionnaire

Please complete all of the following questions.

1. What is your name?

2. How old are you? (Please circle one)

18-25 26-33 34-41 42-50

3. Are you male or female? (Please circle one)

Male Female

4. Approximately how many hours do you spend using a computer every day? (Please circle one)

< 1.0 1.1-2.0 2.1-3.0 3.1-4.0 4.1-5.0 5.1-6.0 > 6.0

5. What do you use the computer to do (e.g. write email)?

6. Do you have previous experience using haptic (touch-sense) devices? If so, what kind of devices? (e.g. XBox game controller)

Yes No

7. Do you play a musical instrument? If so, what instrument, and for how long?

Yes No

Please write the names of the other three people in your group and indicate how well you know them.

Name	Hardly know each other	An acquaintance	A friend	A good friend	One of my closest friends

1.					
2.					
3.					

Please indicate how much you agree or disagree with the following statements, where 1 = strongly disagree, and 5 = strongly agree.

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5
I am an expert computer user.					
I feel comfortable using a computer.					
With my close friends, I freely express my opinions.					
With acquaintances, I freely express my opinions.					
I prefer working alone.					
When I'm working in a group, I typically take charge.					
Things do not have to be done my way.					

Post-Haptic Condition Questionnaire

Please answer ALL of the following questions. Indicate how much you agree or disagree with the following statements, where 1 = "Strongly Disagree," and 5 = "Strongly Agree."

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5
My group successfully addressed the demands and constraints of the task					
The haptic feedback was too strong					
If multiple people could access Visio at the same time, our solution would have been better					

If we were working face-to-face, our solution would have been better					
I was able to express my opinion					
I obtained control in a reasonable amount of time					
The haptic feedback felt pleasant					
I didn't remain in control long enough					
My group shared control fluidly					
I displayed the User Window when someone asked me for control of Visio					
I displayed the User Window when I was waiting for control of Visio					
I displayed the User Window when I was neither in control nor waiting for control of Visio					
The haptic feedback was too subtle					
My group members listened to my opinion					
I easily recognized what each haptic signal meant					
Sharing control was frustrating					
I easily remembered how to use the two extra buttons on the mouse					
The haptic feedback was distracting					

Please answer ALL of the following questions. Indicate how much you agree or disagree with the following statements, where 1 = "Strongly Disagree," and 5 = "Strongly Agree." If you did not experience the situation described, choose "Does Not Apply"

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5	Does Not Apply
When I moved from <i>waiting for control</i> of Visio to <i>being in control</i> of Visio, I noticed it quickly						
When someone gently requested control from me while I was in control, I noticed it quickly						
When someone urgently requested control from me while I was in						

control, I noticed it quickly						
When I was in control of Visio, and someone took control away from me, I noticed it quickly						

If you have any other comments about this condition, please write them in the space below.

Post-Visual Condition Questionnaire

Please answer ALL of the following questions. Indicate how much you agree or disagree with the following statements, where 1 = "Strongly Disagree," and 5 = "Strongly Agree."

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5
My group successfully addressed the demands and constraints of the task					
If multiple people could access Visio at the same time, our solution would have been better					
If we were working face-to-face, our solution would have been better					
I was able to express my opinion					
I obtained control in a reasonable amount of time					
I didn't remain in control long enough					
My group shared control fluidly					
I constantly monitored the User Window when I was waiting to take control of Visio					
I constantly monitored the User Window when I was in control of Visio					
I constantly monitored the User Window when I was neither in control nor waiting for control of Visio					
My group members listened to my opinion					
Sharing control was frustrating					

Please answer ALL of the following questions. Indicate how much you agree or disagree with the following statements, where 1 = "Strongly Disagree," and 5 = "Strongly Agree." If you did not experience the situation described, choose "Does Not Apply"

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5	Does Not Apply
When I moved from <i>waiting for control</i> of Visio to <i>being in control</i> of Visio, I noticed it quickly						
When someone gently requested control from me while I was in control, I noticed it quickly						
When someone urgently requested control from me while I was in control, I noticed it quickly						
When I was in control of Visio, and someone took control away from me, I noticed it quickly						

If you have any other comments about this condition, please write them in the space below.

Post-Haptic+Visual Condition Questionnaire

Please answer ALL of the following questions. Indicate how much you agree or disagree with the following statements, where 1 = "Strongly Disagree," and 5 = "Strongly Agree."

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5
My group successfully addressed the demands and constraints of the task					
The haptic feedback was too strong					
I relied more on the haptic feedback than the User Window for information					
If multiple people could access Visio at the same time, our solution would have been better					
If we were working face-to-face, our solution would have been better					
I was able to express my opinion					

I obtained control in a reasonable amount of time					
The haptic feedback felt pleasant					
I constantly monitored the User Window when I was waiting to take control of Visio					
I constantly monitored the User Window when I was in control of Visio					
I constantly monitored the User Window when I was neither in control nor waiting for control of Visio					
I didn't remain in control long enough					
My group shared control fluidly					
The haptic feedback was too subtle					
My group members listened to my opinion					
I easily recognized what each haptic signal meant					
Sharing control was frustrating					
I relied more on the User Window than the haptic feedback for information					
I easily remembered how to use the two extra buttons on the mouse					
The haptic feedback was distracting					

Please answer ALL of the following questions. Indicate how much you agree or disagree with the following statements, where 1 = "Strongly Disagree," and 5 = "Strongly Agree." If you did not experience the situation described, choose "Does Not Apply"

	Strongly Disagree 1	Disagree 2	Neutral 3	Agree 4	Strongly Agree 5	Does Not Apply
When I moved from <i>waiting for control</i> of Visio to <i>being in control</i> of Visio, I noticed it quickly						
When someone gently requested control from me while I was in control, I noticed it quickly						
When someone urgently requested control from me while I was in						

control, I noticed it quickly						
When I was in control of Visio, and someone took control away from me, I noticed it quickly						

If you have any other comments about this condition, please write them in the space below:

C.4 Post-Condition Likert Scale Responses

After each condition in Study 3, subjects were asked to fill out a questionnaire. Each questionnaire consisted of a set of 5-point Likert scale questions and some free-response questions. The Likert scale responses are shown here. Sixteen questions appeared in all three conditions, five appeared in the haptic and haptic+visual conditions, and two appeared only in the haptic+visual condition.

Subject ratings are summarized below, organized by group and condition. Each subject within a group is labeled Sx. Group 3 and its subjects (S9 - S12) are not shown because their data were not used; the group struggled to communicate with one another in English during the study and we deemed that it compromised their collaborative efforts. Below the group listings, group averages are listed (Gx), along with the overall average (OV).

My group successfully addressed the demands and constraints of the task

Group	Visual	Haptic	Haptic+Visual
S1	3	4	5
S2	3	2	4
S3	5	4	4
S4	5	4	5
S5	4	2	3
S6	3	3	3
S7	4	4	3
S8	4	4	4
S13	1	4	5
S14	5	5	5
S15	3	4	5
S16	4	4	4
S17	4	5	4
S18	3	5	4
S19	4	5	4
S20	3	5	4
G1	4	3.5	4.5
G2	3.75	3.25	3.25
G4	3.25	4.25	4.75
G5	3.5	5	4
OV	3.625	4	4.125

If multiple people could access Visio at the same time, our solution would have been better

Group	Visual	Haptic	Haptic+Visual
S1	3	2	2
S2	5	5	5
S3	4	4	3
S4	3	4	3
S5	4	4	4
S6	2	4	3
S7	3	3	3

S8	5	5	5
S13	2	2	2
S14	1	1	1
S15	2	2	2
S16	3	2	3
S17	2	2	3
S18	4	2	4
S19	2	2	2
S20	4	2	2
G1	3.75	3.75	3.25
G2	3.5	4	3.75
G4	2	1.75	2
G5	3	2	2.75
OV	3.0625	2.875	2.9375

If we were working face to face, our solution would have been better

Group	Visual	Haptic	Haptic+Visual
S1	3	2	3
S2	5	4	5
S3	4	5	4
S4	3	3	3
S5	5	4	5
S6	4	5	4
S7	3	4	4
S8	5	5	4
S13	5	5	5
S14	3	3	3
S15	4	4	4
S16	4	4	4
S17	2	2	3
S18	2	2	3
S19	4	4	4
S20	4	3	4
G1	3.75	3.5	3.75
G2	4.25	4.5	4.25
G4	4	4	4
G5	3	2.75	3.5
OV	3.75	3.6875	3.875

I was able to express my opinion

Group	Visual	Haptic	Haptic+Visual
S1	4	5	4
S2	2	2	4
S3	5	4	4
S4	4	4	5
S5	4	3	4
S6	4	4	4
S7	4	5	4
S8	4	4	4
S13	4	4	5

S14	4	3	5
S15	4	4	4
S16	5	4	3
S17	4	4	4
S18	4	4	4
S19	4	5	4
S20	5	5	5
G1	3.75	3.75	4.25
G2	4	4	4
G4	4.25	3.75	4.25
G5	4.25	4.5	4.25
OV	4.0625	4	4.1875

I obtained control in a reasonable amount of time

Group	Visual	Haptic	Haptic+Visual
S1	4	4	4
S2	4	4	4
S3	5	2	5
S4	5	4	4
S5	4	5	4
S6	4	3	4
S7	4	4	4
S8	4	4	4
S13	4	4	4
S14	3	3	5
S15	4	4	4
S16	4	4	4
S17	3	4	4
S18	4	4	5
S19	3	5	4
S20	5	5	4
G1	4.5	3.5	4.25
G2	4	4	4
G4	3.75	3.75	4.25
G5	3.75	4.5	4.25
OV	4	3.9375	4.1875

I didn't remain in control long enough

Group	Visual	Haptic	Haptic+Visual
S1	2	4	2
S2	3	4	2
S3	2	4	2
S4	2	2	2
S5	2	2	2
S6	2	4	2
S7	3	3	3
S8	3	3	3
S13	3	3	2
S14	3	4	2

S15	3	2	2
S16	2	2	2
S17	3	2	2
S18	1	1	1
S19	2	1	2
S20	2	2	2
G1	2.25	3.5	2
G2	2.5	3	2.5
G4	2.75	2.75	2
G5	2	1.5	1.75
OV	2.375	2.6875	2.0625

My group shared control fluidly

Group	Visual	Haptic	Haptic+Visual
S1	4	4	4
S2	3	3	4
S3	4	4	4
S4	5	4	5
S5	5	5	4
S6	4	3	4
S7	4	4	4
S8	5	5	5
S13	4	4	5
S14	4	4	4
S15	4	4	4
S16	4	4	4
S17	4	4	4
S18	5	5	5
S19	2	4	4
S20	3	5	4
G1	4	3.75	4.25
G2	4.5	4.25	4.25
G4	4	4	4.25
G5	3.5	4.5	4.25
OV	4	4.125	4.25

I constantly monitored the User Window when someone asked me for control

Group	Visual	Haptic	Haptic+Visual
S1	4	4	5
S2	2	2	2
S3	2	2	2
S4	4	2	4
S5	5	1	5
S6	4	3	2
S7	3	4	4
S8	3	4	2
S13	3	4	4
S14	3	1	5
S15	4	2	3

S16	4	1	4
S17	3	2	4
S18	4	1	2
S19	3	1	1
S20	1	2	3
G1	3	2.5	3.25
G2	3.75	3	3.25
G4	3.5	2	4
G5	2.75	1.5	2.5
OV	3.25	2.25	3.25

I constantly monitored the User Window when I was waiting for control

Group	Visual	Haptic	Haptic+Visual
S1	4	4	5
S2	2	4	2
S3	4	2	2
S4	3	4	4
S5	5	4	3
S6	3	4	2
S7	2	2	2
S8	3	4	2
S13	2	4	4
S14	2	1	4
S15	3	2	2
S16	2	1	4
S17	3	2	2
S18	4	1	2
S19	4	1	1
S20	1	2	3
G1	3.25	3.5	3.25
G2	3.25	3.5	2.25
G4	2.25	2	3.5
G5	3	1.5	2
OV	2.9375	2.625	2.75

I constantly monitored the User Window when I was neither in control nor waiting for control

Group	Visual	Haptic	Haptic+Visual
S1	4	4	5
S2	2	3	2
S3	1	2	2
S4	3	3	4
S5	5	3	4
S6	3	4	2
S7	3	3	3
S8	3	4	2
S13	2	4	4
S14	5	1	4
S15	2	2	3

S16	2	1	2
S17	3	2	2
S18	4	1	2
S19	2	1	1
S20	1	2	3
G1	2.5	3	3.25
G2	3.5	3.5	2.75
G4	2.75	2	3.25
G5	2.5	1.5	2
OV	2.8125	2.5	2.8125

My group members listened to my opinion

Group	Visual	Haptic	Haptic+Visual
S1	4	5	4
S2	4	4	3
S3	4	4	4
S4	5	4	5
S5	4	4	4
S6	5	4	4
S7	4	4	4
S8	4	5	5
S13	4	4	4
S14	5	5	4
S15	4	4	4
S16	5	4	3
S17	4	4	4
S18	4	5	4
S19	4	4	4
S20	4	5	4
G1	4.25	4.25	4
G2	4.25	4.25	4.25
G4	4.5	4.25	3.75
G5	4	4.5	4
OV	4.25	4.3125	4

Sharing control was frustrating

Group	Visual	Post-Haptic	Haptic+Visual
S1	4	3	2
S2	4	4	3
S3	2	5	2
S4	2	2	2
S5	2	3	3
S6	4		4
S7	3	4	3
S8	2	2	2
S13	4	3	3
S14	1	3	1
S15	2	2	2
S16	3	2	2

S17	3	2	2
S18	3	1	
S19	4	2	1
S20	2	2	1
G1	3	3.5	2.25
G2	2.75	3	3
G4	2.5	2.5	2
G5	3	1.75	1.333333
OV	2.8125	2.66666667	2.2

When I moved from waiting for control to being in control, I noticed it quickly

Group	Visual	Haptic	Haptic+Visual
S1	5	2	5
S2	4	4	4
S3	2	4	5
S4	5	5	5
S5	5	5	5
S6	4	4	4
S7	4	4	3
S8	4	4	5
S13	4	4	4
S14	4	4	5
S15	4	4	4
S16	5	5	4
S17	3	4	4
S18	4	4	4
S19	4	4	5
S20	3	4	4
G1	4	3.75	4.75
G2	4.25	4.25	4.25
G4	4.25	4.25	4.25
G5	3.5	4	4.25
OV	4	4.0625	4.375

When someone gently requested control from me, I noticed it quickly

Group	Visual	Haptic	Haptic+Visual
S1	5	5	5
S2	4	4	4
S3	2	5	5
S4	5	4	5
S5	4	2	5
S6	2	4	4
S7	3	4	3
S8	4	5	5
S13	4	5	4
S14	1	4	3
S15	4	4	3
S16	2	2	2
S17	3	2	2

S18	4	3	4
S19	2	4	5
S20	3	4	3
G1	4	4.5	4.75
G2	3.25	3.75	4.25
G4	2.75	3.75	3
G5	3	3.25	3.5
OV	3.25	3.8125	3.875

When someone urgently requested control from me, I noticed it quickly

Group	Visual	Haptic	Haptic+Visual
S1	5	5	5
S2	4	4	4
S3	2	5	5
S4	5	5	5
S5	4		5
S6	2	5	
S7	3	4	3
S8	4	5	5
S13	4	4	4
S14	1	4	4
S15	4	4	4
S16	2	4	4
S17	3	2	2
S18	4	3	4
S19			
S20	3	4	3
G1	4	4.75	4.75
G2	3.25	4.66666667	4.333333
G4	2.75	4	4
G5	3.3333333	3	3
OV	3.3333333	4.14285714	4.071429

(Subjects were instructed to answer this question only if they experienced an urgent request, hence the missing values)

When I was in control and lost control, I noticed it quickly

Group	Visual	Haptic	Haptic+Visual
S1	5	2	5
S2	4	4	4
S3			4
S4	5		5
S5	5	2	5
S6	4	4	
S7	3	4	3
S8	4		
S13	5	4	4
S14	3	3	2
S15	4	4	4
S16		4	4
S17	3	3	4
S18	4	4	4

(Subjects were instructed to answer this question only if they experienced the situation, hence the missing values)

S19			
S20	4		
G1	4.6666667	3	4.5
G2	4	3.33333333	4
G4	4	3.75	3.5
G5	3.6666667	3.5	4
OV	4.0769231	3.45454545	4

The haptic feedback was too strong

Group	Haptic	Haptic+Visual
S1	2	2
S2	2	2
S3	4	1
S4	2	2
S5	2	2
S6	2	2
S7	3	3
S8	2	2
S13	5	3
S14	2	1
S15	2	2
S16	4	2
S17	2	2
S18	1	2
S19	1	2
S20	1	1
G1	2.5	1.75
G2	2.25	2.25
G4	3.25	2
G5	1.25	1.75
OV	2.3125	1.9375

The haptic feedback was too subtle

Group	Haptic	Haptic+Visual
S1	2	3
S2	2	3
S3	4	2
S4	3	2
S5	3	2
S6	2	2
S7	3	3
S8	2	3
S13	2	2
S14	4	3
S15	2	2
S16	2	2
S17	3	3
S18	3	3
S19	2	2

S20	2	4
G1	2.75	2.5
G2	2.5	2.5
G4	2.5	2.25
G5	2.5	3
OV	2.5625	2.5625

The haptic feedback felt pleasant

Group	Haptic	Haptic+Visual
S1	4	4
S2	3	4
S3	2	4
S4	4	5
S5	3	5
S6	3	3
S7	4	4
S8	4	5
S13	3	3
S14	3	3
S15	4	4
S16	4	4
S17	3	3
S18	2	3
S19	4	4
S20	4	4
G1	3.25	4.25
G2	3.5	4.25
G4	3.5	3.5
G5	3.25	3.5
OV	3.375	3.875

The haptic feedback was distracting

Group	Haptic	Haptic+Visual
S1	3	4
S2	4	2
S3	2	2
S4	3	3
S5	3	1
S6	3	2
S7	2	3
S8	4	4
S13	4	2
S14	2	2
S15	2	2
S16	3	2
S17	2	3
S18	3	2
S19	2	2
S20	2	2

G1	3	2.75
G2	3	2.5
G4	2.75	2
G5	2.25	2.25
OV	2.75	2.375

I easily recognized what each haptic signal meant

Group	Haptic	Haptic+Visual
S1	4	4
S2	3	4
S3	4	5
S4	4	4
S5	4	3
S6	5	5
S7	3	4
S8	2	4
S13	5	3
S14	3	4
S15	4	4
S16	1	2
S17	4	3
S18	2	3
S19	4	4
S20	4	4
G1	3.75	4.25
G2	3.5	4
G4	3.25	3.25
G5	3.5	3.5
OV	3.5	3.75

I relied more on the haptic feedback than the User Window for information

Group	Haptic
S1	2
S2	5
S3	4
S4	2
S5	3
S6	5
S7	2
S8	5
S13	2
S14	4
S15	3
S16	1
S17	2
S18	2
S19	5
S20	3
G1	3.25

G2	3.75
G4	2.5
G5	3
OV	3.1

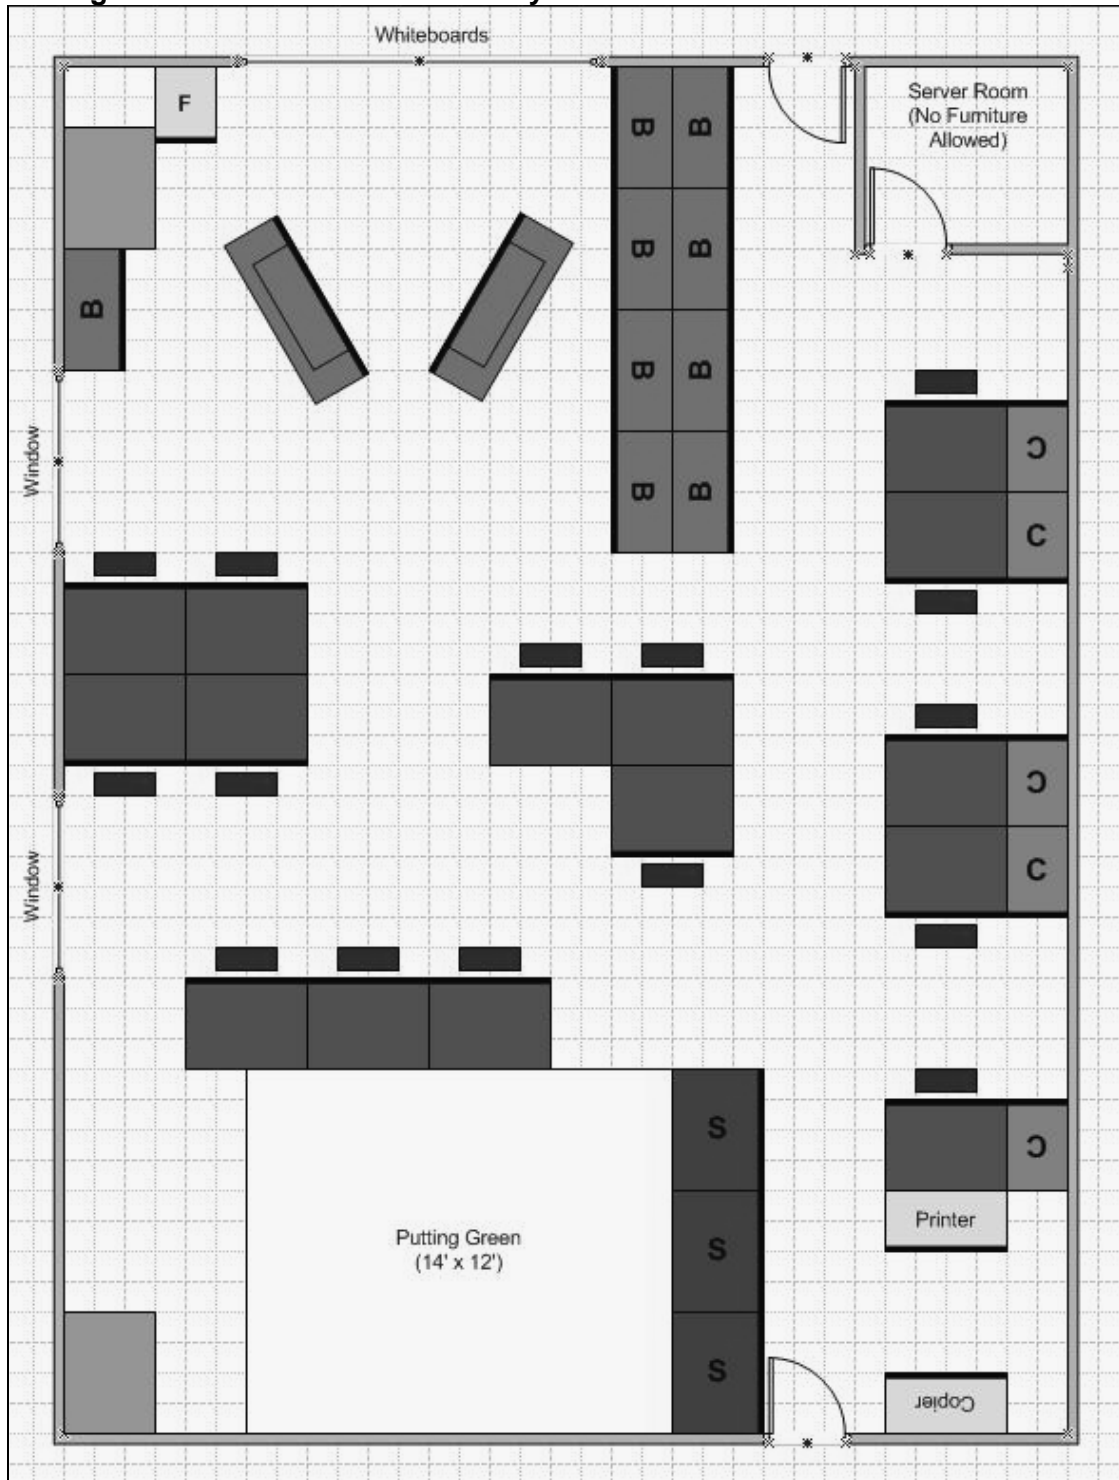
I relied more on the User Window than the haptic feedback for information

Group	Haptic
S1	4
S2	2
S3	4
S4	4
S5	2
S6	1
S7	4
S8	2
S13	5
S14	2
S15	3
S16	5
S17	4
S18	2
S19	1
S20	4
G1	3.5
G2	2.25
G4	3.75
G5	2.75
OV	3.06


C.5 Floor-layout Problems and Solutions

On the following pages the initial floor-layout and reference solution for each of the three tasks used in Study 3 are shown. The scoring keys used to evaluate groups' solutions are also provided.

Adding Workstations Task – Initial Layout



Adding Workstations Task – Reference Solution



Adding Workstations Task - Scoring Scheme

Demands (upper bound: 84 points)

- Adding Desks:
 - -10 points for *each* workstation missing, if there are fewer than the required 20 workstations
 - +10 points for *each* workstation added beyond the required 20 workstations, up to a limit of 50 points
- Pop Pool Layout:
 - +10 points for having fridge, coffee table, bookshelf in close proximity; +5 points if somewhat spread apart; 0 points if items are dispersed
- Image Measurement Area:
 - 0 points if IMA reduced to 6' x 6'; 1 point for every "foot" preserved (e.g. 7' x 6' = 1 points; 7' x 7' or 8' x 6' = +2 points; 14' x 12' = +14 points);
 - -10 points if IMA removed
- Social Area:
 - +10 points for a "viable" social area; +5 points for attempting to have a social area; -10 points if Social Area removed

Must-Satisfy Constraints (lower bound: -80 points)

- -10 points for each violation of the *must-satisfy* constraints


Try-to-Satisfy Constraints (rough upper bound: 50 points)

- -1 point for every pair of workstations that are directly across from one another and touching
- +1 point for every foot of whiteboard that is accessible by a walkway
- +10 / +7 / +4 / 0 points, depending on how well workstations are separated from noisy areas
- +1 point for every foot of window that is accessible by a walkway
- -1 point for every workstation that has does not have two routes to each entrance (within reason - workstation immediately beside door is ok)
- +10 / +7 / +4 / 0 points, depending on how well workstations are clustered
- +10 / +7 / +4 / 0 points, depending on how well bookshelves, workshelves, and filing cabinets are placed
- -3 points for each item of furniture that is removed

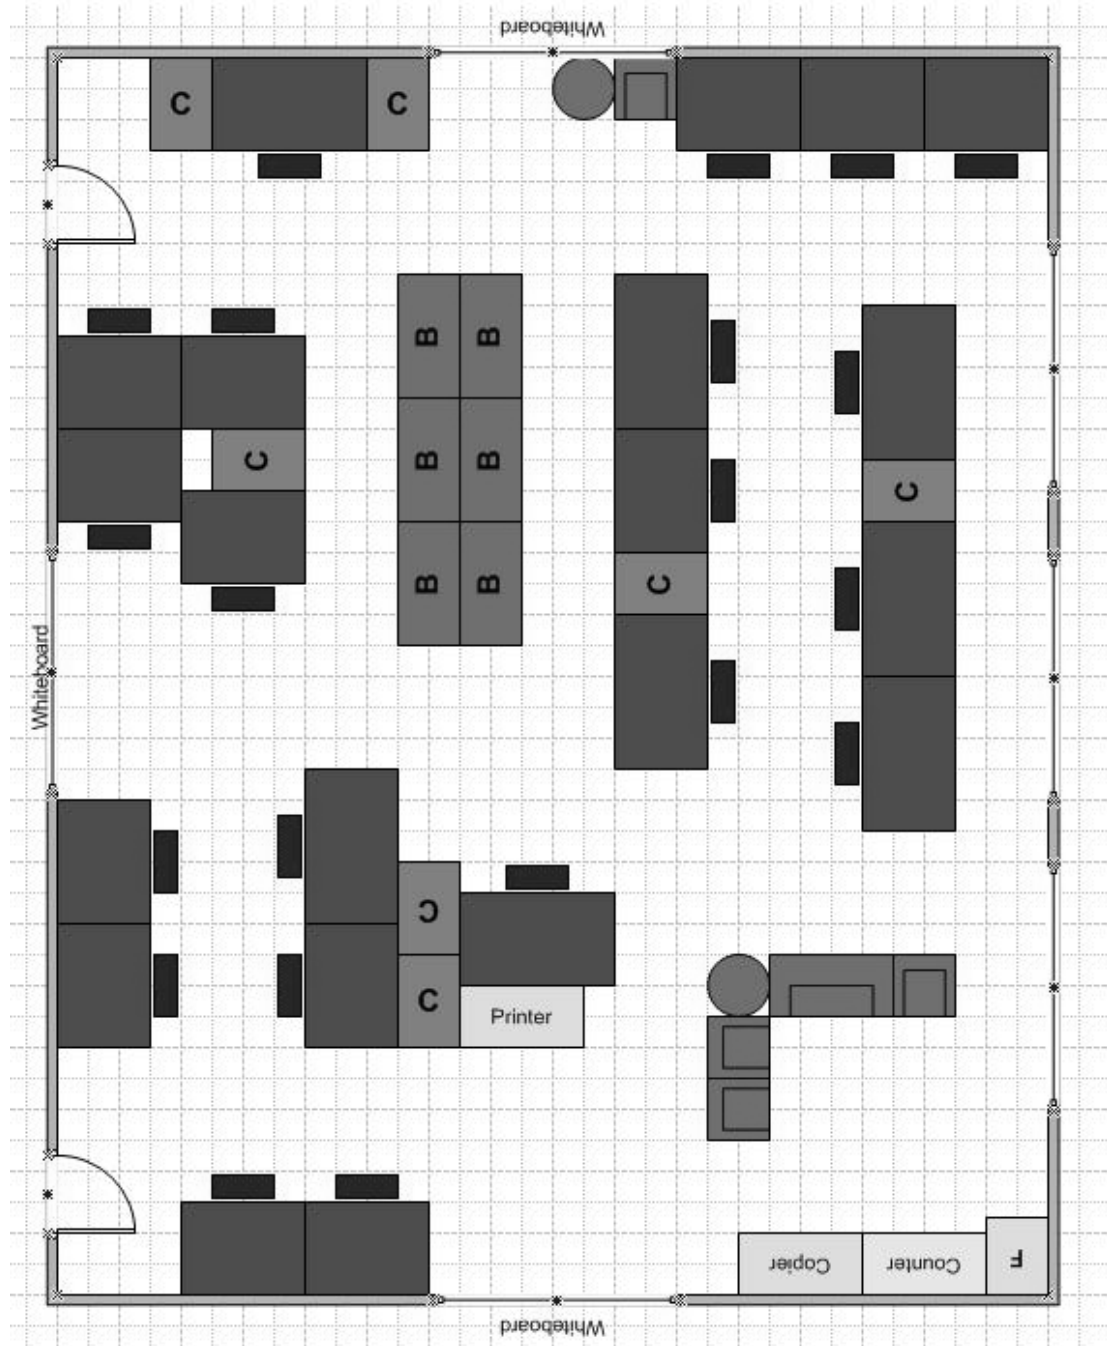
The sample solution for the lab would have a score of $69+0+30 = 99$:

- +40 points for 4 extra workstations
- +5 points for pop-pool layout
- +14 points for keeping IMA the same size
- +10 points for a viable social area
- (satisfied all must-satisfy constraints)
- -9 points for workstations that have a facing, touching neighbour
- +12 points for whiteboard accessibility
- +4 points for fair separation
- +12 points for window accessibility
- +4 for clustering
- +7 points for placement of bookshelves, etc.

Rearranging Workstations Task – Initial Layout



Rearranging Workstations Task – Reference Solution



Rearranging Workstations Task – Scoring Scheme

Demands (rough upper bound: 110 points)

- Keeping Desks:
 - +5 points for *each* workstation kept beyond the required 16 workstations, up to a limit of 25 points
- Social Area:
 - +10 points for a "viable" social area; +5 points for attempting to have a social area
 - -10 points if Social Area removed
- Access to Windows:
 - +5 points for each window that is fully accessible, up to 15 points (duplicates try-to-satisfy constraint)
- "Personal Space"
 - +3 points for each workstation that has no neighbours
 - +1 point for each workstation with only one (non-facing) neighbour

Must-Satisfy Constraints (lower bound: -80 points)

- -10 points for each violation of the *must-satisfy* constraints


Try-to-Satisfy Constraints (rough upper bound: 70 points)

- -1 point for every pair of workstations that are directly across from one another and touching
- +1 point for every foot of whiteboard that is accessible by a walkway
- +10 / +7 / +4 / 0 points, depending on how well workstations are separated from noisy areas
- +1 point for every foot of window that is accessible by a walkway
- -1 point for every workstation that has does not have two routes to each entrance (within reason - workstation immediately beside door is ok)
- +10 / +7 / +4 / 0 points, depending on how well workstations are clustered
- +10 / +7 / +4 / 0 points, depending on how well bookshelves, workshelves, and filing cabinets are placed
- -3 points for each item of furniture that is removed

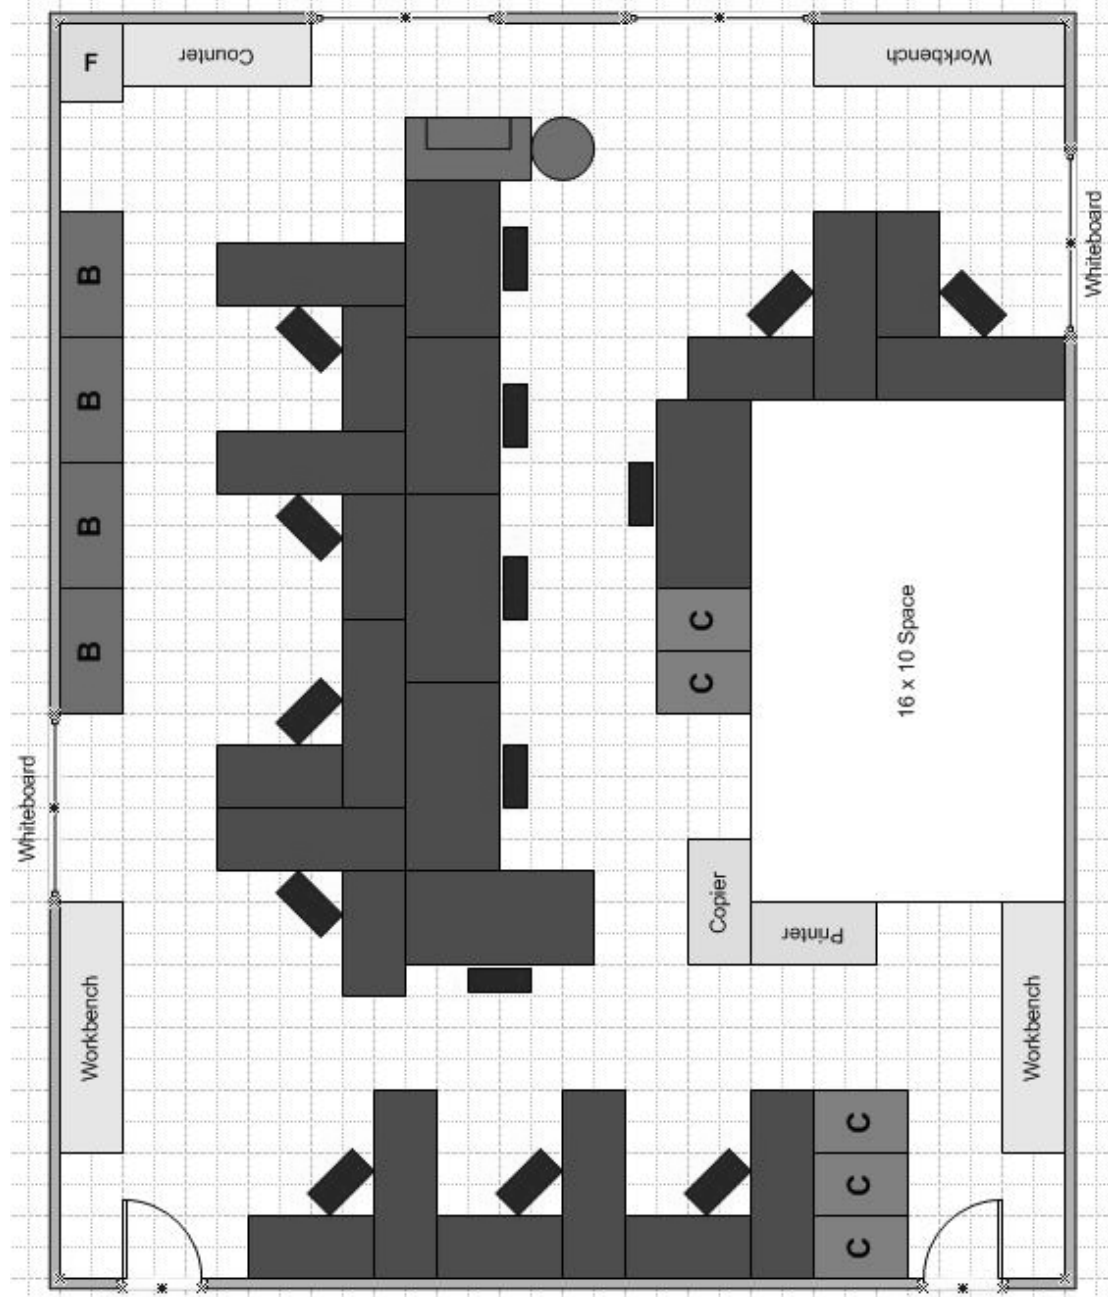
The sample solution would have a score of $74+0+66=140$:

- +25 points for keeping all the extra workstations
- +10 points for a viable social area
- +15 points for access to 3 windows
- +12 points for no-neighbour workstations
- +12 points for one-neighbour workstations
- (satisfied all must-satisfy constraints)
- -1 point for facing/touching workstation pair
- +19 points for accessible whiteboard
- +7 for separation from noisy areas
- +24 points for accessible window
- +10 for clustering
- +7 for filing cabinets

Replaiing Workstations Task – Initial Layout



Replacing Workstations Task – Reference Solution



Replacing Workstations Task – Scoring Scheme

Demands (rough upper bound: 85 points)

- Replacing Desks: (rough upper bound: 75 points)
 - +5 points for *each* 4x4 workstation replaced by a 6x6 (L-shaped) workstation
 - +3 points for each 4x4 workstation replaced by a 6x4 workstation
 - +2 points for each 4x4 workstation replaced by a 5x4 workstation
 - +1 point for each 5x4 workstation replaced by a 6x4 workstation
- Prototyping Area
 - -10 points if it isn't against a wall
 - -5 points for each missing entrance
 - +10 / +7 / +4 / 0 points for isolating the area from the workstations

Must-Satisfy Constraints (lower bound: -80 points)

- -10 points for each violation of the *must-satisfy* constraints

Try-to-Satisfy Constraints (rough upper bound: 60 points)

- -1 point for every pair of workstations that are directly across from one another and touching
- +1 point for every foot of whiteboard that is accessible by a walkway
- +10 / +7 / +4 / 0 points, depending on how well workstations are separated from noisy areas
- +1 point for every foot of window that is accessible by a walkway
- -1 point for every workstation that has does not have two routes to each entrance (within reason - workstation immediately beside door is ok)
- +10 / +7 / +4 / 0 points, depending on how well workstations are clustered
- +10 / +7 / +4 / 0 points, depending on how well bookshelves, workshelves, and filing cabinets are placed
- -3 points for each item of furniture that is removed

The sample solution would have a score of $61+0+19=80$:

- +45 points for replacing 9 - 4x4 workstations with L-shaped ones
- +12 points for replacing 4 - 4x4 workstations with 6x4 workstations
- +4 points for fair isolation
- (no violation of constraints)
- -8 for workstations across
- +12 points for whiteboard access
- +4 points for separation
- +12 points for window access
- +4 points for clustering
- +7 for bookshelf (etc) placement
- -12 for removing 4 items