The Effects of Network Delay on Task Performance in a Visual-Haptic Collaborative Environment

Benjamin M. Lambeth
Information Technology
Armstrong Atlantic State University
markuslambeth@gmail.com

James LaPlant
Computer Science
Armstrong Atlantic State University
jl8111@gmail.com

Elena Clapan
Computer Science
Armstrong Atlantic State University
symonny@gmail.com

Felix G. Hamza-Lup
Computer Science
Armstrong Atlantic State University
fham76@gmail.com

ABSTRACT
Computer networks have grown considerably over the past decade. Faster and cheaper Internet connections have brought millions of PCs into a domain where rich content and fast downloads have become a necessity. New technology such as haptics must integrate into the existing infrastructures if it is to be considered a viable resource. As with visual and audio that preceded it, haptics too will find its home on the Internet. This research investigates the problems inherent in networks that haptic technology must overcome to make the step from a fascinating technology to a practical one.

Categories and Subject Descriptors
H.5.2 [**User Interfaces**]: Haptic I/O; C.2.1 [**Network Operations]**: Network Architecture and Design – Network Communications

General Terms
Measurement, Performance, Experimentation, Human Factors

Keywords
Haptics, Collaborative Virtual Environments, Network Delays

1. INTRODUCTION
Haptics, the science of tactile senses, has been a topic of research interest for over a decade. Although the interaction between haptic devices, such as Sensable's PHANToM Omni (Figure 1), and computer applications has been studied extensively, many aspects of haptic performance are still unknown. The ability to transmit data over networks, and extract meaningful information from haptic technology is one aspect that requires more research. Most of the fields where haptic technology has potential applications (e. g. medical, education, and gaming) require high performance in respect to consistency and accuracy.

ACMSE '09 March 19-21, 2009, Clemson, SC, USA.

Our research investigates the effects of network delays on a visual-haptic task performance in a collaborative virtual environment (CVE). The collaborative visual-haptic task requires the use of both visual and tactile senses to successfully complete the task. This balanced design consideration guided the implementation, assessment and analysis of the results.

The paper is structured as follows. In Section 2 we provide a review of related work. In Section 3 we discuss the experimental design by describing the visual-haptic task, the participants, assessments used, and methods of implementation. In section 4 we analyze the results of the experiment. Section 5 concludes our research and summarizes future work.

Figure 1. PHANToM Omni haptic device

2. RELATED WORK
The sense of touch in shared virtual environments has been investigated from several viewpoints. Researchers at MIT explored the communication between remote users interacting with haptic technology while measuring the contribution of haptics in user's task performance [1]. The examined task consisted of moving a ring back and forth on a randomly-curving wire without allowing the ring to touch the wire. The experiment required collaboration to move an object in a generally fixed line of motion along a single path. The experiment relied solely on subjective assessments to understand the user's perception of the sense of togetherness. The results showed that haptic feedback enhanced performance and increased the feeling of togetherness. Our research investigates the sense of togetherness by allowing participants to interact without guidance in addition to focusing on the measurable outputs of task performance in the presence of network delays.

Researchers from BTexact Technologies accomplished an early investigation of effects of network latency on a haptic task, demonstrating that latencies around 40ms, a rather small latency by Internet standards, has a significant impact on haptic feedback [2]. Our research broadens the range of measurable delays to 400ms, providing us with a better insight into the relationship between network delays and performance.

A more recent investigation of the effects of the network delay in a distributed haptic application examined the sense of togetherness [3]. In their experiment two subjects moved in the simulated environment while they attempted to maintain haptic contact with each other. A peer-to-peer network was used, and ten latencies were added to packets, ranging from 0 to 400ms. Delays of 50ms or higher were found to affect the collaboration of the subjects in the experiment. The experiments showed that the subjects became aware of the delay once it exceeded 50ms and were able to adjust their rate of progression through the experiment to compensate for the delay at 100ms. Our research takes an alternative approach to studying network delays. We utilize a client-server model and use a smaller set of delays by studying only five levels of latency. In Jay et al's experiment, augmented projection of the objects is employed so that the users will mostly rely on the haptic feedback to accomplish the task. In our research the analyzed task of stacking the cubes is more complex, and users are required to rely on both visual and haptic feedback. In addition, we created shadows to help users perceive the location of the object in the virtual scene, while in Jay's experiment the movement on the Z plane is not easy to perceive visually.

Research has been conducted on solving other problems that network delays might cause. At the Haptics and Virtual Environments Laboratory at the University of Southern California, mutual exclusion is used as a potential solution to scalability problems in collaborative haptics [4, 5]. Dead reckoning algorithms are explored by applying linear filters on a haptic signal. As our research continues to progress, we will consider such algorithms to improve the task performance in visual-haptic environments.

3. EXPERIMENTAL DESIGN

We designed a visual-haptic task consisting of stacking a set of cubes using two haptic pointers as described in the next section.

3.1 The Visual-Haptic Task

We developed a collaborative visual-haptic stacking cube task, where two remote users manipulate four cubes (Figure 2) displayed randomly in the CVE. Each of the users manipulates a haptic device that has a representation in the virtual environment of a sphere. The goal of the task is to stack all four cubes on a central stacking pedestal (Figure 3) by applying forces to opposite sides of the cubes with two haptic devices and simultaneously lifting the cubes using simulated friction. In order to accomplish this, one participant must apply force to one side of the cube with a haptic device. The other participant must apply an equal opposing force on the opposite side of the cube with the other haptic device. Once in the air, the participants must move the cube towards the stacking pedestal while maintaining the appropriate amount of force. The task is not complete until all four cubes are stacked.

The CVE monitors the amount of force applied to the cubes. Excess total force from both users (greater than 5 newtons) causes the cubes to break, while too little force causes the cube to slip and drop, potentially causing the cube to break if it lands with too much momentum (greater than 7 kg*m/s). Broken cubes reappear in a new random position in the CVE. The CVE contains boundaries that prevent cubes from leaving the range of motion of the haptic devices. Additionally, gravity and friction are simulated to provide mass and grip to the cubes.

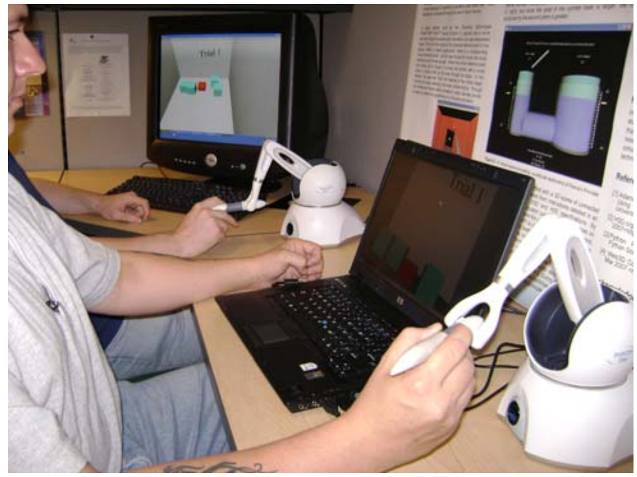

Figure 2. Two participants collaborate to stack a set of cubes. Each participant can manipulate one haptic pointer. In order to pick up the cubes they have to coordinate their moves.

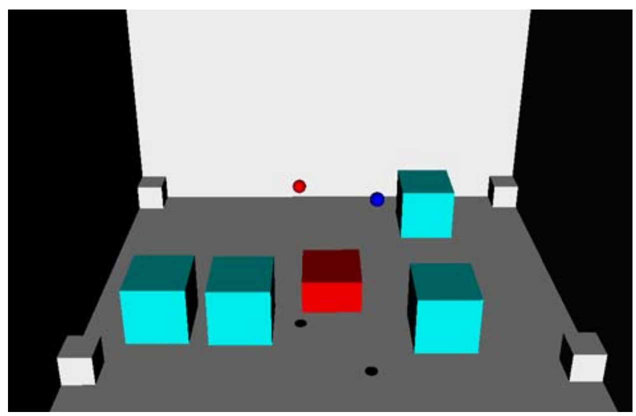

Figure 3. Collaborative Stacking Cube Task

In order to create a reasonable testing session for our participants, we simplified some aspects of the CVE that would require extensive training. The CVE incorporates only 3 degrees of freedom, allowing movement of cubes along the X, Y and Z axis. Cubes do not pitch, yaw or roll in the CVE. Furthermore, all cubes have the same dimensions and weight.

The participants performed the tasks in pairs. Each pair performed six warm-up trials to become familiar with the equipment and task. The warm trials did not incorporate cube breaking and had a near zero delay, the baseline for our experiments. After the warm-up trials, the subjects performed ten measured trials. During the measured trials, the CVE randomly inserted a generated simulated network delay, two each of 0, 50, 100, 200 and 400 milliseconds.

In addition to the simulated network delay, there were three notable, although negligible processing delays: Haptic hardware, PC and Network. The haptic hardware processing delay for the PHANToM Omni consists of the time it takes for the haptic device to register movement of the stylus and transmit the data through the six foot long FireWire cable to the FireWire port on the PC. The PC processing delay consists of the time it takes the haptic data to travel from the FireWire port to memory, the processing of data by the software and system, and the transmission of processed data to the network interface card. The Network delay consists of the time it takes the data to travel the length of the Ethernet cables (approximately 15 feet in length to each node) through the network switch to the destination network interface card.

Given the low complexity of our experiment, we assumed that the total hardware and software delay was negligible (<1ms, the resolution of the haptic device). Therefore, the only non-negligible delays were the ones we simulated. To create these artificial delays, we designed the software to wait a specified number of milliseconds before sending the packet. This delay effect was applied to packets sent by both the clients and the server, creating a total delay set of 0, 50, 100, 200 and 400 milliseconds.

3.2 Participants

22 students participated in the experiment, 15 male and 7 female. Out of the 22 participants, 20 were right handed and 2 were left handed. Five participants took part in one of our previous experiments and were already familiar with the equipment and task. The participants were notified of the assessments before the trials began. They were also located in the same room and allowed to communicate with each other verbally while completing the task.

3.3 Assessments

3.3.1 *Objective Assessments*

We established a set of objective parameters for our experiment:

1. *Time for task completion* is a measure of how well the users worked together to complete the task. By studying this parameter, we were able to identify the overall performance of the task using visual and haptic senses. Also, we determined which users worked better with each other and how their performance improved over time.
2. *The number of broken cubes* is a measure on how well the users responded to each other during the task. Because the CVE required that the users apply forces to the opposite sides of the cube, each user felt the force the other applied to the cubes. Each user had to apply an appropriate amount of force to lift the cube, without breaking it. The number of broken cubes influenced the completion time, because after breaking a cube, it was re-rendered on the floor at a random location.
3. *The quality of stacking* (Figure 4) is a measure of how well the users worked together and responded to each other. This parameter balances the previous two. While users could quickly stack the cubes without breaking them, the quality of the stack determines how well the two users collaborated within the time to stack the cubes.

3.3.2 *Subjective Assessments*

We defined a single subjective parameter to help us gauge the quality of the user's interaction with each other in the CVE.

1. *Togetherness* is the measure of how participants perceived each other both visually and tactilely in the CVE. The sense of togetherness may improve collaboration if users can perceive each other.

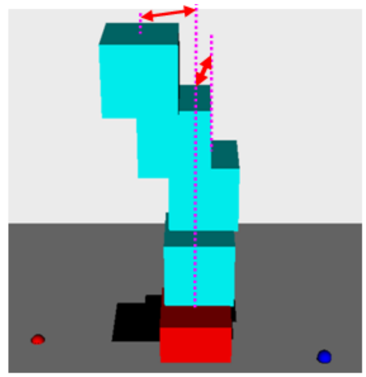

Figure 2. Quality of stacking is measured by the distance of cube center to platform center

3.4 Implementation

We used Sensable's H3D API, which combines numerous open-source standards into one software development platform, to implement the design. The scene was implemented using X3D [6, 7]. The Python scripting language was used to define the behavior of the haptic devices.

Shadows were implemented to give the users additional depth cues. They were implemented as black squares that followed the cubes along the floor, with lengths and widths equal to the cubes they represented, and positions defined by the x and z coordinates of that object.

The distributed system was implemented on a client-server platform. The client package provided the user interface and user-centric logic for haptic interaction. The client sent streaming data to the server using the User Datagram Protocol (UDP). The packets contained pointer coordinates and forces. The server processed the information received by the client and sent streaming data back to the clients, providing continuous synchronization.

We utilized three computers located in our research lab and a 5-port gigabit Ethernet switch to build up a LAN (Figure 5). The two clients were installed on comparable desktop PCs. The server and Wireshark were installed on a laptop PC. The PCs were coupled together using Cat 5e Ethernet cables ranging from 15 to 20 feet and configured in a Class C Network [8].

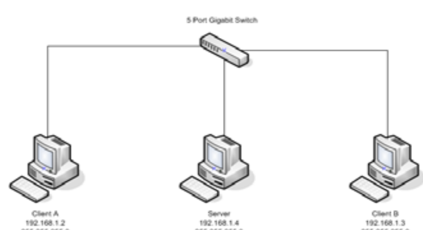

Figure 3. Network configuration

4. RESULTS

Our first assessment showed that the mean time to complete the task with 50, 100, 200, and 400 ms delays was 82, 95, 138, and 171 seconds respectively (Figure 6). The baseline was 95 seconds. These data formed a quadratic trend, demonstrating that as the network delay increased, the time to complete the task increased. The results for the second assessment showed a strong correlation with the first assessment, where the mean number of cubes broken was 1, 3, 4, and 6 for the same delays and a baseline of 2 (Figure 7). The final assessment yielded unexpected results. The data formed a cubic trend with a small variance between network delays with a mean distance of 3.3, 3.8, 3.9, and 3.7 cm with a baseline of 3.5 cm (Figure 8).

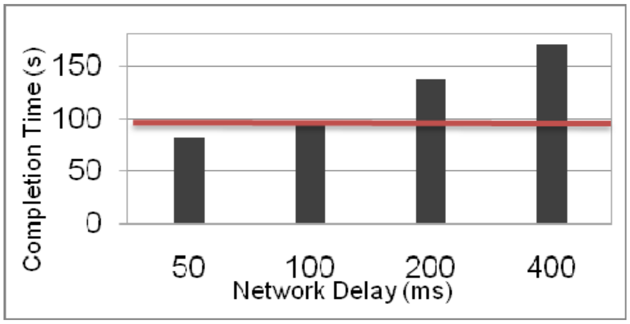

Figure 4. Mean Time to Complete Task

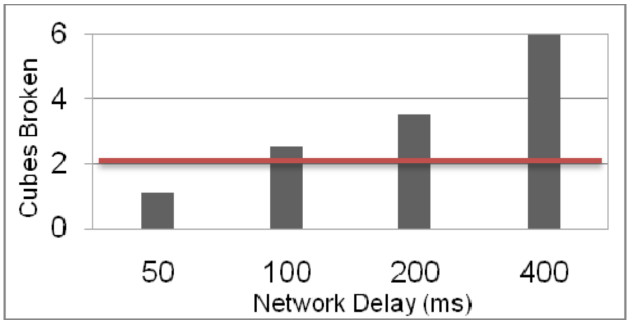

Figure 5. Mean Number of Cubes Broken

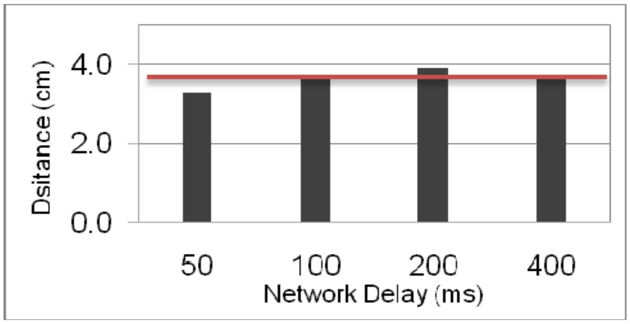

Figure 6. Mean Distance from Center

Following the experiment, we presented each participant with a survey asking them to rate how well they perceived the other user visually and tactilely, and to indicate which sense they felt the most helpful in completing the task. The results showed that 21 of 22 participants felt that they perceived the other participants through the use of haptics. However, only 5 of the 22 participants selected the tactile sense as the most useful. Most users (14) predominantly relied on the visual sense to accomplish the task, while the remaining 3 chose the auditory sense as the most helpful.

5. CONCLUSIONS AND FUTURE WORK

We presented a visual-haptic collaborative environment in which two users collaborate in order to achieve a common goal. We simulate network delay in a client-server architecture deployed over the LAN in order to investigate the latency effect on task performance. Preliminary results suggest that delays negatively affect reaction time and task completion time, but has little affect on the user's accuracy.

It is probable that the quality of stacking assessment is skewed by the fact that the cubes are locked once they reach the pedestal, Participants have only one opportunity to stack the cube properly. This will be addressed in future work to see if there is an inverse correlation between the quality of the stack and the time to complete the task.

Seemingly, there is a delay threshold (of about 100ms) for this task. Passing this threshold will negatively and substantially impact task completion. Our findings corroborate others research groups' results. Regardless of the complexity of the task, small delays can have significant impact on task performance. Even though the visual sense is still the dominant sense in virtual environments, the use of tactile enhances task performance and creates a greater sense of togetherness.

In future work, we will investigate network delays by introducing network jitter. Beyond that, we will use our CVE to explore solutions that will help minimize the negative impact that network delays create to provide a better understanding of the capabilities and limits of haptic collaboration in environments where time and distance are important.

6. REFERENCES

[1] C. Ho, C.B., M. Slater, N.Durlach, M. A. Srinivasan. *An Experiment on the Influence of Haptic Communication on the Sense of Being Together*. in *British Telecom Workshop on Presence in Shared Virtual Environments*. 1998: ACM Press.

[2] Stephen Furner, A.H., David Hands. *Haptic Transmission over IP networks: Image perception in response to increasing latency*. in *EuroHaptics 2001*. 2001. University of Birmingham, UK.

[3] Caroline Jay, M.G., Roger Hubbold, *Modeling the Effects of Delayed Haptic and Visual Feedback in Collaborative Virtual Environments.* Transactions on Computer-Human Interaction, 2007. **14**(2).

[4] McLaughlin, M.L., Sukhatme, G., Hespenha, J., Shahabi, C., Ortega, A. *Touch in Immersive Environments*. in *Conference on Electronic Imaging and the Visual Arts*. 2000. Edinburgh, Scotland.

[5] McLaughlin, M.L., Sukhatme, G., Hespenha, J., Shahabi, C., Ortega, A., Medioni, G. *The Haptic Museum*. in *Conference on Electronical Imaging and Visual Arts*. 2000. Edinburgh, Scotland.

[6] H3D, http://www.h3dapi.org/. Nov. 2008.

[7] X3D, http://www.web3d.org/x3d/. Nov. 2008.

[8] WireShark, http://www.wireshark.org/. Nov. 2008.

7. Appendix

7.1 Client A Configuration

Model:	SilverStone SG01B-Evolution
CPU:	Intel Pentium D 3.0 GHz / 1.0 GB
Network Card:	Intel PRO/1000 PL Ethernet
Connection:	Cat 5e to Port 1
Haptic Device:	Phantom Omni
OS:	Windows XP Pro v. 2002 SP2
Software:	H3D API / BACH v1.0 Client

7.2 Client B Configuration

Model:	Dell Precision 530
CPU:	Intel Xeon 2.00 GHz / 1.00 GB
Network Card:	3COM 3C920 Fast Ethernet
Connection:	Cat 5e to Port 2
Haptic Device:	Phantom Omni
OS:	Windows XP Pro v. 2002 SP2
Software:	H3D API / BACH v1.0 Client

7.3 Server Configuraton

Model:	HP Compaq nc8430
CPU:	Intel Core2 T5600 1.83 GHz / 512 MB
Network Card:	Broadcom NetXtreme Gigabit Ethernet
Connection:	Cat 5e to Port 3
OS:	Windows XP Pro v. 2002 SP2
Software:	H3D API / BACH v1.0 Server / Wireshark

7.4 Switch Configuration

Model:	Netgear GS105 PRO Safe
Type:	5 port gigabit switch