

Effects of network delay on a collaborative motor task with telehaptic and televisual feedback

Robert S. Allison*
Dept. of Computer Science
York University

James E. Zacher†
Centre for Vision Research
York University

David Wang, Joseph Shu‡
Handshake Interactive
Technologies

Abstract

The incorporation of haptic interfaces into collaborative virtual environments is challenging when the users are geographically distributed. Reduction of latency is essential for maintaining realism, causality and the sense of co-presence in collaborative virtual environments during closely-coupled haptic tasks. In this study we consider the effects of varying amounts of simulated constant delay on the performance of a simple collaborative haptic task. The task was performed with haptic feedback alone or with visual feedback alone. Subjects were required to make a coordinated movement of their haptic displays as rapidly as possible, while maintaining a target simulated spring force between their end effector and that of their collaborator. Increasing simulated delay resulted in a decrease in performance, either in deviation from target spring force and in increased time to complete the task. At large latencies, there was evidence of dissociation between the states of the system that was observed by each of the collaborating users. This confirms earlier anecdotal evidence that users can be essentially seeing qualitatively different simulations with typical long distance network delays.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces – haptic I/O; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism – virtual reality.

Keywords: teleoperation, haptics, collaborative virtual environments, virtual environments, delay

1 Introduction

Collaborative virtual environments (CVE) have been implemented for a variety of applications including distributed training, multiuser tele-operation, gaming, tele-medicine, computer-mediated social interaction, entertainment and computer supported collaborative work. In these environments, users share a common virtual space though they may be geographically separated. Ideally, the ability to share a distributed virtual environment facilitates enhanced communication and collaboration between users. In a CVE, active exploratory touch can provide an intimate

and interactive means for communication between users. Touch is thought to be particularly potent contributor to co-presence due to the closeness and intimacy that is characteristic of interpersonal touch [Durlach and Slater 2000]. Collaboration in a virtual environment often incurs significant ‘cooperation overhead’ due to the requirement to coordinate actions without the typical affordances of the real world. It has been proposed that coordination could be facilitated by appropriate haptic collaboration minimizing this overhead [Ruddle et al. 2002].

However, implementers of CVEs, and especially haptically-mediated CVEs, face many of the problems of single user virtual environments. In addition, there are significant challenges to provide interpersonal interaction and to maintain a consistent and synchronized virtual environment for all users.



Figure 1. Photograph of the robot used for presenting haptic display to the subject.

2 Delay in haptically-mediated CVEs

In virtual environments and tele-operators, latency is universally detrimental resulting in mismatch between motor action and simulated sensory feedback. For example, in manual control tasks in a virtual environment, delay can interfere with coordination and planning of motor actions [Ellis et al. 1997; Park and Kenyon 1999].

In shared CVEs, latency entails further complications as users act in the simulated environment and the results of these actions must be observable to other users (and themselves). System response latency can introduce discord between users’

*e-mail: allison@cs.yorku.ca

†e-mail: zacher@cvr.yorku.ca

‡e-mail: dwang,jshu@handshakeinteractive.com

experiences in the shared VE [Vaghi et al. 1999]. In a distributed CVE, local system latency may be low for each user but transmitting VE updates between users over long distance network links entails transport, queuing and processing delays. Delay between users can result in disconnect between the simulations experienced by the users that is manifest as a number of qualitative effects: disruption of cause and effect relations [Pantel and Wolf 2002; Vaghi et al. 1999], loss of sense of co-presence [Mortensen et al. 2002], difficulty in coordinating actions [Park and Kenyon 1999], and discontinuities in VE state [Pantel and Wolf 2002; Sharkey et al. 1998].

Arguably, the sensory modality most affected by sensorimotor discord is touch [Durlach and Slater 2000]. One reason is that we can only touch things that are close to us making touch more local and intimate than vision or audition. Also the hand (or other effector) is both a sensor and an end-effector. Conversely, the haptic device is typically both a display and an input device. In natural environments, this tight input-output coupling normally implies a lawful synchrony between the user's action and the sensory feedback.

2.1 Dependence on type of interaction

Interpersonal interaction in shared virtual environments can be classified in a number of ways [Margery et al. 1999]. For example, one can ask whether the users collaborate to achieve a goal, act independently, act while avoiding interference with others, or compete with each other in the VE. Synchronization of the virtual environment is vital to ensuring that the virtual environment is consistent and causal relations are maintained.

Users can also collaborate (or compete) at a variety of levels in terms of interpersonal coupling. Many applications require a mix of closely and loosely coupled interactions. At a high level they can plan behaviours by communicating information verbally, haptically or visually or by performing subtasks of a composite task in a coordinated but somewhat decoupled manner. In these collaborative situations consistency of the CVE is important but some modest asynchrony between the states of the VE observable by the users can be tolerated [Mortensen et al. 2002].

On the other hand, close and dynamic interaction between the users, for example simultaneous cooperative manipulation of a shared object, requires a tight synchronization of the mutual experiences of the VEs to prevent discontinuities, distortion or loss of causality during interaction. A simple form of this type of interaction is when one user tracks the movements of another (the leader) who acts independently to accomplish some goal. More complicated scenarios require users to bilaterally communicate or monitor each others actions to coordinate their behaviour. Coordination of this type of interaction can be accomplished haptically as in ballroom dancing, handshaking, passing of objects between individuals, etc. In this paper we study closely-coupled, symmetric cooperative manipulation.

2.2 Studies of the effect of latency on telehaptic collaboration

Haptically mediated telecollaboration has been the subject of relatively few studies compared to visually based CVE research. McLaughlin et al [2003] studied information

transmission and co-presence in a shared haptically-mediated VE. They found that the transmission of messages haptically (using coded touch signals) was possible in a CVE although the relationship between co-presence and information transmission effectiveness depended in a complicated manner on a number of factors. Thus an effective CVE in terms of information transmission may not be effective in terms of co-presence measures (and vice-versa).

Basdogan [2000] studied the effect of haptic communication on the experience of a shared CVE. Addition of haptic feedback on a collaborative visuomotor task that required precise inter-user coordination – Ellis's ring tracing task [Ellis et al. 1997] – improved performance and the sense of togetherness and co-presence. In a later study [Jordan et al. 2002], haptic feedback improved sense of co-presence in tele-collaboration even in the presence of realistic, 90ms roundtrip, transport delay.

Chong et al [2002] discussed tele-operated multi-operator, multi-robot systems. They reported simulation and field experiments to coordinate the planning of a placement tasks. Predictive simulation of the remote robots, which were controlled through a central relay station, allowed for improved coordination and reduced chance of collisions. Note that the authors considered coordinated planning of composite tasks rather than the coordinated, close-coupled manipulation of a shared object.

3 Rationale and outline of the experiments

Despite the importance of close cooperative work in many human endeavours there has been little study of cooperative manipulation of a common object in virtual environments, especially in haptically-mediated CVEs. Furthermore, most studies to date have considered what haptic information adds to a visual display. There has been little consideration of whether haptic feedback can facilitate cooperative work and communication in a CVE in its own right. In this paper we consider the effects of simulated network delay on the ability of two remote users to collaborate based on visual and haptic communication. We used a simple symmetric collaborative task that could be performed using either haptic or visual information alone and studied the effects of varying simulated network latency and of predictive compensation for delay.

4 Methods

4.1 Apparatus

Two custom-built haptic workstations (Handshake Interactive Technology, Kitchener, Canada) were configured for tele-collaboration. Each workstation consisted of a robotic manipulator and associated drive electronics, a monitor for graphics display and a Pentium class computer (see Figure 1). Each manipulator was a single direct-drive 5-bar-linkage haptic device. This device has a full 3 degrees of freedom and each of the links measure 18 inches. The planar joints were counterbalanced through a spring mechanism to reduce gravitational effects

The robot was locked into a configuration where it could pivot in yaw about its base but other motions were restricted. Haptic effects are achieved by exerting torques on the joints of the robot which, in turn, creates forces on the user's hands. Force and graphics rendering C code was custom designed and

written to run under Windows 2000 augmented by additional real-time operating-system support. The robots were similar but not identical so, for consistency, the weaker robot was always assigned to the expert user (see below).

The workstations were connected via a 10BaseT Ethernet LAN connection and synchronized under software control. Baseline delay between the two devices was 10ms. The graphic and haptic displays were updated at 100Hz. Additional constant transport delay was simulated using a fixed-size buffer to store the time-stamped incoming packets. The participants were physically located in the same room but a sheet of black cloth hung between the two stations preventing the subjects from watching each other and verbal communication about the progress of each trial was not allowed.

In the conditions where prediction of network delay was incorporated, Handshake's Hard Real Time Control Centre (HRTCC) was used to compensate for the network delay (see Appendix). This software resides on both nodes in the telehaptic system. As input data is generated and transmitted from the collaborating node a (simulated) network delay is incurred. At the receiving end this delay is measured. The goal is to generate data which approximates the data that would have arrived in the absence of any delay. Thus, a prediction of the undelayed data is generated based on the current network delay and the data time series. The latency used by the compensator was near but slightly smaller than the measured delay to allow for better matching between predicted and actual inputs. Some tuning was required to control overshoot and instability.

4.2 Subjects

Two individuals performed the experiment at a time. One individual was one of two of the experimenters who served as a confederate or an 'expert' user. The other individual was the experimental subject who was naïve with respect to the purposes of the experiment.

Twelve individuals participated in the experiment as experimental subjects (8 men and 4 women, mean age 30.6 years). All subjects had normal visual acuity with their habitual optical correction (7 corrected, 5 uncorrected) and no had history of sensorimotor or visual problems. Three subjects reported that they had experience with online gaming, distributed virtual environments or teleoperation environments, three had reported some experience, while the other six had reported little or no such experience. Subjects gave their informed consent to participate in the study.

4.3 Procedure

Each trial proceeded as follows: In the simulation, the two workstations were connected through a virtual linear spring. Subjects attempted to move the manipulator between two predefined end points (separated by 90 degrees) as rapidly as possible, subject to the constraint that the simulated spring force between the two devices be maintained within a target range. A trial could be started when both the expert and subject had their arm located behind the appropriate start point. The subjects then commenced the movement and timing and data collection began when the arm moved past the start point. The simulation software simulated a flexible link between the two computers that approximated an elastic rope (i.e. it could pull but not push) with a linear spring force characteristic. On the right to left movement the expert user led the motion and pulled the subject

who was required to follow the leader while maintaining spring force. In practice, both leader and follower needed to adjust their movements to complete the task. On the return pass (left to right) the subject led experimenter.

During the trial, the relative position of the two manipulators was monitored and used to calculate the simulated spring force. If the force was within the desired range then an indicator bar on the display was coloured green; if it was not in range, then the bar was coloured red. This only provided information about whether force was within the desired range. Trials differed in the type of feedback provided about the *direction and magnitude* of the error: haptic versus visual. The error in simulated force was communicated to the user either i) haptically by rendering the simulated force or ii) visually by displaying two triangles on the screen, one for each robot, with the relative separation of the triangles indicating the relative position of the two robots and hence, indirectly, the force.

The simulated transport delay between the two workstations was varied on a trial by trial basis between 0, 100 and 200ms for both haptic and visual feedback conditions (3 delays by 2 feedback conditions). An additional three conditions were formed by running the three visual feedback conditions with network prediction enabled. Thus, there were a total of nine conditions, which were repeated five times in each of two sessions for each subject (for a total of 90 trials for each subject). Trials were blocked in groups of three haptic or six visual trials, randomized within blocks and counterbalanced over the blocks and sessions.

5 Results

A repeated-measures Multivariate Analysis of Variance (MANOVA) was conducted on a factorial model that included the following independent variables: Delay (0, 100 & 200 ms), Feedback Mode (Visual & Haptic), Prediction (Yes & No), Role (Lead & Follow) and Person (Expert & Subject). Gaming/teleoperation experience was treated as a between subjects effect (Experienced, Some & None). Time to complete the trial, number of overforce errors, number of underforce errors and time spent executing an overforce or underforce error were included as dependent measures. For the purpose of clarity only significant results will be discussed.

Increasing system delay increased the overall time to complete the task. This was true for both haptic and visual feedback for leaders and followers (Figures 2 and 3). Increased delay also significantly increased the number of overforce and underforce errors as well as the time spent in the overforce condition. Thus increasing time delay made the task more difficult whether measured as time to complete the task or degraded ability to stay within the target force range.

The most influential factor in the analysis was whether an individual had to lead or follow during the task. The effects of delay on overforce and underforce errors and on time spent in overforce and underforce conditions for both leaders and followers are plotted in Figures 4 and 5. According to their own local simulation, leaders tended to make predominantly overforce errors that increased in duration and number with increasing time delay. Conversely, followers tended to make predominantly underforce errors that increased in duration and number with increasing time delay. This was true both in the overall analysis and in individual trials. In individual trials the leader's and follower's simulations could be dissociated by




Figure 2. Effect of delay on time to completion for visual and haptic feedback conditions

delay such that the leader experienced a large number of overforce errors while the follower perceived underforce errors.

The frequency of underforce errors and time executing underforce errors increased with delay for followers while it decreased for leaders (who exhibited few underforce errors in any case). The opposite is true for number of overforce and the amount of time in the overforce condition which increased for leaders and decreased for followers as delay increased.

Subjects took less time to complete a trial with visual than haptic feedback, especially at zero delay. However, there was an interaction with delay so that no difference remained between haptic and visual feedback trials at 200 ms (Figure 2). The Multivariate analysis indicated a Delay x Mode interaction with a low η^2 term indicating this was a rather weak effect. Subjects had fewer underforce errors with visual feedback than with haptic feedback. In turn, with visual feedback, subjects had more overforce errors and spent more time executing overforce errors than with haptic feedback.

During the visual feedback trials, the prediction algorithm reduced the increase in task completion time with delay.

Finally, the results indicated that there were no significant




Figure 3. Effect of delay on task completion time for individuals who led or followed.




Figure 4. Effect of leading or following on the mean number of overforces and underforces.

differences between the performance of the confederate and experimental subjects. However, subjects with some or extensive online gaming/teleoperation experience showed shorter error durations and less under or overforce errors than subjects with no prior experience.

6 Discussion

In this study we demonstrated that delays, typical of those encountered in long distance telecollaboration, inhibit the ability of users to collaborate using haptic or visual communication at a distance. Furthermore, we determined that delay can result in dissociation between the state of system at the two end stations and this can result in marked differences in the response of operators compensating for perceived errors. For example, in our study there were significant differences in error rates and qualitative differences in the type of error depending on whether the subject was leading or following. These context dependent differences in the effect of latency are a distinctive feature of distributed shared virtual environments with delay between users. We used a simple but realistic collaborative task that requires users to coordinate forces or positions of input devices to achieve a common goal. In this discussion we'll consider the implications of our results for visuo-haptic collaborative virtual environments and for telehaptic operation of remote machines.

6.1 Haptic communication

Our task could be classified as a symmetric manipulation although one subject took the lead in guiding the movement [Ruddle et al. 2002]. In these experiments the changing of the leadership role was experimentally controlled. Despite the presence of a nominal leader the haptic communication was bilateral – the leader could not simply move as fast as possible but needed to monitor the following ability of the follower and regulate his/her speed accordingly. This type of situation is quite common in symmetric tasks such as the furniture moving or stretcher carrying scenarios that have been studied in the CVE literature [Hubbold and Keates 2000; Mortensen et al. 2002]. During these activities in real environments, leadership roles are determined naturally and dynamically, often with little verbal communication. In asymmetric manipulation the temporal relationships between the cooperating users actions can be even more subtle and complex and prediction of




Figure 5. Effect of leading or following on the mean time for overforces and underforces.

another's movements cannot be as easily predicted from one's own motion. It is known that latency can interfere with turn taking in conversation [Ruhleder and Jordan 2001]. Given that haptic cooperation is more tightly synchronized, simultaneous and dependent on subtle force feedback cues, we expect delay to be especially disruptive to turn taking and coordination during haptic tasks.

Often in CVE's haptic information is used to augment visual displays. However in everyday tasks it is often the case that users cooperate on a haptic task with little visual attention given to it. In this experiment we looked at the ability to perform a cooperative manual task using visual or haptic feedback alone (proprioceptive feedback in the visual case was not informative about the relation between the users' input devices). We have shown that a cooperative task can be performed based solely on haptic communication between the users even in the presence of simulated network delay. It is important to note that the use of a spring connection between the users rather than a rigid link improved the feedback to the user allowing for sensing of direction of error using the magnitude of the force between them.

6.2 Instability in haptic teleoperation and telecollaboration

Collaborative haptic virtual environments with network delay have features in common with teleoperation of a remote machine with haptic feedback. In both cases the control is 'bilateral'. From the point of view of the user, the response of the haptic device is not strictly determined by their actions and the modeled entity but by unpredictable interactions of the controlled device at the remote site. The effect of the movements of the user on the remote device, and hence the feedback, can be modified by contact with either the physical environment or the remote user, in the telehaptic and haptic CVE respectively.

The latency incurred by transmitting the commands and force feedback can destabilize the displays. In a single user haptic VE, Choi and Tan [2003] found that the occurrence of buzzing and chattering instability depend on the type of surface simulated (e.g. its stiffness), the rendering method and the control characteristics of the device. Presumably the effects of latency on stability and performance in a haptic CVE would depend on the simulated haptic coupling between the users and

the method used to render it. For example, in our experiments a stiffer spring between the users should make the task less forgiving and more prone to instability.

Stability of haptic displays in a CVE is an important property for safety, usability and fidelity issues. However, it is important to note that in the present study, latency had effects on performance even in the absence of apparent chattering or other overt instability. Presumably this reflects a cost associated with compensating for the poorer control characteristics of the coupled haptic system.

6.3 Dissociation of VE state

Designers of CVEs face choices about how to distribute the data and simulation processing amongst the nodes in the CVE. For example, one can update the local simulation immediately in response to user input and maintain crisp interaction or can delay the local interaction to wait for contemporary data from the remote site. One possibility is to maintain a central server that controls the simulation and maintains temporal consistency at the expense of increased latency and sluggish local interaction. To reduce network demands rendering is typically done locally.

Due to stability and fidelity considerations, haptically-mediated CVE simulation must be distributed so that local haptic interactions can be simulated with low latency [Buttolo et al. 1997; Hespanha et al. 2001]. For example, users can interact with local replicas of the shared VE. Changes in the state of the VE due to the action of a user are communicated to the other nodes via the network. To minimize transmission requirements, updates are often made only when there is a change in state (e.g. [Bogsanyi and Miller 2002]). This architecture has the advantage of crisp local interaction but inconsistencies between the states of the local copies of the VE can arise due to network delay or lost packets [Mortensen et al. 2002].

With local haptic simulation, the states of the simulations can differ qualitatively as well as simply in terms of temporal offset. For example in our experiment the simulations subjects experienced often differed in whether underforce or overforce errors were being signalled, depending on the leadership context. Previously, this type of disconnect has been demonstrated for online gaming and distributed graphical simulations (eg [Pantel and Wolf 2002]) but, to our knowledge, this is the first report of this phenomenon in a haptically-mediated or purely haptic display.

6.4 Adaptation

The human motor system has remarkable ability to adapt to changing relations between sensory input and motor output. These presumably have developed to allow individuals to cope with demands imposed by growth, aging, or injury and for short term demands for flexibility to compensate for changes in load and muscle fatigue.

Neural signals to control the muscles take time to process and propagate from the brain to the muscle fibres. Similar delays occur for the sensory inputs. As people develop or grow, the temporal relationships between sensory inputs and motor commands change. To maintain sensorimotor coordination, adaptation to changing time delay is a necessary behavioural capability [Foulkes and Miall 2000].

It is not clear whether this type of adaptation can take compensate for latency in interpersonal haptic communication in collaborative environments. In the single user case, the user has access to the error in a closed loop control system and can use this to guide visuomotor adaptation. In the collaborative case it is not clear that the user can access the appropriate error term and in any case it is difficult to predict the volitional acts of the collaborating user. Furthermore it is not known what the sensorimotor adaptation capabilities of humans are under variable latency conditions which are typical of delay in CVEs due to network latency.

6.5 Compensation

Various strategies have been proposed to compensate for lag related deficiencies in CVEs including: tuning of the control response [Adams and Hannaford 2002; Andersen and Spong 1989], predictive filtering, specialized data compression, communication architectures or protocols (for review see [Shirmohammadi and Georganas 2001]), limiting interaction to one user at a time [Buttolo et al. 1997], distorting virtual objects in an attempt to allow simultaneous fast local interaction and maintain global temporal consistency [Sharkey et al. 1998], constraining the effects of user actions or movements of virtual objects [Ruddle et al. 2003], and restricting close coupled interactions [Hespanha et al. 2001].

We found some benefits of predictive filtering on performance in cooperative manipulation of virtual objects based on visual feedback. The effects of predictive filtering with a balanced experimental design that includes both visual and haptic feedback needs to be examined. We showed the benefits of prediction for simulated constant delay and the results need to be verified with more realistic network models. In general, it may also be necessary to consider the effects of the signal processing involved in predictive filtering. Predictive filtering introduces complex changes in the simulated motion and the dynamics of the error remaining are strongly dependent on the predictor [Azuma and Bishop 1995].

7 References

Adams, R. J., and Hannaford, B. 2002. Control Law Design for Haptic Interfaces to Virtual Reality. *IEEE Transactions on Control Systems Technology* 10, 1-12.

Andersen, and Spong. 1989. Bilateral Control of Teleoperators with Time Delay. *IEEE Transactions on Automatic Control* 34,5, 494-501.

Azuma, R. T., and Bishop, G. 1995. *A Frequency Domain Analysis of Head-Motion Prediction*. In Computer Graphics (Proceedings of SIGGRAPH), Los Angeles, CA, 401-408.

Basdogan, C., Ho, C.-H., Srinivasan, M. A., and 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.

Bogsanyi, F., and Miller, M. L. 2002. *Tool and Object Based Synchronization in Collaborative Haptics*. In Proceedings HAVE 2002 - IEEE International Workshop on Haptic Virtual Environments and their Applications, IEEE, 109-113.

Buttolo, P., Oboe, R., and Hannaford, B. 1997. Architectures for Shared Haptic Virtual Environments. *Computers & Graphics* 21.

Choi, S., and Tan, H. Z. 2003. *An Experimental Study of Perceived Instability During Haptic Texture Rendering: Effects of Collision Detection Algorithm*. In Proceedings 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2003, IEEE, 197-204.

Chong, N. Y., Kawabata, S. i., Ohba, K., Kotoku, T., Komoriya, K., Takase, K., et al. 2002. Multioperator Teleoperation of Multirobot Systems with Time Delay: Part I - Aids for Collision-Free Control. *Presence* 11,3, 277-291.

Durlach, N., and Slater, M. 2000. Presence in Shared Virtual Environments and Virtual Togetherness. *Presence* 9, 214-217.

Ellis, S. R., Breant, F., Menges, B., Jacoby, R., and Adelstein, B. D. 1997. *Factors Influencing Operator Interaction with Virtual Objects Viewed Via Head-Mounted See-through Displays: Viewing Conditions and Rendering Latency*. In Proc. IEEE Virtual Reality Annual International Symposium, 138-145.

Foulkes, A. J., and Miall, R. C. 2000. Adaptation to Visual Feedback Delays in a Human Manual Tracking Task. *Exp Brain Res* 131,1, 101-110.

Hespanha, J., McLaughlin, M., and Sukhatme, G. S. 2001. *Haptic Collaboration over the Internet*. In Proc. of the 5th Phantom Users Group Workshop,

Hubbold, R., and Keates, M. 2000. Real-Time Simulation of a Stretcher Evacuation in a Large-Scale Virtual Environment. *Computer Graphics Forum* 19.

Jordan, J., Mortensen, J., Oliveira, M., Slater, M., Tay, B. K., Kim, J., et al. 2002. *Collaboration in a Mediated Haptic Environment*. In Presence 2002: The 5th annual international workshop on presence,

Margery, D., Arnaldi, B., and Plouzeau, N. 1999. *A General Framework for Cooperative Manipulation in Virtual Environments*. In Proceedings of EGVE'99: 5th Eurographics Workshop on Virtual Environments,

McLaughlin, M. L., Sukhatme, G., Peng, W., Zhu, W., and Parks, J. 2003. *Performance and Copresence in Heterogeneous Haptic Collaboration*. In Proceedings of 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2003., IEEE, 285-291.

Mortensen, J., Vinayamoorthy, V., Slater, M., Steed, A., Lok, B., and Whitton, M. 2002. *Collaboration in Tele-Immersive Environment*. In Eighth Eurographics Workshop on Virtual Environments,

Pantel, L., and Wolf, L. 2002. *On the Impact of Delay on Real-Time Multi-Player Games*. In NOSSDAV 2002, ACM, Miami, FL, 23-29.

Park, K. S., and Kenyon, R. V. 1999. *Effects of Network Characteristics On Human Performance in a Collaborative Virtual Environment*. In Proceedings of IEEE Virtual Reality 1999, IEEE, 104-111.

Ruddle, R. A., Savage, J. C. D., and Jones, D. M. 2002. Symmetric and Asymmetric Action Integration During Cooperative Object Manipulation in Virtual Environments. *ACM Transactions on Computer-Human Interaction* 9.

Ruddle, R. A., Savage, J. C. D., and Jones, D. M. 2003. Levels of Control During a Collaborative Carrying Task. *Presence* 12,2, 140-155.

Ruhleder, K., and Jordan, B. 2001. Co-Constructing Non-Mutual Realities: Delay-Generated Trouble in Distributed Interaction. *Computer Supported Cooperative Work: The Journal of Collaborative Computing* 10.

Sharkey, P. M., Ryan, M. D., and Roberts, D. J. 1998. *A Local Perception Filter for Distributed Virtual Environments*. In Proc. IEEE Virtual Reality Annual International Symposium (Cat. No.98CB36180), 242-249.

Shirmohammadi, S., and Georganas, N. D. 2001. An End-to-End Communication Architecture for Collaborative Virtual Environments. *Computer Networks-the International Journal of Computer and Telecommunications Networking* 35,2-3, 351367.

Vaghi, I., Greenhalgh, C., and Benford, S. 1999. *Coping with Inconsistency Due to Network Delays in Collaborative Virtual Environments*. In ACM VRST 99, London, UK, 42-49.

Appendix

The technique for time delay compensation in this paper utilizes proprietary time delay techniques (patent pending) in order to maintain stability. Although the details of the algorithms cannot be discussed, the net effect of the techniques are described in this section




Figure 6. Closed loop haptic effects are implemented on the input device to create a real-time interaction between the destination device and its environment

A Handshake latency compensation module is located at both the input device node and the destination device node. The commands from the user are modified by this module to provide the destination device with a latency compensated command signal. In addition, the feedback signals are processed by another latency compensation module at the input device node so that the applied haptic effect is consistent with

what the destination device is experiencing at a given point in time.

Referring to Figure 6, the process is as follows:

- The User inputs a command to the control system by slewing or actuating the force feedback input device. Consider a command entered at time $t = t_0$.
- After traveling through a network delay of Δ_0 , the command input reaches the destination device at time $t = (t_0 + \Delta_0)$
- Using the $t = (t_0 + \Delta_0)$ command signal, a prediction horizon of Δ_1' where $(0 < \Delta_1' \leq \Delta_1)$, and the current measure of network latency, the latency compensation module modifies the command input. The operational objective is to modify the command input to make it appear that no communication delay was incurred. This signal is applied to the destination device via its controller.
- Sensors on the destination device generate feedback signals at $t = t_1$ (e.g. position). After traveling through a network delay of Δ_1 , the command input reaches the destination device at time $t = (t_1 + \Delta_1)$.
- Using the $t = (t_1 + \Delta_1)$ feedback signals, a prediction horizon of Δ_1' where $0 < \Delta_1' \leq \Delta_1$, and the current measure of network latency, the latency compensation module modifies the feedback signals. The operational objective is to modify the feedback signals to make it appear that no communication delay was incurred. This signal is used by the force command generator to apply an appropriate force on the user via the force feedback input device.

Using this technology, the user is able to feel in real time the effect of the destination device's interaction with its environment. This is achieved, in effect, by closing the force control loop over the network. It is for this reason that the bilateral case is especially sensitive to latency. However, a properly implemented bilateral application will lead to unprecedented remote interactive capabilities.

The performance characteristics depend a great deal on the prediction horizon Δ' . It is known that:

- As the Δ' approaches 0, performance is sluggish and approaches the performance of the case where no time delay compensation has been employed.
- As Δ' approaches the network time delay, the performance is better in that the predicted command more closely matches the actual command that occurred at $t = t_0$. However, the overall increase in performance comes at the expense of increased overshoot and noise effects. In the bilateral case, instability can also occur if the prediction is too aggressive. In general, however, this characteristic is application specific.