RING: A Client-Server System
for Multi-User Virtual Environments

Thomas A. Funkhouser
AT&T Bell Laboratories [‡]

Abstract

This paper describes the client-server design, implementation and experimental results for a system that supports real-time visual interaction between a large number of users in a shared 3D virtual environment. The key feature of the system is that server-based visibility algorithms compute potential visual interactions between entities representing users in order to reduce the number of messages required to maintain consistent state among many workstations distributed across a wide-area network. When an entity changes state, update messages are sent only to workstations with entities that can potentially perceive the change – i.e., ones to which the update is visible. Initial experiments show a 40x decrease in the number of messages processed by client workstations during tests with 1024 entities interacting in a large densely occluded virtual environment.

CR Categories and Subject Descriptors:
[**Computer Graphics**]: I.3.7 Three-Dimensional Graphics and Realism – *Virtual Reality.*

Additional Key Words and Phrases: Visual simulation, multi-user systems, virtual reality, 3D virtual environments, real-time graphics, client-server design, distributed systems.

1 Introduction

In a multi-user visual simulation system, users run an interactive interface program on (usually distinct) workstations connected to each other via a network. The interface program simulates the experience of immersion in a virtual environment by rendering images of the environment as perceived from the user's simulated viewpoint. Each user is represented in the shared virtual environment by an entity rendered on every other user's workstation, and multi-user interaction is supported by matching user actions to entity updates in the

shared virtual environment. Applications for these systems include distributed training simulations, collaborative design, virtual meetings, and multiplayer games.

A difficult challenge in multi-user visual simulation is maintaining consistent state among a large number of workstations distributed over a wide-area network. Since three dimensional rendering at interactive rates requires fast access to the geometric database, shared portions of the virtual environment (including dynamic entity states) are replicated on every participating workstation. As a result, whenever any entity changes state (e.g., moves) or modifies the shared environment, an appropriate update must be applied to every copy of the database in order to maintain consistent state (see Figure 1).


Figure 1: Multi-user systems must maintain consistency between entities (A, B, C, and D) replicated on multiple workstations.

Implementing visual simulation systems for large numbers of users is especially challenging because updates can occur at extremely high rates. If N entities move through a shared virtual environment simultaneously, each modifying its position and/or orientation M times per second, then $M * N$ updates are generated to a shared database per second. Moreover, updates must be propagated to participating workstations in near real-time since large variances or delays in updates can result in visually perceptible jerky or latent motion, and thus may be disturbing to users. As a result, general-purpose distributed database systems are not adequate for use in multi-user visual simulation applications, and special-purpose messaging protocols are typically used to maintain consistent state in multi-user visual simulation systems [9, 13].

2   Previous work

Numerous experimental virtual reality systems and multi-player games have been developed for real time interaction in shared virtual environments. Unfortunately, most existing systems do not scale well to large numbers of simultaneous users.

Reality Built For Two [2], VEOS [4], and MR Toolkit [14] are multi-user virtual reality systems that maintain consistent state among N workstations by sending a point-to-point message to each of N-1 workstations whenever any entity in the distributed simulation changes state. This approach yields $O(N^2)$ update messages during every simulation step (see Figure 2), and thus does not scale to many simultaneous users before the network gets saturated.


Figure 2: Systems using point-to-point connections pass $O(N^2)$ update messages (labeled arrows) during each simulation step.

SIMNET [5], NPSNET [17], and VERN [3] use broadcast messages to send updates to all other workstations participating in a virtual environment at once. Although, this approach cuts down on the total number of messages transmitted to $O(N)$, every workstation still must process a message whenever any entity in the distributed simulation changes state (see Figure 3). Since every workstation must store data and process update messages and/or simulate behavior for all N entities during every simulation step, these systems do not scale beyond the capabilities of the least powerful participating workstation. Experiences with SIMNET and NPSNET show that a significant percentage of every workstation's processing capability is used just to read update messages from other workstations during large simulations; and, therefore, broadcast protocols are not practical for more than a few hundred users on inexpensive workstations [17].

In order to support very large numbers of users (> 1000) interacting simultaneously in a distributed virtual environment it is necessary to develop a system design and communication protocol that does not require sending update messages to all participating hosts for every entity state change. Kazman has proposed a system design, called WAVES, in which message managers mediate communication between hosts, possibly culling irrelevant messages [10, 11]. His approach is very similar to the one presented in this paper. One difference is that this paper presents algorithms and experimental results for visibility-based message culling during large simulations.


Figure 3: Systems using broadcast messages pass only $O(N)$ updates each simulation step. But, every workstation still must process every update message.

3   Overview of Approach

This paper describes a system (called RING) that supports interaction between large numbers of users in virtual environments with dense occlusion (e.g., buildings, cities, etc.). RING takes advantage of the fact that state changes must be propagated only to hosts containing entities that can possibly perceive the change – i.e., the ones that can see it. Object-space visibility algorithms are used to compute the region of influence for each state change, and then update messages are sent only to the small subset of workstations to which the update is relevant.

The key idea is illustrated in Figure 4. Although entities A, B, C, and D (filled circles) all inhabit the same virtual environment, very little visual interaction (hatched polygons) is possible due to the occlusion of walls (solid lines). In fact, in this example, only one visual interaction is possible – entity A can see entity B. Therefore, only one update message must be sent for each update to entity B's position in real-time (to the workstation with entity A). All other entities need not distribute any update messages in real-time since they are not visible to any other entity. From this example, we see that it is possible to greatly reduce the number of messages passed in real-time to maintain consistent state among multiple entities in a densely occluded environment using line-of-sight visibility to determine the region of influence for each update.


Figure 4: A system that culls messages based on entity-entity visibility may be able to reduce the number of messages processed by each workstation in densely occluded environments.

The following section describes the RING system design. Results of experiments with the system are presented in Section 5, while a discussion of alternate approaches and possible future work appears in Section 6. Finally, Section 7 contains a brief summary and conclusion.

4   RING System Design

RING represents a virtual environment as a set of independent *entities* each of which has a geometric description and a behavior. Some entities are static (e.g., terrain, buildings, etc.), whereas others have dynamic behavior that can be either autonomous (e.g., robots) or controlled by a user via input devices (e.g., vehicles). Distributed simulation occurs when multiple entities interact in a shared virtual environment by sending messages to one another to announce updates to their own geometry or behavior, modifications to the shared environment, or impact on other entities.

Every RING entity is managed by exactly one *client* workstation. Clients execute the programs necessary to generate behavior for their entities. They may map user input to control of particular entities and may include viewing capabilities in which the virtual environment is displayed on the client workstation screen from the point of view of one or more of its entities. In addition to managing their own entities (local entities), clients maintain surrogates for some entities managed by other clients (remote entities). Surrogates contain (often simplified) representations for the entity's geometry and behavior. When a client receives an update message for an entity managed by another client, it updates the geometric and behavioral models for the entity's local surrogate. Between updates, surrogate behavior is simulated by every client.

Communication between clients is managed by *servers*. Clients do not send messages directly to other clients, but instead send them to servers which forward them to other client and server workstations participating in the same distributed simulation (see Figure 5). A key feature of this client-server design is that servers can process messages before propagating them to other workstations, culling, augmenting, or altering them. For instance, a server may determine that a particular update message is relevant only to a small subset of clients and then propagate the message only to those clients or their servers. In addition, a server may send clients auxiliary messages that contain status information helpful for future client processing. Finally, a server may replace some set of messages intended for a client with another (possibly simpler) set of messages better suited to the client's performance capabilities. The aim of this client-server design is to shift some of the processing burden away from the client workstations and into servers so that larger, more affordable, multi-user visual simulation systems can be built using primarily low-cost client workstations.

In the current implementation, RING servers forward update messages in real-time only to other servers and clients managing entities that can possibly "see" the effects of the update. Server-based message culling is implemented using precomputed line-of-sight visibility information. Prior to the multi-user simulation, the shared virtual environment is partitioned into a spatial subdivision of *cells* whose boundaries are comprised of the static, axis-aligned polygons of the virtual environment [1, 15]. A visibility precomputation is per-


Figure 5: RING servers manage communication between clients, possibly culling, augmenting, or altering messages.

formed in which the set of cells potentially visible to each cell is determined by tracing beams of possible sight-lines through transparent cell boundaries [15, 16] (see Figure 6). During the multi-user simulation, servers keep track of which cells contain which entities by exchanging "periodic" update messages when entities cross cell boundaries. Real-time update messages are propagated only to servers and clients containing entities inside some cell visible to the one containing the updated entity. Since an entity's visibility is conservatively over-estimated by the precomputed visibility of its containing cell, this algorithm allows servers to process update messages quickly using cell visibility "look-ups" rather than more exact real-time entity visibility computations which would be too expensive on currently available workstations.


Figure 6: Cell-to-cell visibility (stipple) is the set of cells reached by some sight-line from anywhere in the source cell (dark box) passing only through transparent portals (dash lines) and no opaque walls (black lines). It is a useful, pre-computed overestimate of the visibility of any entity resident in the source cell.

As an example of RING server operation, consider the flow of messages between clients A, B, C, and D for the entities shown in Figure 4 connected to servers in the topology shown in Figure 5. Figure 7 shows the surrogates (small squares labeled by entity) and flow of update messages (arrows labeled by entity) for each of the four entities in this example.

- *If entity A is modified:* client A sends an update message to server X. Server X propagates that message to server Y, but not to server Z because entities C and D are not inside cells in the cell-to-cell visibility of the cell containing entity A. Server Y forwards the message to Client B which updates its local surrogate for entity A.

- *If entity B is modified:* client B sends an update message to server Y. Server Y then propagates that message to servers X and Z, which forward it to clients A and C. Server Z does not send the update message to client D because the cell containing entity D is not in the cell-to-cell visibility of the cell containing entity B.

- *If entity C is modified:* client C sends an update message to server Z. Server Z propagates that message to server Y, which then forwards the message to Client B. Server Z does not send the message to either server X or client D because neither is managing entities in the visibility set for entity C.

- *If entity D is modified:* client D sends an update message to server Z. Server Z does not forward the message to any other server or client because no other entity can potentially see entity D.
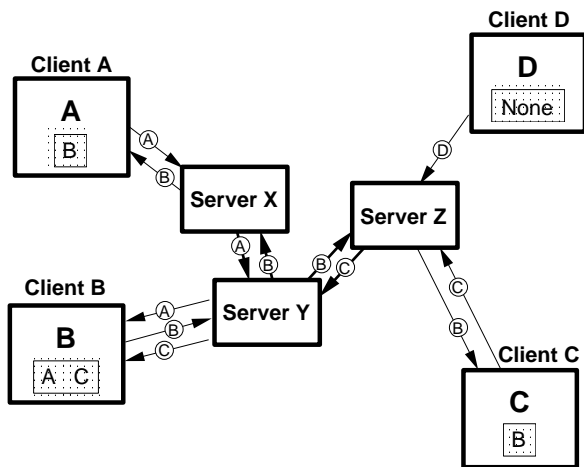

Figure 7: Flow of update messages (labeled arrows) for updates to entities A, B, C, and D arranged in a virtual environment as shown in Figure 4.

RING servers allow each client workstation to maintain surrogates for only the subset of remote entities visible to at least one entity local to the client. All other remote entities are irrelevant to the client so there is no need to waste storage space or behavioral simulation processing for them. To support this feature, servers send their clients an "Add" message each time a remote entity enters a cell potentially visible to one of the client's local entities for the first time. A "Remove" message is sent when the server determines that the entity has left the client's visible region. As entities move through the environment, servers augment update messages with "Add" and "Remove" messages notifying clients that remote entities have become relevant or irrelevant to the client's local entities. Since the system uses an unreliable network protocol, the "Add" and "Remove" messages are considered hints and

need not necessarily be processed by clients. However, they allow a client to store and simulate a small subset of the entities with little additional processing or message traffic.

The primary advantage of the RING system design is that the storage, processing, and network bandwidth requirements of the client workstations are not dependent on the number of entities in the entire distributed simulation. Client workstations must store, simulate, and process update messages only for the subset of entities visible to one of the client's local entities. In densely occluded virtual environments, visible sets tend to be constant size (e.g., how many rooms you can see looking into the hallway from your office usually does not depend on the size of your building or whether your building is surrounded by a large city), so the burden on individual client workstations does not grow as the entire system does.

Another advantage is that high-level management of the virtual environment may be performed by servers without the involvement of every client. For instance, adding or removing an entity to or from the virtual environment requires notification of only one server. That server handles notification of other servers and clients. Also, the client-server design allows use of efficient networks and protocols available between server workstations, but not universally available to all client workstations. For instance, clients may connect to servers via low-bandwidth networks, while servers communicate with each other via high-bandwidth networks.

The storage and processing requirements of RING servers are within practical limits. Unlike clients, servers do not have to store display data (e.g., polygons, textures, etc.). But, they must maintain spatial subdivision and visibility information for the virtual environment (typically < 20MB for large environments) and a surrogate representation for every entity in the environment (currently 48 bytes per entity). As server storage requirements grow linearly with the total number of entities, the size of server workstation memory may theoretically limit the number of entities that are able to share a virtual environment simultaneously. However, this is not likely to be a problem in practice since a workstation with 64MB of memory can accommodate nearly one million entities.

Server workstation processing is also within reasonable bounds. Servers must process messages in real-time only for entities visible to some entity managed by one of their clients; they are not required to simulate entity behavior between updates; and, they do not render images of the virtual environment. As a result, the memory capacity and processing power of standard UNIX workstations are adequate for RING servers in densely occluded virtual environments with very large numbers of simultaneous users.

The disadvantage of the RING system design is that extra latency is introduced when messages are routed through servers. Rather than sending messages directly between clients, RING routes each one through at least one server, and possibly two. Computations are performed in the servers before messages are propagated further adding to latency. So far, the extra latency due to server processing has not been noticeable during experiments. Additional work will have to be done to quantify the latency costs and to determine which types of entity interactions are sensitive to latency issues.

5 Experimental Results

A prototype multi-user simulation system has been implemented with the client-server design described in the previous section. The system runs on Silicon Graphics workstations and uses UDP/IP datagrams for message passing. This section presents results of experiments with this system managing many entities interacting in large densely occluded virtual environments. The virtual environments used in these experiments were mazes of "rooms" connected by "hallways." They were constructed by instancing a simple floor-plan 1, 2, 4, 8, 16, and 32 times in a square tiling pattern. Each tile contained 25 rooms (counting hallways) and had 724 polygons (see Figure 8). The largest environment used in these tests had 23,168 polygons which formed 2,219 cells. The spatial subdivision and visibility information for this environment took 99 seconds to compute and required 11.2MB of storage.


Figure 8: One tile of virtual environment used in tests.

Experiments were run with several environment sizes and various numbers of entities, clients, and servers to characterize the scalability of the system design. During these experiments, entities navigated through the virtual environment "randomly" following piecewise linear paths in randomized directions for randomized distances. Clients sent update messages only for changes in derivatives of entity position and/or orientation (i.e., dead-reckoning) while other clients simulated intermediate positions with linear "smooth-back." Update messages containing 40 bytes (message-type[4], entity-ID[4], target-position[12], target-orientation[12], positional-velocity[4], and rotational-velocity[4]) were generated for each entity once every 2.3 seconds on average with this "random" navigational behavior.

To investigate the message processing requirements of a single client in RING, we performed tests measuring the rates of messages received by clients managing one entity navigating through virtual environments containing 64, 128, 256, 512, and 1024 entities managed by other clients. Each test was repeated in virtual environments containing 25, 50, 100, 200, 400, and 800 rooms. Plates I and II contain images captured during tests with 512 entities in a 400 room environment. Table 1 and Figure 9 show average rates of messages received by individual clients in each test. In Figure 9, points representing the same number of total entities are connected by lines, while points representing the same density of entities are at the same horizontal position in the plot.

Entities Per Room	# Entities	# Rooms	Client↔Server	
			Output	Input
10.24	1024	100	0.44	61.37
10.24	512	50	0.43	70.43
10.24	256	25	0.47	53.68
5.12	1024	200	0.55	55.93
5.12	512	100	0.45	37.37
5.12	256	50	0.44	33.20
5.12	128	25	0.46	27.26
2.56	1024	400	0.50	24.56
2.56	512	200	0.47	19.88
2.56	256	100	0.46	23.19
2.56	128	50	0.41	17.42
2.56	64	25	0.46	13.65
1.28	1024	800	0.50	11.35
1.28	512	400	0.46	14.18
1.28	256	200	0.43	13.28
1.28	128	100	0.45	12.08
1.28	64	50	0.43	8.39
0.64	512	800	0.40	4.62
0.64	256	400	0.45	6.57
0.64	128	200	0.50	6.41
0.64	64	100	0.46	5.37
0.32	256	800	0.35	3.18
0.32	128	400	0.38	3.20
0.32	64	200	0.33	3.35
0.16	128	800	0.38	1.91
0.16	64	400	0.40	1.68
0.08	64	800	0.32	0.52

Table 1: Average message processing rates (messages per second) measured in a single client (managing one entity) during experiments with 64, 128, 256, 512, and 1024 entities in virtual environments with 25, 50, 100, 200, 400, and 800 "rooms."
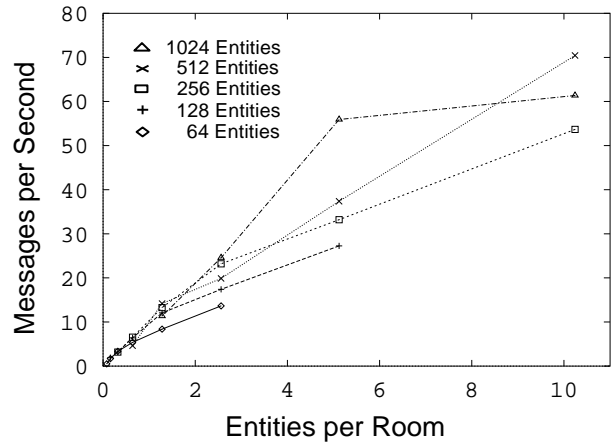

Figure 9: Average rate of messages sent to a single client (managing one entity) during tests with 64, 128, 256, 512, and 1024 entities interacting in virtual environments with 25, 50, 100, 200, 400, and 800 "rooms." Horizontal axis represents the density of entities in the environment.

From the grouping of points in the plot, we see that the rate of messages received by a single client is dependent more on the density of entities in the virtual environment than the total number of entities. This is because each client has a relatively constant sized region of interest (its visible region) which is independent of the total size of the environment or the total number of entities inhabiting it. During the test with 1024 entities simultaneously navigating through an 800 room environment each client processed only 11.35 messages (4360 bits) per second on average (row 13 of Table 1). This was approximately 2.5% of the 450 messages per second that would have been processed by each client in a system without visibility-based culling – a 40x decrease.

To characterize the message processing requirements of a single server in RING, we performed tests with various numbers of servers managing communication for 16 clients and 256 entities distributed evenly across the clients and servers in a virtual environment with 800 rooms. Table 2 lists average server→client, server↔server, and total message rates for a single server during tests with 1, 2, 4, 8, and 16 servers. Figure 10 shows a plot of total message rates per server measured during each of these tests.

	Server↔Client		Server↔Server		Total	
#	Input	Output	Input	Output	Input	Output
1	107.1	634.1	0.0	0.0	107.1	634.1
2	51.8	296.6	46.9	48.5	98.6	345.1
4	26.3	161.1	69.3	70.6	95.6	231.7
8	13.2	78.7	76.5	78.8	89.7	157.5
16	6.7	39.4	78.9	81.5	85.6	120.9

Table 2: Average message processing rates (messages per second) measured in a single server during tests with 1, 2, 4, 8, and 16 servers managing communication for 16 clients and 256 entities distributed evenly across the clients and servers in a virtual environment with 800 rooms.
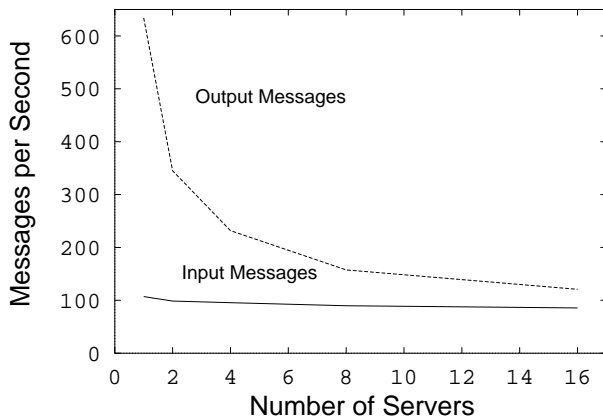

Figure 10: Average rates of messages sent to (input) and from (output) a single server during tests with 1, 2, 4, 8, and 16 servers managing communication for 16 clients and 256 entities distributed evenly across the clients and servers in a virtual environment with 800 rooms.

As the number of servers increases, the total number of messages input and output by a single server decreases. This phenomenon is aided by the fact that update messages are propagated only to servers attached to clients with entities that can potentially see the updated entity. In the test with 16 servers, 74% of the real-time server↔server messages are culled due to visibility (i.e., the 16 entities managed by each of the 16 servers cumulatively see 26% of the environment). These results are encouraging since visibility-based message culling becomes more effective as the number of servers increases and less of the model becomes relevant to each server.

From these results, we conclude that it is possible to build large multi-user visual simulation systems using a client-server design. We have found that server-based message processing algorithms which cull messages based on the three dimensional geometry of the virtual environment can be effective at reducing the network traffic into client workstations. As a result, for sufficiently occluded virtual environments, it is possible to build large, affordable multi-user virtual environments using inexpensive client workstations with low-bandwidth network connections, while higher performance workstations are required only for the relatively few servers.

6 Discussion

Several alternate approaches and future extensions are possible for this system.

Multicast

In our first experiments with multi-user virtual environments, we used IP multicast to send update messages directly between clients. The general idea is to map entity properties into multicast groups, and send update messages only to relevant groups [6]. For instance, Macedonia [12] partitions a virtual world into a 2D grid of hexagonal shaped cells each of which is represented by a separate multicast group. Entities localize their visual interactions by sending updates only to the multicast group representing the cell in which they reside, and they listen only to multicast groups representing cells within some radius.

The multicast approach is similar to the RING client-server approach for wide-area networks. In both cases, intermediate machines may cull messages rather than propagating them to all participating workstations. However, using multicast, message culling is done by routers at the network layer, whereas, in RING, message culling is done by server machines at the application layer (see Figure 11). The advantages of the multicast approach are that: 1) fewer messages must be passed if clients are connected directly to a multicast-capable LAN (e.g., ethernet), and 2) latency is reduced due to faster message routing. The disadvantages are that: 1) delays associated with joining and leaving multicast groups make it impractical to use highly dynamic entity properties for multicast group mappings, 2) the number of unique multicast groups accessible to any one application may not be sufficient for complex virtual environments, and 3) multicast is not generally available across wide-area networks to many types of networked computers (e.g., PCs with modems).

The advantage of the RING client-server approach is that very dynamic and complex *message processing* may be per-
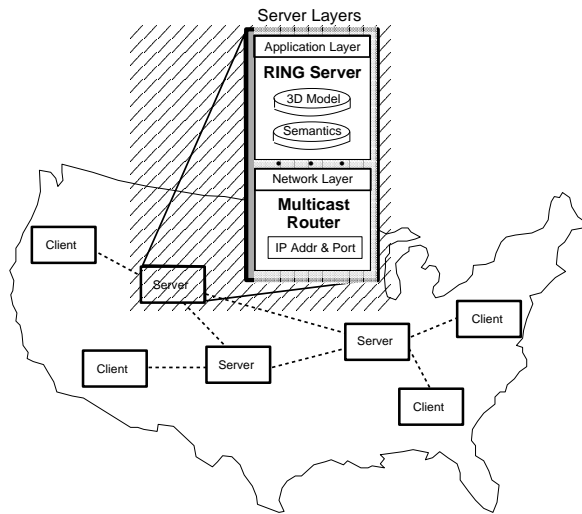
Figure 11: RING servers process messages in the application layer using 3D model and semantic information. Multicast routers use only IP addressing in the network layer.

formed by servers. In contrast to multicast routers, which can only cull messages based on a relatively small, static set of multicast groups, RING servers can cull messages using high-level geometric algorithms and knowledge regarding a multiplicity of highly dynamic entity attributes (e.g., location, orientation, velocity, etc.) and interaction types (e.g., visibility, sound, collision, etc.). Since RING servers can take advantage of knowledge regarding message semantics and the 3D geometry of the virtual environment directly, they can execute more effective and flexible culling algorithms than would be possible using only IP address and port mappings. Furthermore, unlike multicast routers, RING servers may process, augment, and alter messages in addition to culling them. For instance, RING servers already augment update messages with "Add" and "Remove" messages to inform clients that entities are entering or leaving their potentially visible sets.

Server Topology

We have experimented with a variety of topologies for connecting RING clients and servers. For practical reasons, we have focused mainly on arrangements in which clients communicate with a single server. However, depending on the capabilities of available workstations and networks, clients can send messages to server(s) via unicast or multicast. Clients can choose server(s) to manage their messages statically (i.e., all of a client's messages are sent to the same server(s)) or dynamically (e.g., based on the position of the updated entity). Servers have similar choices for distribution of messages among themselves, but can also be arranged in a hierarchy in which some servers manage communication between others.

Perhaps the most promising topologies are those in which servers manage communication between entities in separate regions of the virtual environment. For instance, we have implemented protocols with which entities migrate to a server managing the region of the environment containing the centroid of its enclosing cell. The advantage of this approach is that server-server communication is greatly reduced if there

is relatively little inter-visibility between regions. In such cases, most real-time updates affect only entities managed by the same server, and periodic updates must be passed only to servers whose region is visible to the updated entity. In early experiments, more than 95% of server-server messages are eliminated with regional servers. Further work is required to fully investigate the trade-offs between regional and other types of client-server topologies.

Multiresolution Simulation

An extension to RING currently being investigated is to use multiresolution simulation to reduce network traffic and client behavioral simulation processing. One idea is to allow RING servers to process sequences of messages and elide updates based on the perceptible importance to each client's entities. For example, consider the situation shown in Figure 12. Although A can see both B and E, B is closer to A. Thus, updates to B may be more important to A than updates to E, and could be sent to A at a finer resolution. In fact, E may be far enough away that small updates are imperceptible to A, so they can be elided completely. More generally, RING servers can alter any sequence of update messages for any entity dynamically to meet the perceptible quality required by each client. Finally, time critical computing algorithms can be used to determine an "optimal" set of messages to send to each client based on network connection bandwidths, workstation processing capabilities, and many other real-time performance factors (i.e., in a manner similar to that used in [8]).


Figure 12: RING Servers may propagate sequences of update messages to client A at finer resolutions for entity B, which is nearby, than for entity E, which is far away.

Multiresolution simulation and time critical computing algorithms can also be useful for behavioral simulation in RING clients. Every client simulates behavior for every potentially visible remote entity between updates. If surrogate behaviors can be described at multiple resolutions, simpler behavioral models can be used for entities that are perceptibly less important. For instance, in a flight simulator, very detailed behavioral models might be used to simulate an airplane just off the wing of a local entity, whereas coarse resolution behavioral models can be used to simulate an airplane that is far away and just barely visible on the horizon. By allowing clients to choose a behavioral model to simulate for each remote entity dynamically based on its perceptible importance to local entities, we can further reduce the processing requirements of client workstations.

Interaction Types

Although RING servers currently support only visual interactions, we expect that other types of interactions (sound, collision, etc.) can benefit from similar server-based message processing techniques. We are working on extensions to RING to support more general types of interactions and environments.

7 Conclusion

RING is a system for managing communication between multiple users interacting in a shared three dimensional virtual environment. It uses a client-server design along with visibility-based message culling algorithms to greatly reduce the message traffic required to maintain consistent state during multi-user visual simulations. Each client workstation must store in memory, process update messages, and simulate behavior for only a small subset of the entities participating in the entire distributed simulation – i.e., the ones visible to its entities. Inexpensive workstations with little storage capacity, slower cpus, and low bandwidth network connections may be used for clients, while high performance workstations and high bandwidth networks are required only for the relatively few servers and their interconnections. As a result, this client-server system design scales affordably to very large numbers of users interacting in densely occluded virtual environments.

Acknowledgements

References

[1] Airey, John M., John H. Rohlf, and Frederick P. Brooks, Jr., Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments. *ACM SIGGRAPH Special Issue on 1990 Symposium on Interactive 3D Graphics*, 24, 2 (1990), 41-50.

[2] Blanchard, C., S. Gurgess, Y. Harvill, J. Lanier, A. Lasko, M. Oberman, and M. Teitel, Reality Built for Two: A Virtual Reality Tool. *ACM SIGGRAPH Special Issue on 1990 Symposium on Interactive 3D Graphics*, (Snowbird, Utah), 1990, 35-36.

[3] Blau, Brian, Charles E. Hughes, Michael J. Moshell, and Curtis Lisle, Networked Virtual Environments. *ACM SIGGRAPH Special Issue on 1992 Symposium on Interactive 3D Graphics*, (Cambridge, MA), 1992, 157-164.

[4] Bricken, William, and Geoffrey Coco *The VEOS Project.* Technical Report, Human Interface Technology Laboratory, University of Washington, 1993.

[5] Calvin, James, Alan Dickens, Bob Gaines, Paul Metzger, Dale Miller, and Dan Owen, The SIMNET Virtual World Architecture. *Proceedings of the IEEE Virtual Reality Annual International Symposium*, September, 1993, 450-455.

[6] Carlsson, Christer, and Olof Hafsand, Dive: A Multi-User Virtual Reality System. *Proceedings of the IEEE Virtual Reality Annual International Symposium*, September, 1993, 394-401.

[7] Funkhouser, Thomas A., Carlo H. Séquin, and Seth J. Teller, Management of Large Amounts of Data in Interactive Building Walkthroughs. *ACM SIGGRAPH Special Issue on 1992 Symposium on Interactive 3D Graphics*, (Cambridge, MA), 1992, 11-20.

[8] Funkhouser, Thomas A., and Carlo H. Séquin. Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments. *Computer Graphics (SIGGRAPH '93)*, 27, 247-254..

[9] Institute of Electrical and Electronics Engineers (IEEE), IEEE P1278 - Standard for Information Technology – Distributed Simulation Application – Process and Data Entity Interchange Formats.

[10] Kazman, Rick, Making WAVES: On the Design of Architectures for Low-end Distributed Virtual Environments. *Proceedings of IEEE Virtual Reality Annual International Symposium*, September 1993, 443-449.

[11] Kazman, Rick, Load Balancing, Latency Management and Separation of Concerns in a Distributed Virtual World. *Parallel Computations - Paradigms and Applications*, A. Zomaya (ed.), Chapman & Hall, 1995, to appear.

[12] Macedonia, Michael, R. Michael J. Zyda, David R. Pratt, and Paul T Barham, Exploiting Reality with Multicast Groups: A Network Architecture for Large Scale Virtual Environments. To appear in *Proceedings of IEEE Virtual Reality Annual International Symposium*, 1995.

[13] Pope, Arthur R., The SIMNET Network and Protocols. *Technical Report 9120*, LORAL Advanced Distributed Simulation, Cambridge, MA, June, 1991.

[14] Shaw, Chris, and Mark Green, The MR Toolkit Peers Package and Experiment. *Proceedings of IEEE Virtual Reality Annual International Symposium*, September 1993, 463-469.

[15] Teller, Seth J., and Carlo H. Séquin, Visibility Preprocessing for Interactive Walkthroughs. *Computer Graphics (SIGGRAPH '91).* 25, 4, 61-69.

[16] Teller, Seth J., *Visibility Computations in Densely Occluded Polyhedral Environments.* Ph.D. thesis, Computer Science Division (EECS), University of California, Berkeley, 1992. Also available as UC Berkeley technical report UCB/CSD-92-708.

[17] Zyda, Michael J., David R. Pratt, John S. Falby, Chuck Lombardo, and Kristen M. Kelleher, The Software Required for the Computer Generation of Virtual Environments. *Presence*, 2, 2 (March 1993), 130-140.

Plate I: Top-down view of a multi-user virtual environment with 400 rooms and 512 simultaneously moving entities (represented by yellow spheres with green orientation vectors).


Plate II: Image of same multi-user virtual environment shown in Plate I, rendered from the viewpoint of one entity. Client managing this entity must process updates, simulate behavior, and store surrogates for only 14 remote entities (2.7% of all entities) due to message processing in RING servers.