

## LOCAL AND GLOBAL IMPACT OF MESSAGE CACHING IN SHARED VIRTUAL ENVIRONMENTS

HELMUTH TREFFTZ, IVAN MARSIC  
Rutgers — The State University of New Jersey  
Department of Electrical and Computer Engineering  
Piscataway, NJ 08854-8058 USA  
{trefftz,marsic}@ece.rutgers.edu

### ABSTRACT

The use of Shared Virtual Environments is growing in areas such as multi-player video games, military and industrial training, collaborative design and engineering. At the same time, heterogeneity in computing power and graphics capabilities of the participating computers arises naturally as the number of people/organizations sharing a virtual environment grows. This paper presents an adaptive mechanism to reduce bandwidth usage and optimize the use of computing resources of heterogeneous computers participating in a shared virtual environment based on caching of both outgoing and incoming messages. We also report the results of implementing the proposed scheme in a simple shared virtual environment.

**Keywords:** Shared Virtual Environments, Message Caching, CSCW.

### INTRODUCTION

Shared Virtual Environments (SVEs) allow multiple geographically separated users, to interact in almost real-time [1]. SVEs have been successfully deployed for several years in applications such as military [2,3] and aerospace training [4]. The availability of higher bandwidth, as well as the increasing power and low cost of personal computers with 3D graphic accelerators, make new SVE applications accessible for the general population. The applications include education [5-7], multi-user games [8] and concurrent engineering [9].

As SVEs leave closed and controlled environments, such as military simulations, their designers need to take into account:

- The possibly large number of users that will coexist in the SVE at a given time
- The possibly great heterogeneity of the machines that will run the SVE at a given time.

As the number of users in the SVE increase, so does the bandwidth required to update the position and orientation of moving entities, as well as other types of events that need to be transmitted to maintain a consistent

shared state. If the number of participants is too large or the bandwidth too small, the network might easily become a bottleneck [10]. Since the network is a shared resource, the scope of its optimization is global and should be performed at execution time, since the number of users sharing the virtual world may vary in time.

On the other hand, the SVE should locally adapt to the computing and graphics capabilities of the host machine. There are several parameters that can be fine-tuned in order to maintain an acceptable degree of fidelity while accommodating the limitations of the local machine, such as:

- using less-detailed representations of virtual objects that are outside of the user's area of interest,
- reducing the polling frequency of local devices (such as wands, space mice, trackers, etc.),
- reducing the accuracy and/or scope of collision detection,
- reducing the frequency of local update of remote events.

Generally, the designer of the virtual environment takes into account the type of machines the SVE will run on and creates a program that will run reasonably well on the target computers. The designer is thereby making a static decision with global impact. We argue that the scope of these optimizations should be local and, furthermore, done adaptively in run-time. M. Capps has recently explored tuning fidelity parameters in Shared Virtual Environments according to user preferences [11].

In the current version of our system, we apply our scheme statically to reduce resource utilization both on a local and a global scale.

### Message Caching

One of the most important challenges in implementing an SVE is maintaining a consistent description of the virtual world across the collaborators in the presence of frequent updates. In this sense, the differentiating factor among the architectures is the location where the database describing the current state of

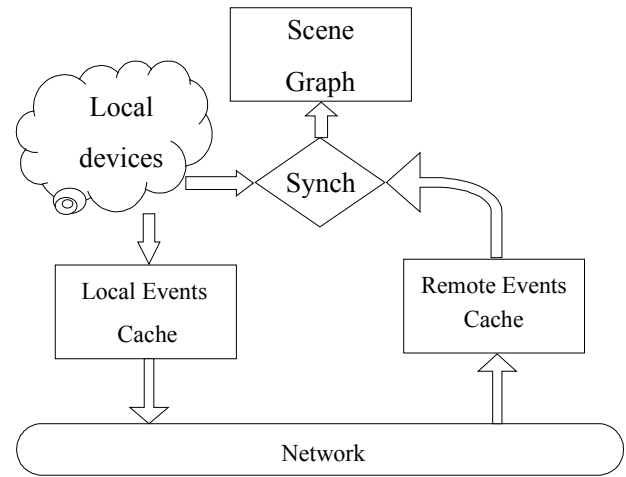
the virtual world is maintained. On one end of the spectrum are the **centralized** systems, in which the description of the virtual world is kept on a server. The server broadcasts the virtual world updates to the clients. On the other end are the **replicated** or **peer-to-peer** architectures, in which the status of the virtual world is kept on every participating site. Each node is responsible for keeping the local copy consistent and up-to-date. We have taken the latter approach for the system described in this paper. There are numerous combinations in between these two extremes, such as systems that assign a “home node” to each entity. In this case, the state of the entity is maintained in the home node only.

The challenge of maintaining a consistent description of the virtual world is similar to many problems in distributed systems, such as the problem of consistency in distributed databases or the cache-coherence problem in multi-processor machines; and designers of SVEs can apply similar solutions.

Aggregation of messages is described in [10] as a technique to reduce bandwidth utilization by grouping together sets of packets sent to the network. Bandwidth saving takes place since only one header is used for the whole group. Singhal reports that the aggregation can be done according to time (sending the update whenever a pre-defined timeout expires) or to packet number (sending the update whenever a certain number of packets be ready).

The **message caching** scheme presented here is different in several ways as will be detailed later. The general architecture of our system is described in Figure 1. The program was written mostly in Java and Java3D [14], except for the driver of the Space Mouse, which was written in C. The main components of the system are:

- **Scene Graph:** The data structure that contains the current state of the shared world. The display thread (Java3D) reads this data structure and renders the scene from the user’s virtual viewpoint.
- **Local Devices:** Drivers for the local devices used for interaction in the virtual world.
- **Local Events Cache (LEC):** Cache for locally generated events. An entry is kept in the outgoing-messages-cache for every entity that has been modified locally. Local updates are displayed and reflected in this cache as soon as they occur, but they are not sent to the network immediately. Instead, an outgoing-time-out (OT) is defined and the updates are sent when the time-out expires. In the meantime, new updates to a local entity overwrite any older update to the same entity existing in the Local Events Cache.



**Figure 1: Architecture of the system.**

- **Remote Events Cache (REC):** Cache for remotely generated events. An entry is kept for every remotely modified entity. Again, remote updates are not displayed immediately. Instead, an incoming-time-out (IT) is defined and the updates are reflected in the local scene when the time-out expires. In the meantime, a new update to a remote entity overwrites any older update to the same entity existing in the IMC.
- **Synch:** Object that guarantees mutually exclusive access to the Scene Graph.

Our approach is different from the one described in [12] in the following essential ways:

- Our scheme handles not only out-going messages but also in-coming ones.
- We propose not only grouping the messages but also sending only the latest update for each modified entity. Our scheme behaves more like a cache.

The global and local impacts of the OMC and the IMC are explained next.

### Caching of Outgoing Messages

Caching the local updates before sending them to the network has a *global* impact on the use of the network. Most of the time, users in a SVE are either moving in the world or manipulating an object. In both cases, the updates will be cached in the LEC and, depending on the value of the OT, the number of updates sent to the network will be reduced considerably. A larger value of OT will avoid having a fast machine flooding the network and the slower machines with too frequent updates.

The frame rate (number of times the scene is rendered in the local computer per unit of time) is important when determining the value of OT and IT. The frame rate of the fastest machine defines a lower

boundary for OT. Sending updates more frequently than the fastest machine can display them would be a waste of resources. Defining OT as the reciprocal of the frame rate of the *slowest* machine would make the updates unnecessarily slow for everybody else. For this reason we decide to define OT as the reciprocal of the median frame rate among the participating machines.

$$OT = \text{median} \{ (1/FR_I), I: \text{for all participating nodes} \}$$

where  $FR_I$  is the Frame Rate of node I.

### Caching of Incoming Messages

Caching the remote updates before updating the scene locally has a *local* impact. The local data structure describing the Virtual World becomes a critical shared resource that can be updated by either local or remote events in order to avoid inconsistencies. Caching remote updates favors local updates, in which the user is probably more interested.

Slower nodes use caching of incoming messages to further reduce the flooding effect of frequent updates generated by faster machines. Messages are received and stored in the local cache, but the scene is not necessarily updated immediately. Participating computers with a faster frame rate than the reciprocal of OT do not need to slow down the rate of remote updates, but slower computers do, in order to favor local updates. For this reason we are using the following formula for IT:


$$IT = \max \{ (1 / FR) \times (N + 1), OT \}$$

where FR is the Frame Rate of the node and N is a factor related to the importance of local events relative to remote events.

In a system that polls the devices once per frame, if N is 1, IT will tend to accommodate for 1 local event for every remote event; if N is 2, IT will accommodate for 2 local events for every remote event, and so on. In a system that has separate threads for handling events and for rendering, N+1 will simply indicate how many frames will elapse, in average, before the local scene is updated with in-coming events.

### Caching and Global vs. Local Optimization

Caching outgoing events reduces the bandwidth usage, since the number of messages sent by unit of time is reduced. The impact is especially meaningful when the participating nodes share a segment or a local area network or when the bandwidth is a scarce resource. Caching outgoing events also reduce the time other nodes will use to process remote events. Obviously the time-out values which determine the caching must be chosen judiciously: if the outgoing-time-out is too large the



**Figure 2: Typical display of the Shared Virtual Environment.**

remote events will look jumpy<sup>1</sup>, if it is too small there will be no benefit.

Caching incoming messages, on the other hand, has only a local effect. If computing resources are scarce, in most applications it is best to give priority to updates generated from local events, since they will most likely be of most importance for the local user.

## EXPERIMENTAL SETUP

In order to evaluate the effects of our message-caching scheme, we implemented a simple virtual environment that allows users to furnish an office in a collaborative manner. Most interactions are accomplished by manipulating a Logitech 3D-mouse [13]. Figure 2 shows a typical display during the process of arranging the furniture of the virtual office. Three computers were used for the experiments as shown in Table 1.

**Table 1: Experimental setup.**

Computer	Processor(s)	Memory	Graphics Card
BOCHICA	Dual Pentium III 700 MHz	1 Gigabyte	FireGL 1000
MORLAK	Pentium II 400 MHz	256 Megabytes	CirrusLogic GCD 54
DRAGONFIRE	Pentium II 350 MHz	32 Megabytes	Diamond SpeedStar

The computers were connected to a 100 Mbps local area network. All participating nodes shared a multicast group to send and receive updates from the peers.

<sup>1</sup> Introducing dead reckoning could alleviate the jumpiness. If dead reckoning is used, though, the remote events will appear delayed.

The experiment consists of the following tasks:

**Setup time:** The program is started in every participating computer. An object is inserted and selected for manipulation on each computer.

**Stage 1** (0 – 20 secs.): Recording of the different times starts. No manipulation takes place. The Frame Rate baseline is obtained.

**Stage 2** (20 – 40 secs.): The locally selected object is manipulated at the slowest machine only (Dragonfire).

**Stage 3** (40 – 60 secs.): A selected object is manipulated at each machine, generating events as frequently as each computer allows.

**Stage 4** (60 – 80 secs.): Manipulation stops at all nodes. At the end of this stage recording finishes.

At each computer the current time (in milliseconds) was saved for each of the following events:

- Each displayed frame.
- Each update of the Scene Graph with a local event.
- Each update of the Scene Graph with a remote event.
- Each message sent to the multicast group.

The experiment trials were arranged according to Table 2.

**Table 2: Experiment trials.**

	No Caching	OT = IT = 250 ms	OT = IT = 500 ms
2 Nodes	Trial 1	Trial 2	Trial 3
3 Nodes	Trial 4	Trial 5	Trial 6

## RESULTS

We observed a positive impact of our scheme in the following areas:

- Reduced bandwidth requirement.
- Increased frequency of local updates
- Increased frame rate.

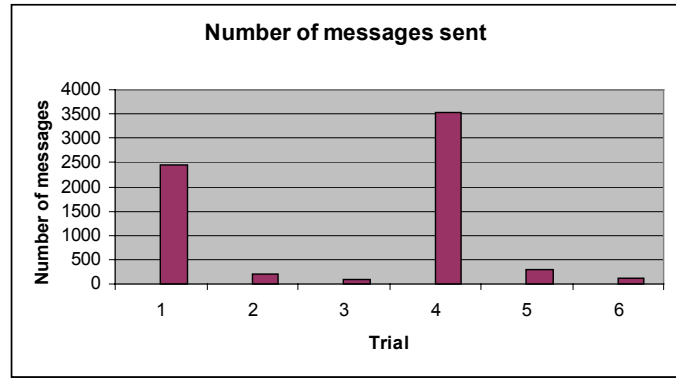
The impact is explained in the following sections.

### Reduced Bandwidth

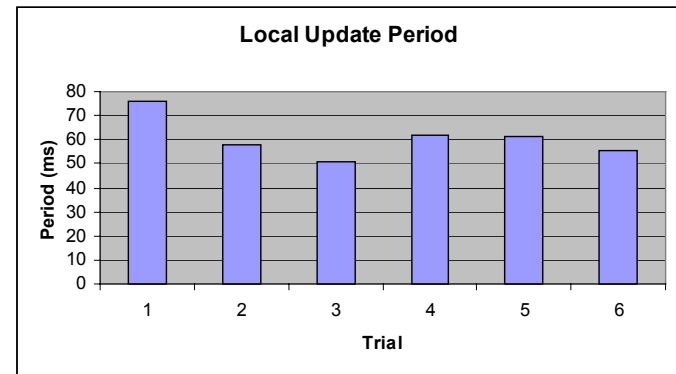
As mentioned before, the effect of caching out-going messages is reducing the amount of utilized bandwidth. This rather evident fact is shown in Figure 3, which describes the number of messages sent by all participating computers during the trials. The number of messages is larger for trials 4, 5 and 6 since three computers are running, instead of two.

### Local Updates

In the presence of caching, the user actions are reflected faster on the local scene graph, as described in Figure 4. Locally generated updates are given priority



**Figure 3: Total number of messages sent during each trial.**



**Figure 4: Average local Update Period in milliseconds for Dragonfire. Period reduces as the caching time-out values increase.**

over remotely generated events as the caching time-out times increase.

Note that the local Update Period is considerably smaller than the Frame Period for Dragonfire, the slowest computer in the setup. In our system, local devices are handled by a separate thread, instead of being polled once per frame. This allows for more frequent generation of events. From a local perspective, this is a waste of resources, but it allows faster computers sharing the same environment to receive more frequent updates.

### Frame Rate

The frame rate of slower machines is affected adversely as the number of remote events increases. Caching reduces the negative impact, bringing the frame rate closer to the baseline level, as described in Table 3 and Figure 5. Table 3 shows the average frame rate period for **Dragonfire** during each stage and Figure 5 shows the percentage increment of stages 2 (only local updates) and 3 (local and remote updates), relative to stage 1 (baseline).

In trials 1, 2 and 3, two nodes were participating, in trials 4, 5 and 6 three nodes were participating. It is easy to see the trend towards larger frame rate periods as the number of nodes increase. The most dramatic increment is 20%, when three nodes take part of the environment without message caching.

On the other hand, the impact of the remote events is reduced as the caching time-out values grow from 0 ms (tasks 1 and 4) to 250 ms (task 2 and 5) and to 500 ms (task 3 and 6).

**Table 3: Average Frame Rate Period in milliseconds. Caching reduces the flooding impact of remote events.**

Trial	Stage 1	Stage 2	Stage 3
1	1042	1133	1174
2	1197	1252	1294
3	1125	1160	1165
4	1170	1255	1415
5	953	1032	1023
6	1111	1160	1180

## CONCLUSIONS AND FUTURE WORK


We have implemented a message-caching scheme that can be used, in run time, to perform global and local optimization of resource utilization. Our initial experiments show the positive impact on several critical factors of SVEs, most interestingly bandwidth requirements and display frame rate.

For this initial set of experiments, the incoming and outgoing message time-out values were statically and globally defined. In the near future we will actually implement the dynamic computation of these parameters, during runtime, based on the Frame Rate of each node.

Future Work: Our main interest is to enable effective collaboration in Shared Virtual Environments even when the participating computers have varying degrees of computing power. We plan to devise similar schemes to work with other parameters such as level-of-detail and frequency to poll local devices, always with the objective of optimizing both local and global usage of resources in run-time.

Further information about the Rutgers University DISCIPLÉ project is available at:

<http://www.caip.rutgers.edu/disciple/>



**Figure 5: Percentage increment of Frame Rate Period. The flooding impact of remote events (stage 3) is reduced by larger caching time-out values.**

## ACKNOWLEDGMENTS

The research reported here is supported in part by NSF KDI Contract No. IIS-98-72995, DARPA Contract No. N66001-96-C-8510 and by the Rutgers Center for Advanced Information Processing (CAIP).

## REFERENCES

- [1] S. Singhal and M. Zyda. *Networked Virtual Environments*. (New York: Addison-Wesley, 1999).
- [2] D. Cohen. DIS—Back to basics. *Proceedings of the 11<sup>th</sup> Workshop on Standards for Distributed Interactive Simulations*, pp.603-617, September 1994.
- [3] NPSNET Web Site:  
<http://www.npsnet.nps.navy.mil/npsnet>
- [4] R.B. Loftin and P. Kenney. The use of virtual environments for training the Hubble space telescope flight team. Available online at:  
<http://metropolis.vetl.uh.edu/Hubble/virtel.html>
- [5] V. Pantelidis and L. Auld. VR in the schools. Available online at:  
<http://www.soe.ecu.edu/vr/pub.htm>
- [6] M. Roussos, A.Janson, J. Leigh, C. Barnes, C. Vasilakis, and T. Moher, The NICE project: Narrative, immersive, constructionist/collaborative environments for learning in virtual reality, *Proceedings of ED-MEDIA/ED-TELECOM 97*, Calgary, Canada, pp.917-922, June 1997.

- [7] Trefftz H., C. Correa, M. A. Gonzalez, J. Restrepo, C. Trefftz. Virtual reality and distance learning in Colombia, *IATED International Conference on Computer Graphics and Imaging*, Halifax, Nova Scotia, Canada, June 1998.
- [8] idSoftware Web site: <http://www.idsoftware.com>
- [9] K. Hartung, S. Muench, L. Schomaker, T. Guiard-Marigny, B. Le Goff, R. MacLavery, J. Nijtmans, A. Camurri, I. Defee and C. Benoit. Development of a system architecture for the acquisition, integration and representation of multimodal information, Deliverable 4, ESPRIT Project 8579 – MIAMI, 1996.
- [10] S. Singhal. Using projection aggregations to support scalability in distributed simulation. *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS'96)*, Hong Kong, 1996.
- [11] M. Capps. Fidelity optimization in distributed virtual environments. PhD thesis, Naval Postgraduate School, Monterey, CA, June 2000.
- [12] S. Singhal and M. Zyda. *Networked Virtual Environments*. pp.191-195, (New York: Addison-Wesley, 1999).
- [13] Magellan / Space Mouse User's Manual. LogiCad3D GmbH, a Logitech company. Gilching, Germany, 1999.
- [14] H. Sowizral, K. Rushforth, M. Deering. *The Java 3D API Specification*. 2<sup>nd</sup> Ed, (New York: Addison-Wesley, 2000).