

Effects of Network Characteristics on Human Performance in a Collaborative Virtual Environment

Kyoung Shin Park and Robert V. Kenyon

Electronic Visualization Laboratory

University of Illinois at Chicago, Chicago, IL 60607, USA

(312) 996-3002 voice, (312) 413-7585 fax

kpark@eecs.uic.edu, kenyon@uic.edu

Abstract

We assessed the effects of network latency and jitter on a cooperative tele-operation task in a collaborative virtual environment. Two remote partners worked together to manipulate shared virtual objects over a network. The task was to minimize the time to transfer a ring through one of four paths with the least number of collisions. The performance of human subjects was measured and analyzed quantitatively as a function of network latency: 10 and 200 msec delays with and without jitter. Jitter had the greatest impact on coordination performance when the latency was high and the task was difficult. These results are discussed in light of current and future CVE tasks.

1. Introduction

Collaborative virtual environments (CVE) are designed to allow people in remote locations to work together over networks. People can share collaborative experiences, learn from their colleagues or teachers, work together on designing systems, or perform a complex group task through these shared virtual environments.

CVEs are applicable to cooperative spatial tasks, such as 3D architectural design and environment planning, car design and modeling, and training to repair the Hubble space telescope [7, 8, 10]. They are also useful for supporting natural spatial social skills [1, 14] (e.g. face-to-face negotiations) - people can make use of the shared virtual space as a mean of interactive negotiations with one another. They can be applied to complex cooperative tele-operation task, such as tele-surgery.

In many current CVE applications, users spend most of their time navigating in 3-D space in order to move to a particular location but spend little time in performing an action, such as manipulating objects or interacting with others. Such navigation tasks generally require neither highly intensive two-/multi-way interactions nor conflict resolution among participants (e.g. colliding each other). In contrast, more complex manipulative tasks need higher interactivity, which requires high network bandwidth and short latency.

Current CVEs transmit information about their local entities to remote sites through the network so that all sites can share the same information [5, 15]. To maintain this consistency, it is important to render remote entities in real-time so that the user will not notice any difference between local and remote entities in the environment. Thus, CVEs demand a high quality of service (QoS) on the network to maintain natural and real-time interactions among users. For example, users expect an accurate visual scene of the remote object's movements to avoid collisions between their objects and those controlled by the remote partners.

As CVEs become widespread, network QoS will be a major issue [4, 5, 9, 12]. Quality of service refers to the performance guarantees on the throughput (bandwidth), network latency, and jitter. Network latency is the time it takes to get information from one site to others through the network. For example, when a user performs an action, the information about the action is transmitted over the network, and remote users will receive the result of the action after some amount of delay. Often, networks exhibit variability in delay, called jitter can which result in a jerky presentation of remote participant's actions. Hence, CVEs mostly run on a local area network (LAN) to insure the required QoS [14, 16]. Some CVE applications use the wide area network (WAN) [10, 11], but the interactivity is reduced. A high-speed network like ATM (Asynchronous Transfer Mode) has been used in CVEs [7, 8], yet access to such a high performance network is still limited to a few research institutes and companies.

This study compares two commonly used networks, Ethernet and Integrated Services Digital Network (ISDN), to examine the tradeoff of network QoS and interactivity in a CVE. Ethernet is relatively fast and routinely used to connect CVEs. ISDN is a slow but inexpensive WAN to connect CVEs. To evaluate the effects of jitter in these networks, a constant (no jitter) latency network was simulated using a fiber-optic local network, called Scramnet.

A set of motor control tasks was developed to measure the coordination between two participants in the CVE (Figure 1). The task required cooperative manipulation of objects and conflict resolution. The task contained four levels of difficulty aimed at providing the subjects with easily negotiable interactions, and inten-




Figure 1. A cooperative tele-operative task in a networked virtual environment

sive hand-guidance interactions between users to assess the impact of the network on cooperative performance.

The purpose of this study is to measure changes in human performance generated by latency and jitter in the network connecting the CVEs. Furthermore, our focus is to identify human behaviors, adoptions and adaptations in relation to the constant or variable latency, and to understand the dynamic nature of the human-to-human coordination process.

2. Methods

Five teams of two subjects were formed. Each subject within the team controlled one object in each VE. All virtual objects were visible within both sites. However, the manipulation of a particular object was under the strict control of only one individual, thereby eliminating interference between the two subjects and allowing simultaneous manipulation of their object.

2.1. Subjects

Ten subjects from the University community volunteered as participants in the experiment. Their ages ranged from 23 to 60 with a median age of 27. All subjects were right-handed, had normal visual acuity and stereo perception. Subjects from the participant pool were randomly assigned to two person teams according to their schedules or were allowed to choose to work with a friend. All subjects were naive to the task and the purpose of the experiment.

2.2. Apparatus

Two tele-immersive environments (CAVE and ImmersaDesk) were connected using one of three differ-

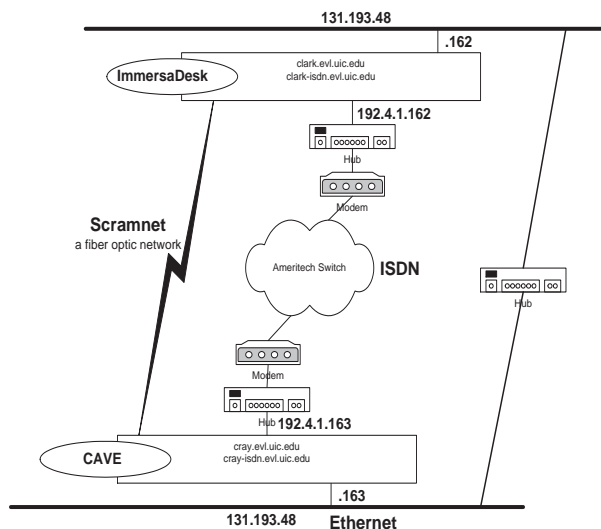


Figure 2. Illustration of computers and networks used for the experiment. CAVE and ImmersaDesk were on the same Ethernet subnet, and the IP address for ISDN was separated from Ethernet IP address. Scramnet was connected with a fiber optic.

ent kinds of networks: Scramnet, Ethernet, and ISDN (Figure 2). The configuration of ISDN was 2B-channels and the dedicated services, provided by two Ameritech ISDN phonelines. The CAVE and the ImmersaDesk are high resolution, large field of view, and projection-based immersive virtual environment systems [2]. The current configuration of CAVE displays 1028 x 768 resolution stereoscopic images at 96 Hz on each surface. The ImmersaDesk is a drafting table format VE display. It features a 67x50-inch rear-projected screen at a 45-degree angle. The screen has a sufficiently wide-angle view - e.g. 110 degrees horizontal field of view when the user stands close (within 1-foot away) to the screen.

The scenes were rendered on a Silicon Graphics Onyx Infinite Reality Engine, and the position data for the user's head and hand was obtained by using Ascension Motion Star Extended Range tracking system. Users use a wand (3D equivalent of a mouse) to interact with and control virtual entities in the CAVE and ImmersaDesk. A direct voice communication was established using wireless headset microphones and speakers in the VE systems.

The shared visual scene within each VE system was rendered at the frame rate of that individual VE system. Each system maintained consistent local and remote object models by transmitting state updates across the networks. The system used the network communication of TCP/IP based client-server distributed model for Ethernet and ISDN. However, the exchange of information over Scramnet was different from the TCP/IP network protocols. Scramnet is fiber optic reflective memory network, which consists of

two Scramnet VME cards (one in each computer) that has 128 kilobytes of memory and connected to each other over optical fiber. Scramnet has a very low and consistent latency (average round-trip latency of 300 microseconds). Data was exchanged between the two computers by writing information into the memory-mapped range on the Scramnet cards, such that the data was almost instantly shared between two machines over the fiber optic connection.

The Ethernet and ISDN one-way network latencies between the two tele-immersive systems were measured during each experimental trial. Ordinarily, the system clock between two systems must be synchronized to get precise one-way network latency. We used Scramnet to overcome the system clock synchronization problem. The actual one-way network latency over Ethernet or ISDN was measured by subtracting the arrival time of the packet over the Scramnet from the arrival time of the identical packet over the regular network.

3. Experimental Design

Subjects were examined using a 2 x 4 x 4 within-subjects factorial design in which the three main factors were the effect of exposure to the task (D1/D2), the effect of the network (i.e., latency and jitter) (NETWORK), and the task difficulty (PATH). *Exposure* refers to the first day (DAY1) or the second day (DAY2) of the two-day experiment. NETWORK refers to the case where the network connecting two VE systems had a 10-msec or 200-msec latency with or without jitter. PATH represents a different level of task difficulty, due to the length and shape (degrees of orientation) of the path.

The four network conditions studied were:

- *Scramnet-10-msec* simulates the average Ethernet latency as measured within the Lab using our tasks. The 10-msec network delay was achieved using the Scramnet network to transmit the data with a queuing delay buffer to produce a constant 10-msec delay.
- *Ethernet* refers to the condition whereby the LAN network within the Lab was used to transmit data between the two systems. The average delay was found experimentally to range from 7-msec to 18-msec. The highest jitter was found to be approximately 500-msec.
- *Scramnet-200-msec* simulates the average ISDN latency as measured over our ISDN network using our tasks. The technique to produce the Scramnet-10-msec condition delay was used generated this constant 200-msec delay.
- *ISDN* refers to the condition whereby the two systems were connected over the Ameritech ISDN network. The average delay for this connection was measured experimentally to vary between 150-msec to 300-msec. The highest jitter was found to be 2 seconds.




Figure 3. Straight, Curve, U-Turns, and Spiral

The task (based on [3]) was to minimize the time to transfer a ring through one of four paths (of varying difficulties) with the least number of collisions. Figure 3 illustrates the four paths: Straight, Curve, U-Turns, and Spiral. Team members might adopt somewhat different strategies for the interaction and cooperation to complete the tasks successfully based on the shape of each path.

- *Straight* path task was the easiest (level 1), as it required the subjects to move the ring and the path in a straight line over a 36 inch distance.
- *Curve* path task was slightly more difficult as it required a smooth movement of the objects on a gradual curved path along two axes (level 2). The length of this path was 42 inches.
- *U-Turns* path task was one of the hardest (level 3), as it required the subjects to alternate 90 degree turns to the left and right. Four sharp turns were required to complete the task such that there was a high probability of collisions at the corners. The total length of this path was 54 inches.
- *Spiral* path task was the most difficult (level 4) as it required the subjects to guide the objects around a spiraling path along with three axes. The length of this path is 67.8 inches.

Audio and *Avatar* communication channels were provided for conveying verbal and non-verbal cues between the two subjects. Subjects were encouraged to talk to each other during the experiment trial or breaks between trials using headset microphones and speakers in the CAVE and ImmersaDesk. Avatar (the remote cursor; a virtual representation of the other's wand in this context) enhanced the awareness between team members, such that a member could know what the other was doing.

3.1. Procedure

The inner diameter of the path was 1.2 inches. The diameter of ring was 9 inches. Path lengths were discussed in the experimental design above. The task was to transfer the ring through the path. However, it required each subject to move his/her object, Ring or Path, from its starting point to the end point. The task trial was considered *completed* when both partners reached the end marker with their respective objects (Figure 4). Each starting location was either on the left or right 18 inches away from the end point. The end point was located at the upper middle center




Figure 4. To complete the Straight path task, one person would need to move the ring 1.5 ft to reach the end marker and other person would need to move the rod 1.5 ft to reach the same end marker.

of the VE systems. That is, both subjects were given an equal amount of work.

Subjects in the CAVE had control of *starting* the task and *resetting* the objects to their original locations at the end of a trial. When the task was started or reset, both subjects were notified by graphical display and sound cues. The start or reset command signals might be delayed to the ImmersaDesk with the delay depending on the network latency. Subject could negotiate when to start the task with the voice communication.

The task required both subjects to grasp and move their object as fast as possible with as few collisions as possible, thus requiring both speed and accuracy. Subjects were instructed that collisions could be avoided by keeping the path in the center of the ring while they were moving their object, which should result in better performances. Accuracy was scored by using the number of collisions, i.e., number of errors. The speed was determined using the completion time of the task. The completion time and the number of collisions for each person in the CAVE and the ImmersaDesk were collected locally. The tracked hand and head motions were saved to examine the optimal movement for each trial.

At the beginning of the experiment, subjects received a standard set of instructions that described the experimental conditions, the task, and emphasized the coordination process, i.e. both high speed and no collisions. Subjects had several familiarization trials using all paths in random order in the pre-exposure stage. However, subjects were given a larger ring diameter (12 inches) than that used in the experiments (9 inches) to reduce training effects from these sessions.

A within-subjects design was used in this experi-

Table 1. Average Completion Time (sec) with four tasks over various network conditions

	Scramnet 10 msec	Scramnet 200 msec	Ethernet 10 msec	ISDN 200 msec	All Network
Straight	1.4671	1.7913	1.5844	1.8930	1.6839
Turns	4.2845	4.9647	4.7488	5.6876	4.9214
Curve	2.3063	2.7577	2.2591	3.1054	2.6071
Spiral	5.2201	5.7173	5.1431	6.7448	5.7063
All Path	3.3195	3.8078	3.4339	4.3577	

ment. Each group of subjects was tested on the same day with all of the four network conditions: Scramnet-10-msec, Scramnet-200-msec, Ethernet, ISDN. Each experimental session took about 20 minutes per condition to complete and consisted of a short discussion, three or more warm-up trials, and 4 paths with 30 trials for each path, totaling 480 trials (4 network conditions x 4 paths x 30 trials). Straight path began each session followed by U-Turns, Curve and Spiral, respectively. After completing each condition, groups took a short break (10 minutes) before proceeding to the next condition.

With two consecutive days of testing, each group was randomly assigned to one of four possible sequences of network conditions on the first day. Groups were assigned to the reverse order on the second day. At the end of the experiment, subjects completed a post-test questionnaire to give feedback about tasks and network conditions.

4. Results

4.1. Performance Analysis of the Overall Results

The data was analyzed with multivariate analysis of variance ¹ with two dependent variables : mean completion time (CT) and mean number of collisions (ERR) for all five subject groups. The experimental design in this analysis was CT and ERR by 2 (D1/D2) x 4 (NETWORK) x 4 (PATH).

Exposure to the task improved performance. The mean completion time and number of collisions were significantly lower for tasks performed on the second day compared to the first day. The average of the mean completion time on the first day was 3.8556 sec and 3.6038 sec on the second day [F(1,128)=4.092, p<0.05]. The average of mean number of collisions on the first day was 4.9755 and 3.8396 on the second day [F(1,128)=4.585, p<0.05]. However, there was no significant interaction with the other factors.

The different network conditions produced significant differences in completion times [F(3,128)=14.114, p<0.001] and number of collisions [F(3,128)=18.1702, p<0.001]. On average, the mean completion time was

¹Multivariate analysis of variance (MANOVA) is the procedure for testing any of a wide variety of null hypotheses about the effect of other variables on the mean value of several correlated dependent variables.

Table 2. Average Number of Collisions with four tasks over various network conditions

	Scramnet 10 msec	Scramnet 200 msec	Ethernet 10 msec	ISDN 200 msec	All Network
Straight	0.1806	0.1454	1.0620	0.3710	0.4398
Turns	5.4986	7.7014	8.0807	14.693	8.9935
Curve	0.6812	1.6342	1.0909	4.1512	1.8894
Spiral	3.0048	4.5744	6.1947	11.456	6.3075
All Path	2.3413	3.5139	4.1071	7.6680	

shortest for Scramnet-10-msec, followed by Ethernet, then Scramnet-200-msec, and finally ISDN (Table 1).

Notice that subjects using ISDN took more than 1 second longer to complete the task compared to Scramnet-10-msec condition. The mean number of collisions also followed the same pattern as completion times where collisions progressively increased from the Scramnet-10-msec, to Scramnet-200-msec, to Ethernet, to ISDN (Table 2). Post hoc multiple comparison tests revealed that completion times for Scramnet-10-msec were significantly less than Scramnet-200-msec and ISDN. Subject's completion times for ISDN were clearly longer than other three networks. However, completion times for Ethernet and Scramnet-200-msec were not significantly different. Also, Post hoc tests on number of collisions showed that ISDN was significantly different from other three networks.

The four paths with their different levels of difficulty produced significantly different mean completion times [$F(3,128)=231.687, p<0.001$], and mean number of collisions [$F(3,128)=55.346, p<0.001$]. Post hoc tests revealed that the mean completion times for each path were significantly different from one another. Post hoc tests on the mean number of collisions revealed that Straight and Curve paths are not distinct from each other, but these two paths as a group along with Spiral and U-Turns as a group are significantly different from one another. The mean completion time of Straight path task was shorter than Curve, U-Turns, and Spiral (Table 1). U-Turns task produced the highest collision rate because of its four sharp corners. Spiral also proved to be difficult as seen by the high number of collisions. In contrast, Straight and Curve paths did not produce as many collisions as did U-Turns and Spiral paths (Table 2).

Our analysis also showed that NETWORK x PATH interaction on the dependent variable of the mean number of collisions was also significant [$F(9,128)=3.189, p<0.05$]. This indicated that easier network conditions and easier tasks produced better performance (less number of collisions). For example, the least number of collisions resulted when using Scramnet-10-msec with a Straight path while the most resulted using




Figure 5. Task performance over networks and tasks. The mean completion time over networks was depicted as the line and the mean number of collisions over networks was drawn with the bar.

ISDN with the U-Turns path. In general, the longer latencies affected the performance to a greater degree when subjects were working on more difficult tasks.

D1/D2 had a significant main effect, but there was no significant interaction with other factors (i.e. D1/D2 x NETWORK, D1/D2 x PATH, and D1/D2 x NETWORK x PATH). Therefore, we combined the D1/D2 data for each network and path condition. Figure 5 (based on Table 1 and 2) shows the mean completion time and number of collisions for the four paths using the four network conditions from all five subject groups.

The one-way ANOVA using the mean completion time from each path condition showed that there was a significant main effect of network when Curve [$F(3,36)=6.986, p<0.05$], U-Turns [$F(3,36)=4.756, p<0.05$], and Spiral [$F(3,36)=4.994, p<0.05$] paths were used. In other words, for the simple Straight path task, there was no significant difference in mean completion time among the different networks used. Post hoc tests revealed that the mean completion time using Ethernet and Scramnet-10-msec networks is significantly different from ISDN for both Curve and Spiral path conditions.

Scramnet-10-msec is significantly different from ISDN for the U-Turns path condition. The one-way ANOVA using the mean number of collisions for each path showed that there was a significant main effect of network when Straight [$F(3,36)=3.276, p<0.05$], Curve [$F(3,36)=7.288, p<0.05$], U-Turns [$F(3,36)=6.627, p<0.05$], and Spiral [$F(3,36)=8.289, p<0.05$] paths were used. Post hoc tests on the mean number of collisions revealed that Scramnet-10-msec and Scramnet-200-msec networks are significantly dif-




Figure 6. Histogram of average network latency for all trials.

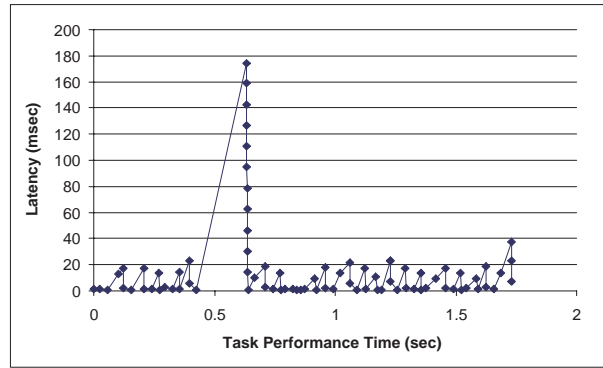


Figure 7. Individual packet latencies as a function of time during one trial: Group 2 Ethernet Curve Trial 3.

ferent from Ethernet for the Straight path condition. ISDN is significantly different from other three networks in Curve, U-Turns, and Spiral path conditions.

4.2. Effects of Variable Network Latency on Performance

Our previous analysis showed that performance was affected by the network conditions used. However, the effects of jitter on performance using Ethernet or ISDN conditions need to be examined further. In this section, the average and variance of the network latency was calculated from the latencies accumulated for each trial. Then, the correlation between the average and variance latency and the performance (on CT and ERR) in each trial was analyzed. The correlation coefficient for Ethernet latency vs. performance was not significant, indicating that there was no evidence of linear relationship between latency and performance. However, there was a significant correlation between completion times and average ISDN latencies ($r = 0.208$) or its variance ($r = 0.135$). Also, the correlation between number of collisions and average ISDN latencies ($r = 0.229$) or its variance ($r = 0.185$) was significant.

The histogram (Figure 6) of average network latencies for each trial using Ethernet and ISDN showed large variations in their distributions, indicating how variable latencies were from trial to trial. The distributions show a tail toward higher latencies in both networks. The average latency within a single trial for Ethernet ranged from 1.824 msec to 65.065 msec (the mean of 12.5352 msec) and for ISDN ranged from 76.091 msec to 2.052 sec (the mean of 263.5807 msec). Note that the maximum latency on Ethernet was still much smaller than the minimum latency on ISDN.

4.3. Effects of Peak Latency on Performance

Given the effect of jitter on performance shown in the previous section, the relationship between performance and the latency over 30 trials was analyzed. When we examined a sequence of average network latencies over 30 trials we found that several trials had very high average latencies compared to the rest of the trials. These peaks occurred sporadically within 30 trials depending on the network traffic. It occurred more frequently and more erratically with ISDN than Ethernet due to the low bandwidth of ISDN.

When we plotted the individual packet latencies for these high latency trials as a function of time we found the pattern of individual packet latencies shown in Figure 7. The saw-tooth behavior displayed in Figure 7 shows that one of the packets experiences a very long delay before it is received, perhaps due to a packet loss or other kinds of network errors. This delayed packet effectively delays all subsequent packets. Once the packet is successfully transmitted, all the subsequent packets follow in a short period of time. This continues until the latency reaches the baseline or until another packet is significantly delayed. As the number of these events increased the average latency increased.

Task performance suffered significantly when high average latencies occurred, compared to the average performance across 30 trials. The magnitude of performance degradation associated with the occurrence of peak latencies was highest using ISDN. In particular, some latency peaks for ISDN were 1 or 2 seconds and resulted in a high number of collisions and long completion times. Basically, subjects could not continue their actions when this happened.

Examining completion times across all trials showed that stable performance was disrupted by peak (very high average latency). This would take a form of a sharp rise in the completion time. Subsequent completion times would improve until the next peak was observed, or reached the baseline when there was no peak. In contrast, the number of collisions disappeared almost completely after the peak because subjects seemed to slow down and become more deliberate, thereby lowering the completion time by reducing the number of collisions.

5. Discussions

Overall performance on our set of tasks was affected by the characteristics of the network used to connect the CVEs. The largest effects were seen when the long latency, high jitter ISDN network was used. However, long latency without large jitter showed a much lesser affect on performance. For instance, we found no significant difference in overall performance between networks having latencies of 200 msec without jitter and 10 msec with jitter (Ethernet). In addition, our results also reveal that accuracy is degraded when jitter is present in the network, as shown by the higher number of collisions for the Ethernet and ISDN networks (both of whom have jitter) compared to the two Scramnet conditions without jitter.

The magnitude of the effects of network latency and jitter on performance was found to be a function of task difficulty, which is similar to the results from the single operator ring/path task [3]. Networks with long latencies and large amounts of jitter had their greatest impact on performance during difficult hand guidance tasks that required constant visual feedback for successful completion. For instance, for the straight path using ISDN, the subjects could ignore the delayed presentation of the remote object and move their objects to the target using ballistic movements. Note that the straight path represents a navigation task that is used mostly in current CVEs. Collaborative work using such simple tasks can be done successfully over ISDN since the task does not require many partner dependent time critical interactions among other people [6, 10]. However, when the task is very difficult (e.g., Spiral), even with a short delay, the subjects tend to be more deliberate in order to avoid collisions. With large network latency and jitter, such a task is much more difficult because the information about the remote object's position is delayed. As a result, subjects adopt the move-and-wait control strategy that appears like micro sequential actions on their movements.

The correlation between performance and latency

across trials indicates that the large variations in ISDN latencies are indeed a very critical factor in the resulting user performance; however, the variations in the short Ethernet latencies in our laboratory's LAN are not significant. Furthermore, examining the time history of individual packet latencies for Ethernet and ISDN shows a dominant sawtooth behavior occurring sporadically during the trials, which also influence to the performance. Consequently, the delay experienced by the user is not always white noise-like but has structure, which is visible to the user. We believe that the initiation of this pattern is the result of an error in the transmission of the packet caused by a network collision or lost packet condition. Due to the fact that TCP/IP guarantees that packets will arrive in the correct sequence, all subsequent packets are pending until the error is cleared and the offending packet is successfully transmitted. Then all the pending packets are sent unless or until there is another network error. We observed this behavior more frequently in ISDN possibly due to its narrow bandwidth.

Greater than 1 or 2 seconds of the highest jitter, producing high average and variance values for latency within a trial, can often be found in ISDN. It appears that high average latencies strongly correspond to performance results. When the average latency in a trial was a 1-second or higher, either the completion time was very long or the number of collisions increased. Interestingly, performance on tasks following this trial with long network delay was also degraded although performance improved with time similar in form to a learning curve.

6. Conclusions and Future Work

The results using our CVE tasks demonstrate the role of task difficulty and the characteristics of network latency and jitter on cooperative performance. The task difficulty quantifies the level of coordination needed between two users in the CVE for a given network's characteristic. The overall task performance shows that high latency can impact users' coordination. The variability of the network latency reduces the ability of the subjects to use prediction in performing the task.

Task visualization is a key problem in tele-operation since most of the operational control decisions in VE are based on visual feedback, i.e. hand-eye coordination [13]. With large network latency and jitter in the CVE, the information of the remote entities is delayed inconsistently. As a result, users may have different views in the shared workspace. The results show that users wait several seconds to synchronize

their views first then continue to perform the task, i.e. move-and-wait control strategy in time-delayed teleoperation without predictive display. Thus completion time increases with poor network conditions because of reduced predictability available for the visual feedback and consequently increasing the amount of effort needed to maintain a constant level of performance.

The difficulty of defining service of CVE networks stems from the lack of empirical data on the QoS requirements. The results of this study provide some empirical data, which will be used to define service of CVE networks in the near future. In general, Ethernet provided a reasonable quality of service in terms of throughputs and latencies to be used in a collaborative virtual environment. ISDN with its limited bandwidth often shows high latencies, but the cooperative work can be done successfully when the task does not involve intensive interactions among partners, as shown in [6, 10]. Thus, the QoS issues should be discussed in terms of the kind of task to be accomplished as well as its difficulty.

The results of this study suggest that future designers of CVE should consider how to increase the user's confidence in their motor control coordination during task execution. Then it will enhance the training capabilities of CVEs by reducing the trial-and-error approaches users apply to successfully interact with their partners. One of the continuing goals will be to address the quality of interaction over various networks with other types of collaborative tasks, such as spatial orienteering and awareness.

7. Acknowledgements

We thank all subjects, who participated in this study, for investing their time and efforts to conduct the successful human performance experiment in the collaborative virtual environment. Our thanks to Stephen Ellis for suggesting the ring/path task used in this study. This research was supported by NSF Grant # IRI-9424272. CAVE and ImmersaDesk are trademarks of the Regents of the University of Illinois.

References

- [1] S. Benford, J. Bowers, L. E. Fahlen, C. Greenhalgh, J. Mariani, and T. Rodden. Networked virtual reality and cooperative work. *Presence*, 4(4):364–386, 1995.
- [2] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The cave automatic virtual environment. *Communications of the ACM*, 35(2):64–72, June 1992.
- [3] S. R. Ellis, F. Breant, B. M. Menges, R. H. Jacoby, and B. D. Adelstein. Operator interaction with virtual objects: effects of system latency. In *Proceedings of HCI '97 International*, pages 973–976, 1997.
- [4] T. A. Funkhouser. Network topologies for scalable multi-user virtual environments. In *Proceedings of the Virtual Reality Annual International Symposium*, pages 222–228. IEEE Computer Society, IEEE, March 1996.
- [5] R. Gossweiler, R. J. Laferriere, M. L. Keller, and R. Pausch. An introductory tutorial for developing multiuser virtual environments. *Presence*, 3(4):255–264, 1994.
- [6] H. Ishii, K. Arita, and T. Yagi. Beyond videophones: Teamworkstation-2 for narrowband isdn. In *Proceedings of the Third European Conference on Computer Supported Cooperative Work*, pages 325–340, September 1993.
- [7] V. D. Lehner and T. A. DeFanti. Distributed virtual reality: Supporting remote collaboration in vehicle design. *IEEE Computer Graphics and Applications*, in press, 1997.
- [8] J. Leigh, A. E. Johnson, and T. A. DeFanti. CALVIN: An immersimedia design environment utilizing heterogeneous perspectives. In *Proceedings of IEEE International Conference on Multimedia Computing and Systems '96*, June 1996.
- [9] J. Leigh, A. E. Johnson, and T. A. DeFanti. CAVERN: A distributed architecture for supporting scalable persistence and interoperability in collaborative virtual environments. *Virtual Reality Research, Development and Applications*, 1997.
- [10] R. B. Loftin. Hands across the atlantic: <http://www.vet1.uh.edu/sharedvir/handat1.html>. Technical report, Virtual Environment Technology Laboratory, 1995.
- [11] M. R. Macedonia, D. P. Brutzman, and M. J. Zyda. NPSNET: A multi-player 3D virtual environment over the internet. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics*, pages 93–94. ACM, ACM, 1995.
- [12] M. R. Macedonia and M. J. Zyda. A taxonomy for networked virtual environments. In *Proceedings of the 1995 Workshop on Networked Realities*, Oct 1995.
- [13] T. Poston and S. Serra. Dextrous virtual work. *Communications of ACM*, 39(5):37–45, 1996.
- [14] M. Roussos, A. Johnson, J. Leigh, C. R. Barnes, C. A. Vasilakis, and T. G. Moher. The NICE project: Narrative, immersive, constructionist/collaborative environments for learning in virtual reality. In *ED-MEDIA/ED-TELECOM*, 1997.
- [15] S. K. Singhal and D. R. Cheriton. Exploiting position history for efficient remote rendering in networked virtual reality. *Presence*, 4(2):169–193, 1995.
- [16] H. Takemura and F. Kishino. Cooperative work environment using virtual workspace. In *Proceedings of the Conference on Computer Supported Cooperative Work*, pages 226–232, 1992.