

Preliminary STAR TAP Tele-Immersion Experiments between Chicago and Singapore

Jason Leigh

National Center for Supercomputing Applications &
Electronic Visualization Laboratory

Kyoung Park, Robert V. Kenyon,
Andrew E. Johnson, Thomas A. DeFanti
Electronic Visualization Laboratory,
University of Illinois at Chicago

Hong-Yee Wong

Institute of High Performance Computing
National University of Singapore

Abstract

Tele-Immersion is the merging of audio and video conferencing with collaborative virtual reality, data-mining and significant computation systems. As Tele-Immersion often involves distantly located participants all interacting with each other in real-time, it relies heavily on the quality of the underlying networks used to distribute information between participants. This paper presents our preliminary findings of testing a number of tele-immersive applications over the STAR TAP link between Chicago and Singapore. The goal is to be able to prescribe and deploy a suite of Tele-Immersion networking experiments around the world for testing and comparing the quality of their links so that predictions can be made of the suitability of the link for supporting different classes of Tele-Immersion applications.

Introduction

Tele-Immersion (Leigh97b) is the integration of audio and video conferencing with collaborative virtual reality (VR) in the context of data-mining and significant computation. When participants are tele-immersed, they are able to see and interact with each other in a shared environment. This environment persists even when all the participants have left. The environment may control supercomputing computations, query databases autonomously and gather the results for visualization when the participants return. Participants may even leave messages for their colleagues who can then replay them as a full audio, video and gestural stream.

CAVERN (the CAVE Research Network) is a collection of participating industrial and research institutions equipped with CAVEs, ImmersaDesks, and high-performance computing resources all interconnected by high-speed networks for the purpose of supporting tele-immersive- engineering and design; education and training; and scientific visualization and computational steering. In 1992 there was a single CAVE in Chicago, in 1998 there are over 80 CAVE and ImmersaDesk installations around the world. One of the problems facing this growing community, and the VR community as a whole, is how to best provide a mechanism to support long term collaborative work between them. CAVERNsoft (Leigh97a, 97c) is our architecture for achieving the goals of Tele-Immersion. Part of CAVERNsoft's research goals is to develop techniques to evaluate the networks over which the tele-immersive activities take place. This paper describes our preliminary experiences and results of connecting over a collection of high speed/bandwidth networks called STAR TAP, between the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago and the Institute of High Performance Computing (IHPC) at the National University of Singapore.

STAR TAP -- Science, Technology And Research Transit Access Point -- is a model for international high-performance networking that will facilitate scientific collaborations on such computationally demanding and data-intensive topics as global climate change, protein structures, nano-materials, video streaming, and interactive virtual reality. Supported by a three-year, \$1.2 million grant from the National Science Foundation, the STAR TAP project intends to provide a persistent connection point for U.S. and

international high-performance research networks and to improve the speed and performance of the applications that will run over these networks. One of STAR TAP's goals is to incorporate new tools for diagnostics, usage tracking, and direct feedback on performance. Such statistics will help scientists and network engineers understand the bandwidth and latency requirements of various applications so that they may tune both the application and the network for optimal performance (STAR TAP98).

Our Chicago/Singapore STAR TAP experiments were performed on June 17, 1998 between 8pm and 11pm Central Standard Time. The networking route taken was from EVL to MREN (Metropolitan Research and Education Network) to vBNS (very high-performance Backbone Network Service) to SINGAREN (Singapore Internet Next Generation Advanced Research & Education Network) to IHPC and vice-versa.

The goals were as follows:

1. To debug connectivity between EVL and IHPC by working together to temporarily by-pass any firewalls and to get EVL software to run interconnected at both sites.
2. To perform preliminary gathering of qualitative as well as quantitative network-related data. Specifically:
 - NICE (a tele-immersive educational application) was run between the remote sites and the coordination between the avatars observed.
 - COVE (a collaborative tele-operation experiment) was run between the remote sites and the performance of the human subjects was recorded.
 - Raw bandwidth and latency data over the network was recorded.

By gathering this data we will be able to predict the kinds of human tasks, in Tele-Immersion, that are possible across similar trans-oceanic links. These kinds of experiments will also allow us to isolate bottlenecks in the networks and the Tele-Immersion systems, and make appropriate recommendations for improvements. Finally, it will allow us to prescribe an experimental suite that can be deployed at other Tele-Immersion sites around the world, for testing and comparing the quality of their trans-oceanic links. The experiments in this report represent only a preliminary attempt at collecting, building and testing the tools and procedures for such an experimental suite.

The Experiments

1.1 NICE

NICE (Roussos97, 98) is a collaborative environment designed for young children. NICE depicts a virtual island on which grows a virtual garden that children may tend. The children, represented by avatars, collaboratively plant, grow, and pick vegetables and flowers.




The children can shrink down to the size of a mouse and crawl under the garden to see the root system, and can talk with the other remotely located children or other characters in the scene. The children are able to modify the parameters of this small ecosystem to see how it affects the health of the garden.

NICE uses two central servers located at EVL to maintain consistency across all the participating virtual environments. On one server, NICE uses unreliable UDP to share avatar information from magnetic trackers. On another server, NICE employs reliable TCP to share world state information and to dynamically download models from WWW servers using the HTTP 1.0 protocol.

Although NICE was initially intended as a program for children, in the past, we have found it useful for testing connections between remotely networked CAVEs. At Supercomputing 97 we performed just such a test between 17 CAVEs and ImmersaDesks around the world (Johnson98). One standard qualitative test that is typically done when NICE is deployed for the first time at a new CAVE or ImmersaDesk, is the wand-wave test. This test involves each participant repeatedly raising the wand and lowering it. As the wand is brought to the top of the waving motion the participant utters the word "UP". As the wand is brought down to the bottom of the motion the participant utters "DOWN." The purpose of this test is to roughly compare the latency of the audio (via long distance phone call) with the latency of the underlying network used to carry avatar data. In our experiment the perception was that there was little to no difference between what was heard over the audio line to the observed motion of the participant. That is, the latency for the audio appeared roughly the same as that for network traffic. This may mean that the audio and the data may have been delivered over the same networks. With this initial test we went on to perform more quantitative experiments.

1.2 COVE



COVE (Park97) is a cooperative tele-operation task in which two remote participants work together to manipulate shared virtual objects. The task is to minimize the time to transfer a ring through one of four rods (Straight, Curve, Jagged Turns, and Spiral) with the least number of collisions. The task requires each subject to move his/her object, either the ring or rod, from its starting point to the end point. The task trial is considered completed when both partners reach the end point with their respective objects. Each starting point is either on the left or right 1.5 feet away from the ending location. The end point is located at the upper middle center of the CAVE.

The inner diameter of rod is 0.1 feet (1.2 inches). The inner diameter of ring is 0.75 feet (9 inches). The Straight rod trial is the easiest as it only requires the subjects to move the ring and rod in a straight line over a 3.0- feet-distance. The Curved rod trial is slightly more difficult as it requires that the subjects move the ring and rod on a gradual curved path along two axes. The Jagged Turns rod trial is one of the most difficult as it requires the subject to guide the objects

around alternating 90-degree turns. The Spiral rod trial is the most difficult as it requires the subjects to guide the objects around a spiraling path along three axes.

In previous experiments at the Electronic Visualization Lab (Park97), the task was performed between a CAVE and an ImmersaDesk. 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 (to eliminate jitter SCRAMNET reflective memory was used to simulate a constant network delay.) Using reliable TCP, the average latency for local area Ethernet and ISDN was approximately 10 msec and 200 msec respectively. The highest jitter was found to be 500 msec for Ethernet and 2 seconds for ISDN. Results of the COVE experiment under those situations showed that jitter had the greatest impact on coordination performance when the latency is high and the task is difficult. In fact it was found that in general, jitter had a greater impact on coordination performance than latency.

For the Chicago/Singapore/STAR TAP connection, we began by taking a preliminary measurement of the round-trip latency. This was found to be approximately 255ms. Based on our previous experiments we predicted that this STAR TAP connection would yield results similar to those produced by using ISDN or SCRAMNET simulating a 200ms delay.

In our STAR TAP experiment only two (Curve and Spiral) of the four paths were used. These two paths were chosen as they were found to be sufficient to mirror the kinds of 2D and 3D manipulations found in collaborative VR. Since the amount of time allocated to perform the experiment was limited we were only able to gather data for one pair of subjects. However for each complete run of the experiment the entire task was reversed so that the ring bearer became the rod bearer and vice versa. Both remote participants performed the experiment in CAVEs.

Tables 1, 2 and 3 compare previously collected results with the new results from the STAR TAP experiment. These preliminary results suggest that the average completion time for the STAR TAP experiments were similar to those found in previous ISDN experiments with one significant difference- the average number of collisions (and hence errors) was much lower in the STAR TAP results than those found in all the previous experiments. Part of the reason for this low collision count is believed to be because of the low jitter found in the STAR TAP link (see section 1.3.2.) Another reason is that while performing the experiment, it was observed that the subjects were employing a more sequential manipulation strategy than previously observed. That is, the rod bearer would wait until the ring had moved almost completely to the end point before finally ending the trial by moving the rod to its end point. The employment of such a sequential coordination strategy is usually associated with subjects working under high latency conditions. This however was not believed to be the cause here. This occurrence clearly indicates that future experiments of this kind should more carefully enforce simultaneous, tightly coordinated action, than what was required in COVE. A third possible reason for the low collision count could be the fact that both subjects were operating in CAVEs. In previous experiments, one of the subjects operated an ImmersaDesk. Performing the experiments in the CAVEs meant that both subjects could benefit from the improved field of view and freedom of movement within the environment. Despite this, the results do seem to indicate that the STAR TAP connection between Chicago and Singapore is highly useable for low bandwidth Tele-Immersion. The results also seem to indicate that previous results of the COVE experiment can be used to some degree to predict results on other networks.

Additional information about the COVE experiment can be found by consulting (Park98).

	Average Completion Time (s)	Average # of Collisions
Scramnet 10ms (all paths)	3.3195	2.3413
LAN Ethernet (all paths)	3.4339	4.1071
Scramnet 200ms (all paths)	3.8078	3.5138
ISDN (all paths)	4.3577	7.6680
Scramnet 10ms (Curve/Spiral)	3.7632	1.8430
LAN Ethernet (Curve/Spiral)	3.7011	3.6428
Scramnet 200ms (Curve/Spiral)	4.2375	3.1043
ISDN (Curve/Spiral)	4.9251	7.8037
STAR TAP (Curve/Spiral)	4.8011	0.7586

Table 1. A Comparison of Average Completion Time and Number of Collisions with Various Network Media. Scramnet 10 and Scramnet 200 means that the networking connection between the VR devices were created using a reflective memory pipe that simulates a constant latency (without jitter) of 10 and 200 ms respectively.

	Average Completion Time (s)	Average # of Collisions
Straight (all prev networks)	1.6839	0.4398
Curve (all prev networks)	2.6071	1.8894
Turns (all prev networks)	4.9214	6.3075
Spiral (all prev networks)	5.7063	8.9935
Curve (STAR TAP)	3.3621	0.3214
Spiral (STAR TAP)	6.1441	1.1667

Table 2. A Comparison of Average Completion Time and Number of Collisions with Various Rod Shapes.

	Average Completion Time (s)	Average # of Collisions
Curve (ISDN)	3.1054	4.1512
Spiral (ISDN)	6.7448	11.4563
Curve (STAR TAP)	3.3621	0.3214
Spiral (STAR TAP)	6.1441	1.1667

Table 3. A Comparison of Average Completion Time and Number of Collisions with Two Specific Rod Shapes over ISDN and over the Chicago/Singapore STAR TAP link.

1.3 Raw Data Gathering

In this phase of the testing, a number of standard measurements of the bandwidth and latency of the Chicago/Singapore/STAR TAP link, were taken. These include:

1. Traceroute to verify that the networking traffic was indeed going through the ATM link from EVL to Singapore.
2. Latency shortly after the COVE experiment.
3. Latency on the following day. Latency measurements were taken for a period of just over 24 hours at half-hour intervals (from 8am June 18, 1998 to 8am June 19, 1998).
 - One set of measurements involved performing 50 pings with a delay of 1 second between pings.
 - The other experiment involved sending 100 ping packets in rapid succession (flood ping).
4. Bandwidth measurements were taken using `nttcp`. 2048 1K byte-sized packets were sent over reliable TCP from EVL to IHPC and vice-versa.

1.3.1 Results of Tracing Packet Route

The results of `traceroute` are shown below. The results basically confirm that the route taken by the data packets were indeed over STAR TAP and not by regular Internet. In addition the results show that most of the latency occurs in the hop between Chicago and Singapore rather than within their respective local networks.

From Singapore to Chicago:

traceroute to laurel.evl.uic.edu (131.193.48.164), 30 hops max		40 byte packets			
		Probe 1	Probe 2	Probe 3	
		RoundTrip(ms)	RoundTrip(ms)	RoundTrip(ms)	
1	gw1-lis.ar.singaren.net.sg (202.8.95.1)	1	1	1	
2	gw1-vbns.ar.singaren.net.sg (202.8.94.14)	253	253	253	
3	uic-mren.gw.uic.edu (198.32.130.24)	255	254	254	
4	BATM-20.GW.UIC.EDU (128.248.120.20)	254	255	255 (ttl=251!)	
5	lpcard12.eecs.uic.edu (131.193.32.12)	255	(ttl=59!) 256	(ttl=59!) 255	(ttl=59!)
6	laurel.evl.uic.edu (131.193.48.164)	255	(ttl=249!) 255	(ttl=249!) 258	(ttl=249!)

From Chicago to Singapore:

traceroute to 202.8.95.12		(202.8.95.12),	30 hops max	40 byte packets				
			Probe 1		Probe2		Probe3	
			RoundTrip(ms)		RoundTrip(ms)		RoundTrip(ms)	
1	eecsevl.gw.uic.edu	(131.193.48.1)	1	(ttl=64!)	2	(ttl=64!)	1	(ttl=64!)
2	eeecs.gw.uic.edu	(131.193.32.1)	3		3		1	
3	BATM-15.GW.UIC.EDU	(128.248.120.15)	3		2		3	
4	BATM-16.GW.UIC.EDU	(128.248.120.16)	32	(ttl=253!)	7	(ttl=253!)	34	(ttl=253!)
5	aads.vbns.net	(198.32.130.14)	4	(ttl=252!)	4	(ttl=252!)	5	(ttl=252!)
6	gw1-atm8-0-5.ar.singaren.net.sg	(202.8.94.13)	255	(ttl=251!)	256	(ttl=251!)	258	(ttl=251!)
7	202.8.95.12	(202.8.95.12)	257	(ttl=250!)	257	(ttl=250!)	259	(ttl=250!)

1.3.2 Results of Latency Shortly After COVE Experiments

Flood pings were performed shortly after the COVE experiments were completed. The round-trip latency results were:

```

Mean:          255.939ms
Max:           309.156ms
Min:           254.591ms
Packet loss:   3%
4435 packets sent, 4291 received.

```

1.3.3 Results of Just Over 24 Hour Latency Measurements

50 Pings at 1 Second Intervals, every 30 minutes.

```

Mean:          256.3541ms
Max:           378ms
Min:           254ms
Median:        256ms
SD:            3.896739ms

```

Flood Pings of 100 packets every 30 minutes.

```

Mean:          255.74ms
Max:           263.76ms
Min:           254.66ms
SD:            0.8526ms

```

Notice that the flood pings showed relatively stable performance. This lack of jitter is the kind of stability that is optimal for closely coordinated tele-immersive interaction.

1.3.4 Results of Bandwidth Measurements

Bandwidth from Singapore to EVL ranged between:

```

1.3753Mbits/s to 1.5061Mbits/s.
Mean: 1.417Mbits/s.

```

Bandwidth from EVL to Singapore ranged between:

```

0.7828Mbits/s to 1.2567Mbits/s.
Mean: 1.12Mbits/s.

```

1.4 Summary of Results and Discussion

From the NICE experiment we were able to qualitatively determine that casual Tele-Immersion is definitely feasible over the Chicago/Singapore/STAR TAP link. The perceived motion of the virtual participants synchronized well with audio that was delivered simultaneously via a long distant phone call.

Round trip latencies were found to be approximately 255ms with lower jitter than was observed over local area ISDN. Bandwidth was slightly higher from Singapore to Chicago than from Chicago to Singapore at 1.5061Mb/s and 1.257Mb/s respectively.

Based on the measurements of round trip latency the COVE experiment was able to reasonably predict the outcome of the performance of remotely connected participants. The COVE experiment confirmed that, not only is casual interaction possible, but relatively closely coordinated interaction is also possible. On the other-hand performing this experiment over STAR TAP also revealed a limitation in COVE. That is, the COVE task did not always guarantee that both participants were simultaneously manipulating the object. This meant that in some instances coordinated interactions occurred sequentially rather than in parallel. Some provision for enforcing parallel interaction will be needed in future experiments. In addition, it will also be useful to record all COVE interactions with timestamps that are synchronized at both sites so that the perceived location of virtual objects and the real location of the corresponding object can be visualized simultaneously. This will allow us to examine closely how jitter directly impacts interaction.

Finally, since Tele-Immersion usually includes a video conferencing or large database transmission component, it will be important to perform these same experiments in a situation where the network would be partially loaded with data from audio, video, and domain data streams. These kinds of experiments will allow us to establish some base levels of human performance under varying network conditions. And in a manner that COVE has been able to predict the outcome of the STAR TAP experiments, these new experiments will be able to predict human behavior in similar classes of tele-immersive applications. However, in order to guarantee this predictability, performance data from real tele-immersive applications must still be collected. It is therefore necessary to build tools, beyond the obvious networking tools, that can be used to monitor coordinated activity in tele-immersive applications.

We believe that this preliminary set of qualitative and quantitative experiments represents a good template for the kinds of experiments that should be used to evaluate upcoming STAR TAP connections. One of the continuing goals of CAVERNsoft will be to collect, build and test tools and procedures to form an experimental suite that can be deployed at other Tele-Immersion sites around the world, for testing and comparing the quality of their trans-oceanic links and the quality of interaction over those links.

1.5 Appendix: Computer Specifications

Singapore Computer

Silicon Graphics Onyx2

Machine name: omni.ihpc.nus.edu.sg

Machine IP: 202.8.95.12

4 195 MHZ IP27 Processors

Main memory size: 1024 Mbytes

Instruction cache size: 32 Kbytes

Data cache size: 32 Kbytes

Secondary unified instruction/data cache size: 4 Mbytes

Graphics board: InfiniteReality2

Integral Fast Ethernet: ef0, version 1 <- for the first IP addr of omni.ihpc.nus.edu.sg (137.132.15.60)

Chicago Computer

Silicon Graphics Onyx2

Machine name: laurel.evl.uic.edu

Machine IP: 131.193.48.164

FPU: MIPS R10010 Floating Point Chip Revision: 0.0
 CPU: MIPS R10000 Processor Chip Revision: 2.6
 8 195 MHZ IP27 Processors
 Main memory size: 2048 Mbytes
 Instruction cache size: 32 Kbytes
 Data cache size: 32 Kbytes
 Secondary unified instruction/data cache size: 4 Mbytes
 Graphics board: InfiniteReality2
 Integral Fast Ethernet: ef0, version 1

1.6 Acknowledgements

EVL's virtual reality research and collaboration programs would not be possible without major funding from the National Science Foundation (NSF), notably grants CDA-9303433, IRI-9424272, and NCR-9712283. In addition we would like to thank the members of Institute of High Performance Computing of the National University of Singapore for their generous assistance in preparing for these experiments.

1.7 References

Johnson, A., Leigh, J., Costigan, C., "Multiway Tele-Immersion at Supercomputing." In IEEE Computer Graphics and Applications 1998, Pp 6-9.

Leigh, J., Johnson, A., DeFanti, T., "CAVERN: A Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments." In Virtual Reality: Research, Development and Applications, Vol 2.2, December 1997, Pp 217-237.

Leigh, J., DeFanti, T., Johnson, A., Brown, M., Sandin, D., "Global Tele-Immersion: Better than Being There,." In the proceedings of 7th International Conference on Artificial Reality and Tele-Existence. Tokyo, Japan, Dec 3-5, 1997, Pp 10-17.

Leigh, J., Johnson, A., DeFanti, T., "Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments,." In the proceedings of Supercomputing '97 San Jose, California, Nov 15-21, 1997.

Park, K., Kenyon R. V., The COVE Experiment. <http://www.evl.uic.edu/park/cove/>, 1998.

Park, K. The COVE Experiment. Masters Thesis, University of Illinois at Chicago, 1997.

Roussos, M., Johnson, A., Leigh, J., Vasilakis, C., Barnes, C., Moher, T. The NICE Project. <http://www.ice.eecs.uic.edu/~nice>, 1997.

Roussos, M., Johnson, A., Moher, T., Leigh, J., Vasilakis, C., Barnes, C., "Learning and Building Together in an Immersive Virtual World." To appear in Presence, 1998.

STAR TAP Web Page: <http://www.startap.net/>, 1998.