

# Collaboration in a Mediated Haptic Environment

J. Jordan, J. Mortensen, M. Oliveira and M. Slater  
Department of Computer Science, University College London, London WC1E 6BT, UK  
[ j.jordan | j.mortensen | m.oliveira | m.slater ] @cs.ucl.ac.uk

B. K. Tay, J. Kim, and M. A. Srinivasan  
The Touch Lab, Department of Mechanical Engineering and Research Laboratory of  
Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA  
[ jungkim | srini ] @mit.edu

---

Keywords: Collaborative Virtual Environments, Internet-2, Presence, Haptics

---

## 1. Abstract

This paper examines an experiment in which pairs of people interact directly via a haptic interface over a network path that has significant physical distance and number of network hops. The aim of the experiment is to evaluate the use of haptics in a collaborative situation mediated by a networked virtual environment. The task of the experimental subjects was to cooperate in lifting a box together under one of four conditions in a between-groups design. Questionnaires were used to report about the ease with which they could perform the task, and the subjective levels of presence and co-presence experienced. This extends the work reported in [1] to consider the possibility of haptic collaboration under real network conditions, and [7] in the addition of the use of haptics.

## 2. Environment

A three dimensional 'room' containing a cube and two pointers is displayed to both participants on their separate desktop systems. The application software is networked such that each participant has a view into the same virtual environment. Each of the pointers is controlled by a haptic device with 6-degrees-of-freedom in terms of motion and 3-degrees-of-freedom force feedback – one being attached to each of the disparate participant's desktop systems. The pointers can be independently moved to any location in the environment (including those outside the room) and may be rotated to point in any direction. Through manipulation of the pointer and consequently the application of force (via the haptic device), the cube may be moved around within the room. The cube is axis aligned and does not respond to torque. A snapshot of the application can be seen in Figure 1.

The walls of the room constrain the cube, gravity is 'on' (at the natural value of  $9.8\text{m/s}^2$ ), and there are unnaturally high dynamic and static coefficients of friction between the cube and the room, and the cube and the pointers.

The actual value of the frictional coefficients is 1.0 (which in practice feels somewhat similar to contact between medium-hard rubber↔rubber surfaces.) The mass of the cube is 0.1kg, which empirically tends to suggest that the cube is light for its size. These values were set such that the mechanical work required to interact with the cube and perform the task would be low.



Figure 1: The application

## 3. Method

Of the two sites involved, the first is situated in the UK (UCL VECG Lab – London), and the second in the USA (MIT Touch Lab – Massachusetts).

The experiment was conducted with twenty subjects ( $n=20$ ) recruited by advertisement at UCL. The collaborator at MIT was a confederate – this being unknown to the subjects at UCL.

The subjects were randomly assigned to one of four groups. Half of the subjects experienced haptic plus visual feedback, and independently half of the subjects saw directional information associated with the pointing device. Thus there was a 2 by 2 factorial design, such that the experiment they undertook either

contained or did not contain pointer directional information, and either provided or did not provide haptic feedback. All tasks used the same virtual environment and procedure, to be described. The same experimental condition was applied for both the subject and confederate ensuring they had the same experience.

## 4. Materials and Implementation

### 4.1 Network conditions:

Between the two sites, the Internet2 network is in place providing a relatively fast connection compared to that of typical Internet speed. There were 13 hops, and there was an approximate round-trip-time of 90ms. However, it is important to note that the inter-arrival times of the data packets would vary, which ultimately affected the performance of the task. The network latency was recorded for each experiment and a graph of one session is shown in Figure 2. Packets were only transmitted when required; that is, when one of the participants was in contact with the cube.

In order to test the system, we initially ran the set-up over a LAN. Networked haptic experiments have been run over LANs in the past, an example of which can be found in [2]. However, for the actual experiments the Internet2 was utilised. A diagram of the network topology is shown in Figure 3.

### 4.2 Software:

The software utilised the PHANToM's GHOST Software Development Kit (SDK) for haptic interaction, and OpenGL for graphical display. It was developed using MS Visual C++.

### 4.3 Hardware:

#### UCL:

- A SensAble PHANToM Desktop haptic device.
- 1GHz PC 512Mb RAM with an NVidia GeForce2 based graphics card running Microsoft Windows 2000 operating system.
- A 19" monitor.

#### MIT:

- A SensAble PHANToM Desktop haptic device.
- Dual 0.9GHz PC 256Mb RAM with an NVidia GeForce2 based graphics card running Microsoft Windows NT operating system.
- A 19" monitor.

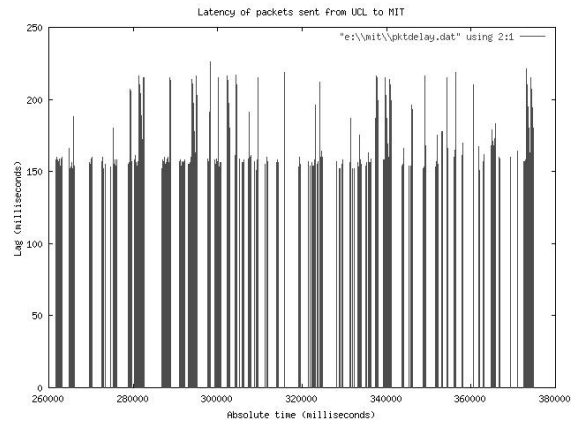


Figure 2: Network Packet Latency



Figure 4: MIT Confederate

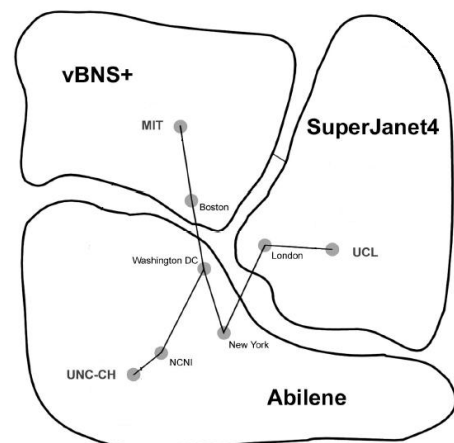


Figure 3: Network Topology

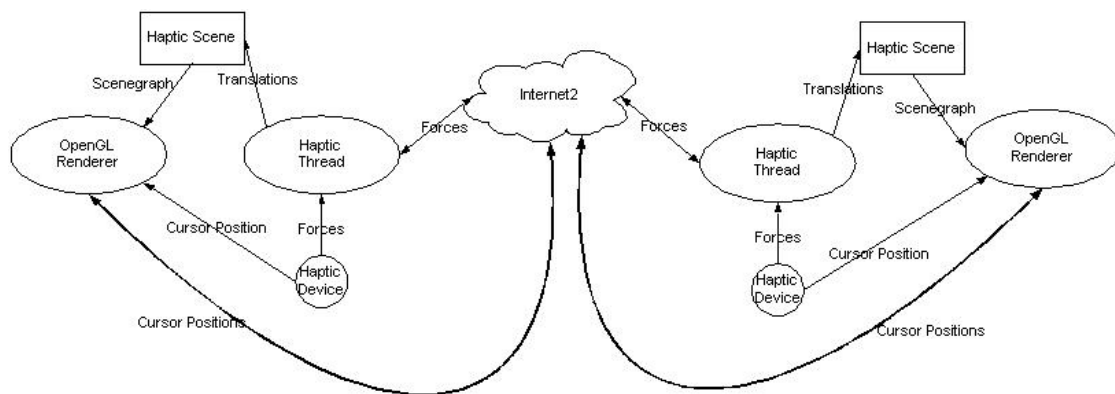


Figure 5: Software Architecture

## 5. Software Design

The software was written in the form of a multithreaded application, which enabled the haptic subsystem to run 'concurrently' (one machine had a single processor, the other having two) with the graphical component. This approach was a requirement, since the initial approach using a particular Distributed Virtual Environment system (DIVE - see [7]) did not yield satisfactory results having its network subsystem coupled to the rendering thread. This coupling prevented the transmission of haptic events over the network from being as frequent as the occurrence of the events themselves. The architecture is shown in Figure 5.

The two machines ran the code independently, and had their own copy of the environment. The system has an inherent peer-peer architecture. This meant that any environmental change made at a machine had to be communicated and then applied at the remote instance. The use of a peer-peer rather than client-server architecture meant that there was no intermediate server introducing further transmission delay between clients. In addition there was no possibility of there being a server bottleneck. Such a (client-server) system has been implemented by [5].

To keep the two systems perfectly in synchrony it would be necessary to apply all changes to the state of the system at exactly the same time at both machines. In reality this was found to be impossible to do with our approach, as it would require all of the equipment to have ideal characteristics (zero latency). The clocks on both systems were synchronised using the Network Time Protocol (NTP) [6].

The network I/O code that related to the haptic subsystem was written into the haptic event loop.

The haptic subsystem was run at exactly 1000Hz due to the requirements of the GHOST SDK. It was decided that for this phase of our work data packets would be sent over the network at 1000Hz providing direct sampling and transmission.

We found that the most effective method to connect the two distant instances of the environment was to send the forces applied to the local cube to the remote system. Conversely, each instance applied any forces received over the network to its cube in addition to any forces applied by the local user.

The same protocol was implemented over both UDP [8] and TCP [9], allowing us to select one of these at run-time. UDP is a basic connectionless transport protocol while TCP provides connection oriented, guaranteed, in-sequence delivery. However, to support a service with these characteristics, TCP has elaborate buffering mechanisms with a complex control protocol with significant overhead that enables it to slowly adapt to network conditions.

Empirically it was found that TCP was inadequate to support user interaction successfully but UDP yielded a workable solution. This did have the side effect that packets of information could be lost. In practice, loss of packets would affect the environment by reducing the force, and therefore desynchronise the location of the cube instances in the shared environment.

## 6. Procedure

All trials of the experiment followed exactly the same procedure:

The application was initialised on both machines simultaneously, and the same 3D environment was displayed to both users (i.e. the confederate, and the subject.)

A simple procedure ensured both the subject and confederate could see each other's pointers, and also the effect of the force(s) the other person could apply to the cube.

The subjects were then allowed a short period to practice (lasting a few minutes) until they were comfortable with the haptic device. This also ensured the subject was competent in manipulating the cube within the virtual environment.

Once this was completed, the cube was centred in the environment, and each participant was instructed to approach the cube from a particular predetermined side in order to lift the cube off the 'ground'. They were to do this by exerting pressure upward and toward one of the vertical sides of the cube only. They were also told to keep the cube off the 'ground' for as long as possible once it was lifted.

The subjects were given approximately 2 minutes to accomplish the task.

## 7. Questionnaire

Upon expiration of the allotted time, regardless of whether the task was accomplished or not, the subject was asked to complete a questionnaire.

The questionnaire recorded data regarding the following criteria:

- Presence (in the environment)
- Co-Presence (with the confederate)
- Task Performance (of self and confederate)
- Demographics

Each was measured on a scale of 1-7, 7 being the highest level for the particular variable being measured. The questionnaire was identical to the one used in [1]<sup>1</sup>.

The main response variable of interest for this paper was co-presence, measured by seven different questions, each on a 1 to 7 scale, where '7' always meant the highest level of co-presence. The response variable was, as in the previous paper [1] taken as the number of high scores out of the seven questions, where 'high' is taken as a score of 6 or more.

The co-presence questions were:

*To what extent, if at all, did you have a sense of being with the other person?*

*To what extent were there times, if at all, during which the computer interface seemed to vanish, and you were directly working with the other person?*

*When you think back about your experience, do you remember this as more like just interacting with a computer or working with another person?*

*To what extent did you forget about the other person, and concentrate only on doing the task as if you were the only one involved?*

*Think about a previous time when you cooperatively worked together with another person in order to move or manipulate some real thing in the world (for example, shifting some boxes, lifting luggage, moving furniture, and so on). To what extent was your experience in working with the other person on this task today like that other real experience, with regard to your sense of doing something together?*

*During the time of the experience, did you often think to yourself that you were just manipulating some screen images with a pen-like device, or did you have a sense of being with another person?*

*Overall rate the degree to which you had a sense that there was another human being interacting with you, rather than just a machine?*

## 8. Results

---

<sup>1</sup> <http://www.cs.ucl.ac.uk/staff/m.slater/Private/Questionnaires/>

The main response variable was analysed using standard normal ANOVA on the mean co-presence scores and also with a logistic-regression based ANOVA as in the earlier experiment (where the count of the number of high scores out of the 7 questions was used). This is a conservative statistical strategy, that avoids the averaging of ordinal scales, and relies instead on count data which can then be used with logistic regression. This method has been successfully used several times before (two examples being, Garau [3], Slater & Steed [10]).

The analysis shows that in the conditions where force feedback was provided the sense of co-presence increased significantly at the 5% level of significance.

Taking co-presence as the response variable, the logistic regression results in a good fitting model to the data (Chi-Squared = 20.43 on 16 d.f.), where the model includes 'haptics' (whether or not haptics was used) 'directional' (whether or not the directional pointer was used) and age of the subject. The use of haptics is associated with an increase in co-presence (Chi-Squared = 5.4 on 1 d.f.). The use of the directional pointer reduces co-presence (Chi-Squared = 11.85 on 1 d.f.). Finally age of the subject is negatively associated with co-presence (Chi-Squared = 5.9 on 1 d.f.). In each case the Chi-Squared value indicates the reduction in fit of the model should the corresponding variable be eliminated. Each one is significant at the 5% level (tabulated Chi-Squared at 1 d.f. is 3.841).

## 9. Conclusion

It is clear that the use of haptics was beneficial to the interaction between two persons involved in a collaborative environment. This was found to be true even in this far from ideal network configuration.

Our intention is now to improve the system such that the application is made to be more efficient in network utilisation, and provides a more stable shared environment. A potential solution might be found in utilising the mechanisms found in [4]. Other improvements will also be made to further enhance the application from an interface/usability point of view.

## Acknowledgements

We would like to gratefully acknowledge the funding of the ESPRC, and the Equator IRC.

We would also like to thank especially both N. Melder and Dr W S Harwin of the Department of Cybernetics at Reading University (UK) for their help, time and the use of their resources.

## References

1. Basdogan, C., Ho, C., Srinivasan, M. A., Slater, M. (2000). An Experimental Study on the Role of Touch in Shared Virtual Environments. ACM Transactions on Computer Human Interaction 7(4), 443-460, 2000.
2. Choi, H. R., Choi, B.H., Ryew, S. M. (1997). Haptic Display in the Virtual Collaborative Workspace Shared by Multiple Users. IEEE International Workshop on Robot and Human Communication 1997.

3. Garau, M., Slater, M., Bee, S., Sasse, M.A. (2001) The impact of eye gaze on communication using humanoid avatars. Proceedings of the SIG-CHI conference on Human factors in computing systems, Seattle, WA USA, 309-316, March 31 - April 5, 2001.
4. Jeffay, K., Hudson, T., Parris, M. (2001). Beyond Audio and Video: Multimedia Networking Support for Distributed, Immersive Virtual Environments. Proceedings of 27<sup>th</sup> EUROMICRO conference, 300-307, 2001.
5. Matsumoto, S., Fukuda, I. et al. (2000). The influences of Network Issues on Haptic Collaboration in Shared Virtual Environments. 5<sup>th</sup> Phantom Users Group Proceedings 2000.
6. Mills, D. (1988). Network Time Protocol. RFC 1059, 1988.
7. Mortensen, J., Vinayagamoorthy, V., Slater, M., Steed, A. (2002). Collaboration in Tele-Immersive Environments. Eighth Eurographics Workshop on Virtual Environments 2002.
8. Postel, J. (1980). User Datagram Protocol. RFC 768, 1980.
9. Postel, J. (1981). Transmission Control Protocol. RFC 793, 1981.
10. Slater, M. and Steed, A. (2000) A Virtual Presence Counter. Presence: Teleoperators and Virtual Environments 9(5), 413-434, 2000.