

Can You Feel the Force?

An Investigation of Haptic Collaboration in Shared Editors

Ian Oakley, Stephen Brewster and Philip Gray

Glasgow Interactive Systems Group, Department of Computing Science

University of Glasgow, Glasgow, G12 8QQ, UK

+44 (0) 141 330 3541

io, stephen, pdg@dcs.gla.ac.uk http://www.dcs.gla.ac.uk/~stephen

Abstract

Users of collaborative systems are typically restricted to communication through voice and video links. Users find this difficult – it does not encompass the richness of communication they are accustomed to in the real world. Attempting to address this problem we describe the implementation of a novel mechanism for haptic communication based around interactions between users' cursors. An initial, and mainly observational, evaluation is described, along with some promising results. We show improvements in subjective experience and suggest several, more formal, avenues for future research.

1. Introduction

With the advent of the World Wide Web (WWW), distributed, collaborative technologies have begun to gather momentum. Real-time collaborative tools are becoming more common, and there is growing acceptance that such tools may be able to reduce travel costs and facilitate group and team based work. However, these systems still face problems. Crucially, the communication that takes place between users does not reflect the richness of communication in the real world. Even assuming high quality video and audio links, information pertaining to sources such as bodily gestures, eye gaze, and background noise is typically lost.

Consequently, users of these systems experience difficulty communicating effectively with one another. Users typically find it difficult to maintain awareness of the actions of others, and to meaningfully coordinate their work activities. Users find it hard to infer where a co-worker's attention is directed, and the use of gestures, both as aids to conversational token passing and as indicators of topic or information, is less effective than interactions in the real world.

There is a large body of research that attempts to address these issues. Much of it focuses on the addition of information through novel means. There is extensive work investigating the impact of video systems that allow participants to maintain eye contact [12], and various researchers have looked at the possibilities of displaying images of the hands

and arms of remote users [8]. Audio information, in the form of auditory icons [4], has also been shown to be an effective addition to these systems. However, despite the improvements suggested by this research, the problems faced by collaborative systems still remain.

Here we consider how haptics could be used to address these problems. Previous literature concerning haptic communication in dedicated, standalone devices has suggested that it provides a strong interpersonal link between users, and raises levels of presence [2, 3]. More recent work investigating the cooperative performance of physical tasks in virtual environments has shown significant performance benefits and an increase in the sense of "togetherness" and presence achieved between participants [1, 15].

However the majority of collaborative tasks are not physical. Here we describe an approach that attempts to take the benefits of haptic communication – not only increases in performance, but perhaps more importantly, a greater sense of engagement and presence – and apply and evaluate their effects in a realistic collaborative task. We focus on trying to augment the communication between users of synchronous shared editors. Shared editors are a common collaborative tool and allow multiple users, often situated at different locations, to simultaneously work on the same document.

2. Implementation

2.1. Haptic Communication

Synchronous shared editors typically include graphical avatars, known as telepointers. These are additional cursors used to represent the position of other users in the workspace. In accordance with the haptic communication described by Oakley *et al.* [13], we base our haptic communication on enhancing interactions between these cursors. Oakley *et al.* define five different mechanisms for haptic cursor communication. Firstly, a push effect, which transforms cursors from a purely graphical representation into tangible objects – they can be felt, knocked into and pushed around. Secondly, a gesture effect which allows cursors to pull each other around. One user can pick up another user, causing that user to be constrained to follow an accurate representation of his or her path. Thirdly and fourthly, locate and grab effects,

which let users activate forces to guide themselves towards another user and to guide other users towards themselves. These effects differ from the gesture effect in that they make no attempt to represent a path, merely a position. Finally Oakley *et al.* describe a proximity effect, which varies the viscosity of workspace according to the proximity of other users, and is designed to provide a basic sense of the presence of others.

2.2. Haptic Simulation

The haptic simulation was implemented in C++ under Windows NT and for the PHANToM (from SensAble Technologies). As we focused on desktop editing tasks the PHANToM was given control of the cursor to provide a standard pointer interface. However, as the PHANToM is an absolute position device, there were some discrepancies with normal mouse operation. The most important of these was that stiff walls lined the edges of the PHANToM's range, preventing users from further motion. The PHANToM's button was mapped to a left mouse click. The workspace itself was presented as a vertical plane.

The gesture, locate and grab effects all require an explicit command to initiate and halt. All commands were mapped to the PHANToM's single button. To gesture, a user moved over another user and depressed the button; to end the gesture the button was released. This control mechanism was designed to support the metaphor of taking hold of, and pulling along, another user. The locate effect was activated by pulling backwards from the workspace and depressing the button. Releasing the button stopped the effect. Finally the grab effect was initiated by pulling backwards from the workspace, performing a double click and keeping the button down after the second click. Once more, releasing the button stopped the effect. All other buttons presses were passed to the system as normal mouse events. The rationale for using the PHANToM's button in this way was to ensure that the haptics could only be activated while the PHANToM was being held.

2.3. Implementation of Networked Haptics

Haptic feedback in general requires high update rates, typically of 500 Hz or more [11]. Most modern networks, for instance ethernet or the Internet, do not provide a quality of service at anywhere near this level. Several researchers are working to address this issue [16]. The goal of this research, however, is to investigate the potential, for a user, of collaborative haptics. Consequently, a high performance, loss-less transmission medium was required. To achieve this we first analyzed the refresh rate requirements for the haptic cursor communication.

In general, high update rates are required to support the production of stable stiff objects [11] and, in the haptic cursor communication, only the push effect involves the production of an object at all – the other cursor. The proximity effect does not require any such representation, merely modifying the viscosity of the workspace based on the distance to other

cursors, and as such should be fairly insensitive to low update rates. Changes in the viscosity from one millisecond to another are relatively unimportant and should fall well below the perceptual threshold. The locate, grab and pull interactions all regard the other user's position as the target of a homing force and as such, when the range to the target is small, so is the magnitude of the targeting force. Lower forces can be adequately represented by lower update rates [11]. Conversely when the distance to the target is large, the targeting force is also relatively large. However in this instance the distance that the target must move to cause a substantial change in the targeting force is also large. These factors all combine to make the haptic communication relatively tolerant of low update rates.

Furthermore, an analysis of the total network demands of haptic cursor communication leads to the conclusion that even assuming a high update rate, a low total bandwidth is sufficient. This is because the only remote information of significance to the simulation is the position and state of other cursors – for instance whether they are engaged in a gesture. This can be expressed in a handful of bytes (16 bytes per PHANToM per update in the current implementation).

Finally, to further simplify development some of the requirements of a collaborative system were relaxed. Firstly we choose to consider a system consisting of only two users and secondly we decided to implement the communication through a dedicated streaming connection rather than across a network. Both of these choices are commonly seen in groupware research [e.g. 4, 8] and typically reflect the intention of the work - to evaluate the potential of the communication for a user, not to investigate underlying network performance.

With these analyses in mind, the haptic communication was implemented running over serial cables and providing a dialogue between two machines at 100 Hz. This implementation suffers from the restrictions of being simple for two users, but more complex for more, and of being confined to machines positioned only a few metres from one another. Due to the update-tolerant nature of the cursor communication, this refresh rate provides a subjective experience that is both stable and of an acceptable quality.

2.4. Collaborative Editor

To provide a flexible platform for evaluating this haptic communication, we turned to GroupKit [14], a high level tcl/tk based groupware toolkit. GroupKit provides support for the management, including the creation, manipulation and mediation of access to, shared information, and basic groupware aids such as telepointers.

The product of coupling the haptic communication with GroupKit is CHASE (Collaborative Haptics And Structured Editing), a synchronous, Relaxed What You See Is What I See (RWYSIWIS) structured drawing tool, pictured in figure 1. It provides telepointers and allows users to simultaneously work on a large canvas while each maintaining a separate view of it. CHASE allows users to create and edit four types

of object: text items, rectangular groups, oval groups, and links. Text items can be placed in group objects and links can be made between them. All items can be freely moved, edited, resized and otherwise manipulated.

One consequence of the combination of the haptic cursor communication and a collaborative RWYSIWIS workspace is that users can feel forces that attempt to move them outside of their current view of the workspace. To resolve these forces intelligibly, the haptic workspace was restricted to only allow cursor movement within the window of the GroupKit application; walls were presented at the edge of the CHASE window. When users pushed into these walls, as they would if forces were pulling them off the workspace, the workspace scrolled in the direction that they are pushing, gradually changing their view until their target was on screen. This solution had the consequence of providing a new mechanism for scrolling – simply pushing into the walls of the workspace.

The haptic simulation communicated with CHASE through mouse events and also through a socket connection. The mouse events formed the main part of the interface, and functioned simply as a consequence of using the PHANToM as a cursor control device. The socket connection was used to communicate all other information, such as the scroll position of the CHASE workspace, and to transmit scrolling events when the PHANToM pushed into a wall. A diagram of this architecture is shown in figure 2.

3. Evaluation

3.1. Design

Evaluation in collaborative systems is a more challenging process than the evaluation of single user systems [6]. Frequently a metric such as time to complete a task, or even quality of output, is less important than the somewhat ephemeral “quality” of the communication between the users. Furthermore, the use of collaborative systems is arguably much more dependant on context than single user systems. There is less guarantee that a feature which appears to work successfully in a laboratory situation will transfer that success to real world use.

This evaluation was also further restricted by the preliminary state of this work. Lacking foundations on which to build, and combined with the challenge presented in the evaluation of any collaborative software, we decided that the most important issue to address was one of acceptability. We choose to address the question of whether or not users would adopt this novel form of communication and, if they did adopt it, what effects it would exert on their subjective experience. Furthermore, although the haptic communication was designed with a purpose in mind, we felt that, if users did adopt it, they were likely to develop unforeseen uses for it. A secondary goal was simply to detect any novel uses of the communication. Finally, as this software represents a prototype, we wanted to discover any problems users experienced with the communication.

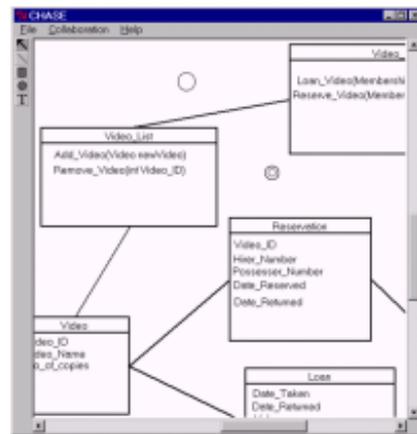


Figure 1. Telepointers and structured objects in CHASE

Designing our initial evaluation to reflect these considerations we chose to perform an observational study of users performing a high level task. This allowed us to observe whether or not participants adopted the communication, what uses they put it to, and the problems they had with it.

3.2. Participants

Sixteen users, all computing science students, in eight pairs, participated in the experiment. Four pairs, forming the Haptic condition, solved a problem with the aid of the collaborative haptic feedback; four pairs, the Visual condition, worked without it. Both sets used the PHANToM in the same restricted workspace, and had access to the same novel scrolling technique. None had previously used a shared editor, nor had anything more than trivial experience with haptic interfaces.

3.3. Task

CHASE was used as a simple CASE (Computer Aided Software Engineering) tool. Participants were required to read a problem statement and design a set of Unified Modeling Language (UML) diagrams to solve it. The advantages of using CASE are that it is an established domain for collaborative tools (e.g. [5]), and that problems, ideal solutions, and semi-expert users are easy to find in an

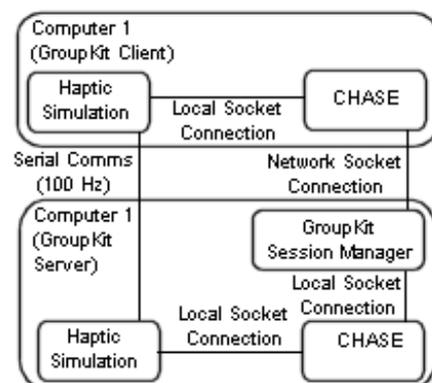


Figure 2. System Architecture for Collaborative Haptics

academic environment. While CHASE cannot represent all the objects in the UML specification, it does allow the development of a basic diagram.

3.4. Measures

Users were observed, with both audio and video recorded, throughout the experiment. No formal dialogue analysis of this data is currently planned. Four questionnaires were administered at the end of the experiment. Firstly, NASA TLX [7], an established questionnaire designed to measure subjective workload. Secondly, QUIS [10], a standard questionnaire for assessing the usability of computer systems. Thirdly, the ITC Presence Questionnaire [9], designed to measure presence in virtual environments. The final questionnaire was created specifically for this experiment to assess collaboration. It consisted of ten items, to be rated on seven-point scales (the questions are included as part of Figure 6, in the results section). The use of this custom questionnaire reflects the lack of an accepted tool for this purpose. This questionnaire data was intended to allow us to gain a measure of the influence of the haptic communication on subjective experience. Objective measures of time to complete task and quality of model were also gathered, but little weight has been attached to them.

3.5. Procedure

For simplicity participants were seated in the same room, separated by a screen. They could not see one another, and no video link was provided. They could talk freely. A disadvantage of this setup is that audio information from the environment, such as the sound of keys presses, which is not typically passed between subjects at different locations, was present. Participants in both conditions went through an extensive training phase. A manual to CHASE, and in the case of the Haptic condition, to the haptic cursor communication was provided and each feature explained and demonstrated. Participants were then required to copy a printed UML diagram into CHASE. At this stage it was clearly explained that the example diagrams consisted of acceptable UML constructs. The recording equipment was then switched on and subjects then began to solve the UML problem. There was no time limit imposed on the task; subjects were instructed to stop when they believed the problem had been solved satisfactorily. This typically took an hour.

4. Results

4.1. General Observational Results

There was substantial variation in the use of the haptic communication. One pair of users in the Haptic condition did not use the communication at all, while others embraced it, using the various effects regularly. This may reflect the fact that touch is a very personal sense. Individual differences and

social factors may well exert a strong influence over the adoption and use of this kind of communication. Furthermore it seems likely that communication of this sort breaks new social ground. Most users prefixed use of the haptic communication with a verbal warning, even when the communication would not affect the other user, as with the locate effect. Users have not previously been able to communicate in this way and appear unsure what protocols should mediate their interaction.

Users also appeared to find the interface to the haptic communication difficult. We attribute this to the fact that the communication was completely controlled through the haptic device; it was simply overloaded. A final crucial point is that the majority of the users appeared to find the haptic communication engaging and helpful, rather than annoying or intrusive.

4.2. Observations of Locate and Grab Effects

The majority of the participants in the haptic condition immediately understood the purpose and applicability of the locate and grab effects. They used the locate effect regularly and the grab effect more infrequently. Reflecting the fact that users found it easier to find one another, pairs in the Haptic condition tended to use much more of the available screen space and created diagrams that spread over a larger area than pairs in the Visual condition. Diagrams in the Visual condition tended to be very compact, with different sections abutting one another.

Participants in the Visual condition also used many different techniques for finding one another or specific objects, which were mainly absent from the Haptic condition. They would describe their position by references to the position of the scrollbars at their location, or by naming nearby objects or diagrams. One participant in the visual condition occasionally instigated a more extreme solution. Upon finding the other user's telepointer he would endeavour to maintain a view on it for as long as possible through rapid scrolling – giving the appearance of pursuing the other user.

One pair in the Visual condition did produce a diagram that occupied a large and diverse area of the screen. They suffered from considerable confusion while discussing where to begin new elements of their solution, and in keeping track of one another when they came to review their diagrams. Such confusion was far less evident in the Haptic condition.

4.3. Observations of Gesture Effect

The haptic gesture was used infrequently. Graphical telepointer gestures, such as when a shape is described by simply moving along its contours, or more commonly, when an object is indicated by moving over it, were far more prevalent. This is probably due to the fact that there is a time cost associated with the haptic gesture. A haptic gesture involves moving over another user, picking that user up, then moving back to the item of interest. The haptic gesture also

provides little enhancement of the most common use of a graphical gesture: indicating a single object.

However, the haptic gesture was used in more complex situations. One pair of participants used the haptic gesture to indicate several objects, spread over an area of the workspace too large to be effectively displayed with a graphical gesture. This could indicate that the haptic gesture is useful for illustrating complex sets of data. Another pair of participants used the gesture in an entirely novel way - one user selected a group object that they were discussing and the other user then began to gesture to the first user, steering them (and consequently the object) towards what they thought was an appropriate location.

4.4. Observations of Proximity and Push Effects

Observing any direct usage of the proximity and push effects was challenging, as neither requires any explicit action to initiate, nor causes an observable motion. However the two effects may have combined to increase a feeling of presence in the workspace. Users appeared to be more confident about using graphical gestures in the Haptic condition, which we suggest may have stemmed from an increased sense of presence brought about by the more tangible representation provided by haptic communication.

4.5. Questionnaire Results

All analyses were conducted using two sample between groups t-tests. Figure 3 contains the results from the TLX questionnaire. Overall workload did not significantly change between the conditions. The Haptic condition, however, was significantly more Physically Demanding than the Visual ($p < 0.05$). The difference in the Frustration Experienced factor approached significance ($p < 0.065$) and we attribute this trend, and possibly also the increase in Physical Demand, to the difficulty users had with the interface to the haptic communication. Most users found the haptic communication hard to invoke, and this is an area that requires substantial improvement before further trials take place.

Figure 4 illustrates the results of presence questionnaire. In keeping with the observation that the haptic communication may increase a user's sense of presence, the Haptic condition

yielded significantly greater subjective ratings of spatial presence ($p < 0.05$). Also, supporting the observation that users found the haptic communication appealing, it achieved significantly higher ratings than the Visual condition in the engagement category ($p < 0.05$). Finally, the Haptic condition was also rated significantly higher on the naturalness factor ($p < 0.05$). This factor attempts to measure how interacting with the system compares with interactions and experiences in the real world.

The results from the usability questionnaire are presented in figure 5. Overall usability was significantly improved in the Haptic condition ($p < 0.05$) as were the individual factors of system usefulness ($p < 0.01$) and interface quality ($p < 0.05$).

Results from the custom questionnaire are shown in figure 6. While this questionnaire was developed simply for this experiment, and as such little trust can be placed in the validity of the data that it produces, it is promising to note the unanimously rated superiority of the Haptic condition.

5. Discussion

The information gained from the questionnaires supports that gained from the observation. On the negative side, users experienced problems with the interface to the haptic communication, and there appears to be substantial variation in the adoption of the communication, probably due to the personal nature of the sense of touch. On the other hand, the majority of the participants appeared to find the haptic communication engaging and used it frequently. It also significantly increased subjective ratings of presence, improved usability and appeared to facilitate collaboration.

The positive results of this experiment suggest future avenues of research. The gesture effect appeared useful when indicating complex information, for instance gesturing to encompass a variety of objects, or a complex shape. We suggest that an evaluation of this technique in such situations would yield interesting results. The locate effect was well received by subjects and it may be interesting to compare this navigation and coordination technique to other, possibly visual, aids. Finally, a direct comparison of the effect of the proximity and push effects on subjective ratings of presence might also prove interesting.

In conclusion, this work has allowed us to assess the

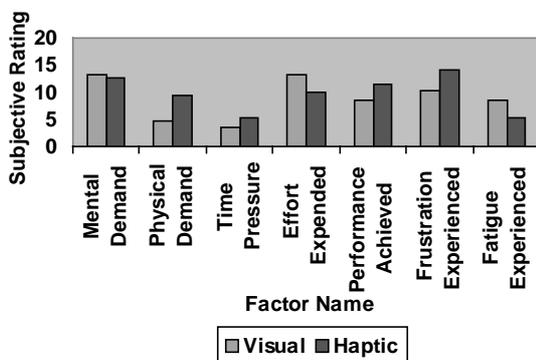


Figure 3. Results from TLX Questionnaire

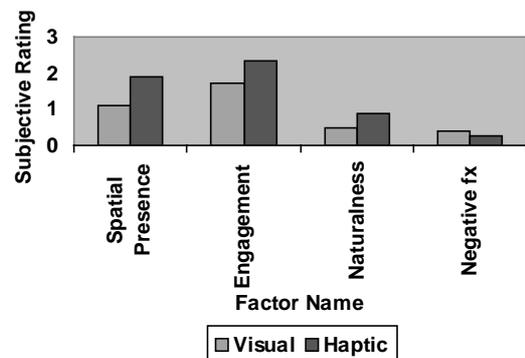


Figure 4. Results from ITC Presence questionnaire

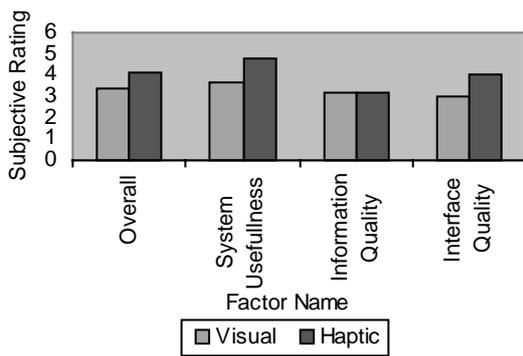


Figure 5. Results from CS Usability Questionnaire

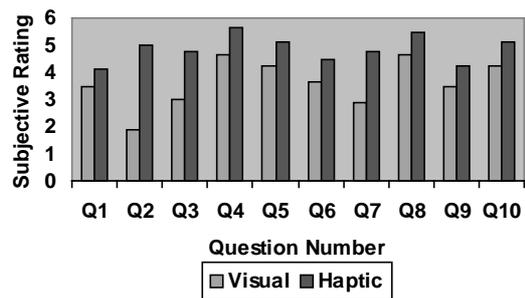
suitability of this novel haptic communication, and to attempt to determine future research possibilities that might stem from it. We have generated several hypotheses for future work, and, more importantly, found that the majority of users regard the communication as appealing and, subjectively at least, gain significant benefits from it.

6. Acknowledgements

This research was supported under EPSRC project GR/L79212 and studentship 98700418. Thanks to the SHEFC REVELATION Project, SensAble Tech and Virtual Presence.

7. References

- [1] Basdogan, C., Ho, C., Srinivasan, M.A., and Slater, M., "An Experimental Study on the Role of Touch in Shared Virtual Environments" *ACM Transactions on CHI*, (2000), 7(4): 443-460.
- [2] Brave, S. and Dahley, A., "inTouch: A Medium for Haptic Interpersonal Communication" in *Proceedings of CHI'97* (1997), ACM Press, 363-364.
- [3] Fogg, B.J., Cutler, L.D., Arnold, P., and Eisbach, C. "HandJive: A Device for Interpersonal Haptic Entertainment" in *Proceedings of CHI'98* (1998), Los Angeles, CA, ACM Press, 57-64.
- [4] Gaver, W., Smith, R.B., and O'Shea, T. "Effective Sounds in Complex Systems: the Arkola Simulation" in *Proceedings of CHI'91* (1991), New Orleans, LA, ACM Press, 85-90.
- [5] Graham, T.C.N., Stewart, H.D., Ryman, A.G., Kopae, A.R., and Rasouli, R. "A World-Wide-Web Architecture for Collaborative Software Design" in *Proceedings of Software Technology and Engineering Practice* (1999), IEEE Press.
- [6] Grudin, J. "Why CSCW applications fail: problems in the design and evaluation of organization of organizational interfaces" in *Proceedings of CSCW'88*, (1988), Portland, OR, ACM Press, 65-84.
- [7] Hart, S. and Wickens, C. (1990) "Workload Assessment and Prediction" in *MANPRINT, an Approach to Systems Integration*, Van Nostrand Reinhold, 257 - 296.
- [8] Ishii, H. and Koyayashi, M. "ClearBoard: A seamless Medium for Shared Drawing and Conversation with Eye Contact" in *Proceedings of CHI'92*, (1992), Monterey, CA USA, ACM Press, 525-532.
- [9] Lessiter, J., Freeman, J., Keogh, E., and Davido, J., "The ITC-Sense of Presence Inventory", (2000), <http://homepages.gold.ac.uk/immediate/immersivetv/P2000-lessiter.htm>



Questions used:

1. There were times when I was unable to communicate effectively with the other user
2. It was easy to find the other user
3. I found the communication in this system to be effective
4. I worked alone
5. The other user and I coordinated our actions together
6. Communicating with the other user was simple
7. I often found it hard to locate the other user
8. I was not aware of the activities of the other user
9. The system did not support my desire to communicate
10. The software made it easy to work as a team

Data has been adjusted so that higher values indicate higher levels of cooperation for all questionnaire items

Figure 6. Questions and Results from Custom Questionnaire

- [10] Lewis, J.R., "IBM Computer Usability Satisfaction Questionnaires: Psychometric Evaluation and Instructions for Use" *Intl Journal of Human-Computer Interaction*, (1995). 7(1): 57-75.
- [11] Minsky, M., Ouh-Young, M., O. Steele, O., Brooks, F.P., and Behensky, M., "Feeling and Seeing: Issues in Force Display" *Computer Graphics, Special Issue on the 1990 symposium of 3D graphics*, (1990). 24(2): 235-243.
- [12] Monk, A.F. and Watts, L. "A Poor Quality Video Links Affects Speech but not Gaze" in *Proceedings of CHI'95*. (1995), ACM Press, 274-275.
- [13] Oakley, I., Brewster, S., and Gray, P. "Communicating with Feeling" in *Proceedings of First International Workshop on Haptics and HCI*.(2000), Glasgow, UK.
- [14] Roseman, M. and Greenberg, S., "Building Real Time Groupware with GroupKit, A Groupware Toolkit" *ACM Transactions on Computer Human Interaction*, (1996). 3(1).
- [15] Sallnas, E., Rassmus-Grohn, K., and Sjostrom, C., "Supporting Presence in Collaborative Environments by Haptic Force Feedback" *ACM Transactions on CHI*, (2000). 7(4): 461-476.
- [16] Wilson, J.P., Kline-Schoder, R.J., Kenton, M.A., and Hogan, N. "Algorithms for Network-Based Force Feedback" in *Proceedings of Phantom User's Group*.(1999), Boston, MA.