

Jung Kim
Hyun Kim
Boon K. Tay
Manivannan Muniyandi
Mandayam A. Srinivasan

The Touch Lab
Department of Mechanical
Engineering and The Research
Laboratory of Electronics
Massachusetts Institute of
Technology
Cambridge, MA

Joel Jordan
Jesper Mortensen
Manuel Oliveira
Mel Slater

Department of Computer Science
University College London
London

Transatlantic Touch: A Study of Haptic Collaboration over Long Distance

Abstract

The extent to which the addition of haptic communication between human users in a shared virtual environment (SVE) contributes to the shared experience of the users has not received much attention in the literature. In this paper we describe a demonstration of and an experimental study on haptic interaction between two users over a network of significant physical distance and a number of network hops. A number of techniques to mitigate instability of the haptic interactions induced by network latency are presented. An experiment to evaluate the use of haptics in a collaborative situation mediated by a networked virtual environment is examined. The experimental subjects were to cooperate in lifting a virtual box together under one of four conditions in a between-groups design. Questionnaires were used to report the ease with which they could perform the task and the subjective levels of presence and copresence experienced. This extends earlier work by the authors to consider the possibility of haptic collaboration under real network conditions with a number of improvements. Using the technology described in this paper, transatlantic touch was successfully demonstrated between the Touch Lab at Massachusetts Institute of Technology, USA and Virtual Environments and Computer Graphics (VECG) lab at University College London (UCL), UK in 2002. It was also presented at the Internet II demonstration meeting in 2002 between University of Southern California and the Massachusetts Institute of Technology.

I Introduction

A multimodal virtual environment enables the creation of a multisensory experience by artificial means through interactions with either computer-generated or remote real environments. The effectiveness of virtual environments has often been linked to the richness of sensory information and realness of the experience. Although the development of virtual worlds that can provide a multisensory experience to a single user is receiving substantial attention from researchers, less attention has been given to shared virtual environments (SVEs). In general, single-user VR applications involve the visualization of a scene and interaction with objects within the scene. However, the fundamental aspect of shared experience is the sensory communication between geographically separated users that enables them to display their actions to each other through a connected network.

One of the sensory modalities that needs a more detailed investigation in SVEs is the “haptic channel.” Haptic interfaces in virtual environments have

several applications in various fields, including medicine, education, entertainment, and CAD (Salisbury & Srinivasan, 1997; Srinivasan & Basdogan, 1997). The amount of research in the area of haptic techniques has grown significantly during the last few years, and initial results have already shown the significant role of haptics in the sensory experience of the user. Moreover, in the area of teleoperators, studies have shown that operator performance improves significantly in telemanipulation of remote objects when haptic feedback is provided in addition to visual feedback (Das, Zak, Kim, Bejczk, & Schenker, 1992). Sheridan (1992a) has conducted experiments on automation and planning of complex tasks through human supervisory control and has shown that the task completion times can be reduced with the addition of haptic feedback. It also appears that touching and manipulating objects improves the subjective experience. Durlach and Slater (2000) emphasize that touch, in comparison to other sensory modalities, is more local and bidirectional, which is also linked to closeness and intimacy. Our research in this area also supports these studies and demonstrates the positive impact of haptics on task performance and the subjective sense of togetherness in SVEs (Basdogan, Ho, Srinivasan, & Slater, 2000).

Recent hardware and software advances in haptic interfaces and faster network speed enable us to integrate force feedback into SVEs over a network such as the Internet. Due to inaccessibility, remoteness, hazardousness, or cost-effectiveness, a human operator may not always be present in a work environment. In some other situations, the human users present close to the work environments are not experts (as in the case where tele-robotic surgery is needed) and they rely on the expertise of a remote specialist. Teleoperation has been proved to be a viable alternative for projecting human intelligence over networks. Sheridan (1992b) describes several paradigms for remote control of physical systems by humans through computers, with applications in teleoperation as well as planning and automation of assembly tasks.

Human-human haptic interaction, the main focus of this paper, is different from teleoperation. In a typical teleoperation setup, the “master” end controls the actions of the “slave” robot end, whereas both ends influ-

ence each other in human-human haptic interaction. In teleoperation, an active user interacts with the real world; however, human-human haptic interaction involves interacting mostly with virtual worlds, as in our case. Though human-human interaction can be characterized as bidirectional active collaboration, the other end need not be restricted only to humans; our future studies will explore the haptic interactions between a human and an active computer.

Although there have been several recent studies focused on the development of multimodal virtual environments to study human-machine interactions, less attention has been paid to human-human and human-machine interactions in SVEs. Only recently have researchers paid attention to the question of to what extent the addition of haptic communication between people would contribute to the shared experience (Basdogan, Ho, Srinivasan, & Slater, 2000).

To our knowledge, no attention has been paid at all to research on issues pertaining to human-human haptic interaction over a network of significant physical distance. This research can have a powerful impact on the development of next-generation human-computer interfaces and network protocols that integrate touch and force feedback technology into the Internet environment; on the development of protocols and techniques for collaborative teleoperation such as space station repair or remote surgery; and on the enhancement of virtual environments for performing collaborative tasks in shared virtual worlds on a daily basis, such as cooperative teaching, training, planning and design, cyber-games, and social gatherings.

Our general objective in this research is to provide a necessary foundation for the direct haptic cooperation between humans over a network of significant physical distance and a number of network hops such as networked VR systems and the Internet. Our specific goal in this current study is to investigate the influence of haptic feedback on the task performance of participants over the Internet. Using the technology described in this paper, transatlantic touch was successfully demonstrated between the Touch Lab at MIT, USA and the VECG lab at UCL, UK in 2002. It was also presented

at the Internet II demonstration meeting in 2002 between USC in California and MIT in Massachusetts.

In the next section, we first discuss the development of a multiuser VR setup that enables human-human haptic interaction over the Internet. Details of hardware and software architecture are also described. Design of the experimental paradigm is discussed in section 3. In section 4, the results and their interpretations are presented. The conclusions and a brief discussion of future work are listed in section 5.

2 Methods

2.1 VR Experimental Setup

A virtual three-dimensional “room” containing a cube and two probes 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. Haptic devices with six spatial degrees of freedom (DOF) and three force DOF are attached to each probe. The users directly control the probe through the haptic feedback devices attached to the probes, which can be independently moved and rotated. Through manipulation of the probes with attached haptic devices rendering force feedback, 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, and there are high dynamic and static coefficients of friction between the cube and the room, and the cube and the probes. The coefficient of friction is set at 1.0 to simulate contacts between medium-hard-rubber and rubber surfaces. The effect of gravity is simulated at the natural value of 9.8m/s^2 and the mass of the cube is set at 0.1 kg, which empirically tends to suggest that the cube is light for its size. These values were selected such that the mechanical work required to interact with the cube and perform the task would be low.

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). We conducted the experiment with twenty randomly selected subjects




Figure 1. View of the experimental environment as displayed to the users. Two cursors represent the position and direction of the force feedback devices. The small window in the bottom displays the incoming packet latency.

($n = 20$) 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. The experiment was a 2×2 factorial design in which the subjects were either provided with directional information or not, and haptic feedback or not. Two groups were experimented with haptic and visual feedback, and independently two groups were experimented with directional information associated with the probe. All of the tasks used the same virtual environment and procedure described in the following section. The same experimental conditions were applied for both the subject and the confederate, ensuring they had the same experience.

The two sites were connected with Internet2 network¹, a relatively fast (2.4 Gbps) connection compared to that of a typical Internet connection. There were 13 hops, and there was an approximate round-trip time of 90 ms. However, the interarrival times of the data pack-

¹Internet2 is a research and development consortium of US universities working in partnership with industries and government.




Figure 2. Network packet latency between MIT and UCL.

ets varied and 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 setup over a LAN. Networked haptic experiments have been run over LANs in the past, an example of which can be found in Choi, Choi, and Ryew (1997). However, for the actual experiments, the Internet2 was utilized. A diagram of the network topology is shown in Figure 3.

2.2 Hardware Setup

We used two PHANToM force-feedback devices from SensAble Technologies at both sites. This is a robotic device that lets the user interact with remote and virtual objects. The device has a stylus grip with which the users can touch and feel a 3D object. The update frequency of the device is maintained at 1000 Hz for stable haptic interactions. The hardware specifications of the two sides are given as follows.

MIT:

- a SensAble PHANToM Desktop haptic device
- dual 0.9 GHz PC 256 Mb RAM with an NVidia




Figure 3. Network topology between MIT and UCL.

GeForce2 based graphics card running Microsoft Windows NT operating system

- a 19" monitor

UCL:

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

2.3 Software Architecture

GHOST Software Development Kit (SDK) for haptic interaction from SensAble Technologies and OpenGL for graphical display were used to develop the application. The software was written in the form of a multithreaded application, which enabled the haptic subsystem to run “concurrently” (one machine having a single processor, the other having two) with the graphical component. This approach was a requirement, since the initial approach using a particular Distributed Vir-




Figure 4. *Software architecture.*

tual Environment system (DIVE) did not yield satisfactory results due to having its network subsystem coupled to the rendering thread (Mortensen, Vinayagamoorthy, Slater, & Steed, 2002). 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 4.

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 a client-server architecture meant that there was no intermediate server introducing further transmission delays between clients. In addition, there was no possibility of there being a server bottleneck. Such a client-server system has been implemented by Mortensen et al. (2002).

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 synchronized using the Network Time Protocol (NTP) (Mills, 1988).

The haptic subsystem was run at approximately 1000 Hz 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 1000 Hz, providing

direct sampling and transmission. The network communication code was incorporated in the haptic event loop, enabling the haptic update every time the haptic event loop is called. We found that sending the forces applied to the local cube to the remote system is more effective than sending the position of the probe. Each user felt the forces received over the network from the remote site in addition to forces he or she applied.

The motions of the cube in both sites was computed by the following dynamic equation.

$$m\ddot{\vec{X}} + c\dot{\vec{X}} + k\vec{X} = \sum F$$

$$\vec{X} = [x \ y \ z]^T$$

m , c , k , $\ddot{\vec{X}}$, $\dot{\vec{X}}$, \vec{X} are mass, damping factor, spring constant, acceleration, velocity, and position vector of the cube, respectively. $\sum F$ is the summation of reaction forces from this site and the remote site.

The same protocol was implemented over both UDP (Postel, 1980) and TCP (Postel, 1981), 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 (Niemeyer & Slotine, 2002). Empirically it was found that TCP was inadequate to support user interaction successfully but that UDP yielded a workable solution. However, the side effect of this choice is that packets of information could be lost. In practice, loss of packets would affect the environment by reducing the force, and therefore desynchronize the location of the cube instances in the shared environment.

2.4 Network Latency

Time delay, an inherent property of network communication, poses serious problems in haptic interaction over networks. As early as the 1960s, Sheridan and Ferrell (1963) started investigating this problem. In es-

sence, time delays result in a loss of the sense of causality between the operators' hands.

In the presence of transmission delays, force feedback has a strong destabilizing effect. In this work, the problem is dealt with by adding damping factors at various stages throughout the system, at the user interaction with the cube, in the dynamics of the cube, and also in the probe trajectory. First, a damping factor is introduced at the dynamics of the cube simulating the cube motion.

The dynamic equation of the cube is rewritten as:

$$m\ddot{\vec{X}} + c\dot{\vec{X}} + k\vec{X} = \sum F$$

c , the damping term, simulates viscous friction forces in the environment. It contributes to stabilizing the motion of the cube driven by the force feedback device from both sites. Second, there is a damping term between the virtual cube and the haptic feedback device. When the interaction forces are computed, the velocity of the hand motion is considered as well as the penetration position between the cube and the haptic device.

$$F_{\text{reaction}} = kP_{\text{penetration}} + cV_{\text{phantom}}$$

The variables F_{reaction} , $kP_{\text{penetration}}$, c , and V_{phantom} are reaction forces against interactions of the haptic device, contact stiffness, penetration vector, damping coefficient, and speed of the force feedback device, respectively. Third, a damping factor is introduced at the level of the users' interaction with the haptic device. The user's hand tremor, causing high frequency noise terms, was filtered before passing it to the other end. The hand tremor removal algorithm (as detailed later in this section) effectively acts as a low-pass filter.

Internet time delay is usually modeled as a random process, and stochastic models are used in predicting future values with some error. Most of the prediction techniques use a local model of the other end and use the incoming packets to actually update the local model. Dead reckoning is a technique to reduce network bandwidth employed by distributed military simulations, where a local entity keeps track of the position and orientation of each surrounding entity of interest through the use of a local model. Updated information from the actual surroundings is used to update the local model.

In this work we have used one of the most widely used models of random processes, the Moving Average (MA) algorithm, to predict forces applied by the other end user from the previous known values.

2.4.1 Collision Prediction. The higher haptic-update-rate requirement of the system leaves a very small window for actually performing haptic feedback in each cycle. However, since the users' hands holding the force-feedback devices have low frequency motion (on the order of 10 Hz or less) compared to the sampling frequency of the system (1 kHz), the computational requirement can be simplified by predicting the approach direction of the probe from the trajectories of previous positions (Kim, De, & Srinivasan, 2002). Moreover, the effect of physiological hand tremor could be reduced using low-pass filters to improve the prediction.

Haptic feedback in virtual environments involves collision detection and collision response in each cycle. Collision detection finds the occurrence and location of contact between the probe and the virtual cube. There is substantial literature on this topic in computational geometry, computer graphics, and virtual environments (Gregory, Lin, Gottschalk, & Talor, 1999). Collision response finds the amount of force to be rendered to the user (Ho, Basdogan, & Srinivasan, 1999). Therefore, it is necessary to consider the relationship of both processes in one cycle as well as the individual processing time. The concept of collision prediction is similar to that of motion prediction algorithms in 3D navigation (Chan, Lau, & Si, 2001). The motion of the hand is predicted from the history of tip positions of the force-feedback device driven by a hand. Because the motion of the device is sampled at 1000 Hz, it can predict the motion of the hand, which has lower frequency movements. A more detailed description of the algorithm can be found in Kim et al. (2002), but we briefly explain the algorithm below.

Before the probe first collides with the virtual cube, collision prediction algorithms compute the distance of the probe point to the virtual cube. The computation time is reduced to one distance calculation and one comparison. Without collision prediction, the time

complexity of the collision detection algorithm is directly dependent on object complexity. Also, with collision prediction it is not necessary to perform collision detection at the force update frequency; it could be updated with slower frequency. Collision prediction starts from the determination of a probe-cube vector that is the vector connecting the current position to the previous position.

$$\bar{P}_{\text{tool path}} = \bar{P}_{\text{current}} - \bar{P}_{\text{previous}}$$

Although this method is simple and computationally efficient, the vibration of the tool imposed by the physiological tremor of the user affects the accuracy of the predicted collision point. According to the literature, physiological tremor has a bandwidth of about 7 Hz to 12 Hz (Jones, 1997). To remove this, we take several points from the tool trajectory and average these points as follows.

$$\bar{P}_{\text{mean tool path}} = \frac{1}{m} \sum_{i=1}^m (\bar{P}_{\text{current}} - \bar{P}_i)$$

\bar{P}_i is the i th previous position in the tool trajectory and m is the total number of previous points. This procedure effectively acts as a low-pass filter applied to the tool path vectors and it improved accuracy of the prediction from more than 5 mm deviation to within 1 mm. At the time of the experiment below, however, collision prediction had not been implemented and therefore was not a part of the algorithm used for the experiment.

3 Experiment

The task was to lift a virtual cube in a collaborative environment. The virtual 3D environment was displayed to both the subject and the confederate. A simple procedure ensured that both the subject and the confederate could see each other's probes, and also the effect of the force(s) that the other person could apply to the cube (see Figure 5).

The subjects were then allowed to practice for a short period (a few minutes) until they were comfortable with




Figure 5. Experiment setup at MIT (left) and UCL (right).

the haptic device. This also ensured that the subject was competent in manipulating the cube within the virtual environment. In the actual experiment, the cube was placed at the middle of 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.

3.1 Questionnaire

Each subject was asked to answer a questionnaire at the end of the allotted time, regardless of whether the task was accomplished or not.

The questionnaire recorded data regarding the following criteria:

- presence (in the environment)
- copresence (with the confederate)
- task performance (of self and confederate)
- demographics

Each was measured on a scale of 1 to 7, 7 being the highest level for the particular variable being measured. The questionnaire was identical to the one used in our previous research (Basdogan et al., 2000). The main response variable of interest for this paper was copresence, measured by seven different questions, each on a 1 to 7 scale, where “7” always meant the highest level of

copresence. The response variable was, as in the previous paper (Basdogan Ho, Srinivasan, & Slater, 2000), taken as the number of high scores out of the seven questions, where “high” is taken as a score of 6 or more.

The copresence questions were:

1. To what extent, if at all, did you have a sense of being with the other person?
2. 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?
3. When you think back about your experience, do you remember this as more like just interacting with a computer or working with another person?
4. 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?
5. 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?
6. 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?
7. Overall rate the degree to which you had a sense that there was another human being interacting with you, rather than just a machine?

4 Results

The main response variable was analyzed using standard normal ANOVA on the mean copresence 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 seven questions was

used). This is a conservative statistical strategy, which avoids the averaging of ordinal scales and relies instead on count data that can then be used with logistic regression. This method has been successfully used in the literature (Garau, Slater, Bee, & Sasse, 2001; Slater & Steed, 2000). The analysis shows 5% more sense of copresence with the conditions where force feedback was provided. Taking copresence as the response variable, the logistic regression results in a good fitting of the model to the data ($\chi^2 = 20.43$ on 16 d.f.), where the model includes “haptics” (whether or not haptics was used), “directional” (whether or not the directional probe was used), and age of the subject. The use of haptics is associated with an increase in copresence ($\chi^2 = 5.4$ on 1 d.f.). The use of the directional probe reduces copresence ($\chi^2 = 11.85$ on 1 d.f.). Finally, age of the subject is negatively associated with copresence ($\chi^2 = 5.9$ 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 χ^2 at 1 d.f. is 3.841).

5 Conclusions and Future Work

Direct haptic interaction over a network of significant physical distance and a number of network hops is presented in this paper and has been successfully demonstrated between MIT and UCL as well as between USC and MIT. Various techniques for reducing the transmission delays have also been described. Experimental results show the efficacy of haptics in such a collaborative interaction. This was found to be true even in this far-from-ideal network configuration.

There are a number of future directions for this research. First, it is well known that the presence of transmission delays has a strong destabilizing effect on all feedback systems. In the current experimental system, the stability of force feedback is maintained by adding large amounts of damping at various stages: at the tool interaction with the cube, in the dynamics of the cube, and also in the tool trajectory (hand tremor removal). Although, practically, this strategy has given a stable system, it provides no formal guarantee of stability and

the large damping limits the performance of the interaction. These approaches can only account for a fixed, known delay; they cannot optimize the performance in the presence of random and unpredictable time delays.

An analytical framework that provides a formal guarantee of stability in collaborative interactions will be our future focus. The effect of time delays on the task performance in an SVE may be an important study. Such a study would give insight into the question of to what extent controlled time delays offer any advantages over an ideal no-delay condition. The effects of time delays in other channels, such as visual and auditory channels, on the task performance would also be interesting to investigate further. Second, although the present work experimented with simple virtual environments, fully immersive environments are highly desirable. Visual immersion along with haptic immersion would greatly enhance the collaboration over vast physical distance. In addition to studying the immersive haptic interactions between humans, our future studies will explore the haptic interactions between a human being and a computer. Third, the present work uses point-based haptic exploration, in which the probe (such as a thimble or a stylus) of the force-reflecting haptic devices is modeled as a single point and allows the user to feel the forces that arise from point interaction with virtual objects. More realistic haptic interaction models the probe as a line segment (Basdogan, Ho, & Srinivasan, 1997). This technique, known as ray-based haptic interaction, would enhance haptic realism because it provides a side contact as well as a point contact. Fourth, the subjects in this experiment were unaware of the personality characteristics of the remote partners. It will be interesting to investigate further whether knowing the personality characteristics (such as gender) of the remote partners affects the task performance.

The haptic interaction model used in this study to simulate the underlying physics of interactions (e.g., feeling the forces through pulling or pushing) was simple, but a more sophisticated interaction paradigm can be developed to understand the underlying motor-control, cognitive, and perceptual processes that occur during interactions. One can explore the utility of various models, such as statistical learning models, game

theory, and fuzzy logic techniques, to better understand the role of motor-control, cognitive, and perceptual processes that take place between participants during haptic interactions in SVEs. For example, learning models can be used to explain how well the subjects learn to perform a collaborative task that involves touching and manipulating objects in shared environments. Game theory can be used in modeling haptic negotiations between subjects or between a subject and the computer. Fuzzy logic can be used to investigate the relation between gender and emotional experiences and different network conditions. For example, the interaction model can include fuzzy variables that range from strong to weak, expert to beginner, shy to aggressive, and adverse to favorable network conditions, and the results of the model can be correlated with a selected characteristic such as the gender of the participant.

References

- Basdogan, C., Ho, C., & Srinivasan, M. A. (1997). A ray-based haptic rendering technique for displaying shape and texture of 3D objects in virtual environments. *Proceedings of ASME Winter Annual Meeting*, 77–84.
- 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.
- Chan, A., Lau, R. W. H., & Si, A. (2001). A motion prediction method for mouse-based navigation. *Proceedings of IEEE CGI*, 139–146.
- Choi, H. R., Choi, B. H., & Ryew, S. M. (1997). Haptic display in the virtual collaborative workspace shared by multiple users. *Proceedings of IEEE International Workshop on Robot and Human Communication*, 478–483.
- Das, H., Zak, H., Kim, W. S., Bejczk, A. K., & Schenker, P. S. (1992). Operator performance with alternative manual control modes in teleoperation. *Presence: Teleoperators and Virtual Environments*, 1(2), 201–218.
- Durlach, N., & Slater, M. (2000). Presence in virtual environments and virtual togetherness. *Presence: Teleoperators and Virtual Environments*, 9(1), 214–217.
- Garau, M., Slater, M., Bee, S., & Sasse, M. A. (2001). The impact of eye gaze on communication using humanoid avatars.

- tars. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 309–316.
- Gregory, A., Lin, M., Gottschalk, S., & Talor, R. (1999). H-COLLIDE: A framework for fast and accurate collision detection for haptic interaction. *Proceedings of IEEE Virtual Reality Conference*, 38–45.
- Ho, C., Basdogan, C., & Srinivasan, M. A. (1999). Efficient point-based rendering techniques for haptic display of virtual objects. *Presence: Teleoperators and Virtual Environments*, 8(5), 477–491.
- Jones, L. (1997). Dextrous hands: Human, prosthetic and robotic: A survey. *Presence: Teleoperators and Virtual Environments*, 6(1), 29–56.
- Kim, J., De, S., & Srinivasan, M. A. (2002). Computationally efficient techniques for real time surgical simulation with force feedback. *Proceedings of the International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 51–57.
- Mills, D. (1988). *Network Time Protocol*. RFC 1059.
- Mortensen, J., Vinayagamoorthy, V., Slater, M., & Steed, A. (2002). Collaboration in tele-immersive environments. *Proceedings of the Eighth Eurographics Workshop on Virtual Environments*.
- Niemeyer, G., & Slotine, J. E. (2002). Toward bilateral Internet teleoperation. In K. Goldberg & R. Siegwart (Eds.), *Beyond Webcams: An Introduction to Online Robots* (pp. 193–213). Cambridge, MA: MIT Press.
- Postel, J. (1980). User datagram protocol. RFC 768.
- Postel, J. (1981). Transmission control protocol. RFC 793.
- Salisbury, J. K., & Srinivasan, M. A. (1997). PHANToM-based haptic interaction with virtual objects. *IEEE Computer Graphics and Applications*, 17(5), 6–10.
- Sheridan, T. B. (1992a). Musings on telepresence and virtual presence. *Presence: Teleoperators and Virtual Environments*, 1(1), 120–126.
- Sheridan, T. B. (1992b). *Telerobotics, automation, and human supervisory control*. Cambridge, MA: MIT Press.
- Sheridan, T. B., & Ferrell, W. R. (1963). Remote manipulative control with transmission delay. *IEEE Transactions on Robotics and Automation*, 9(5), 592–606.
- Slater, M., & Steed, A. (2000). A virtual presence counter. *Presence: Teleoperators and Virtual Environments*, 9(5), 413–434.
- Srinivasan, M. A., & Basdogan, C. (1997). Haptics in virtual environments: Taxonomy, research status, and challenges. *Computers and Graphics*, 21(4), 393–405.