

Characterization of Distributed Virtual Environments Based on Peer-To-Peer Architectures*

Silvia Rueda, Pedro Morillo, Juan Manuel Orduña

Dept. Informática

Universidad de Valencia

{Silvia.Rueda,Pedro.Morillo,Juan.Orduña}@uv.es

José Duato

DISCA

Universidad Politécnica de Valencia

jduato@gap.upv.es

Resumen

Large scale distributed virtual environments (DVEs) have become a major trend in distributed applications, mainly due to the enormous popularity of multi-player online games in the entertainment industry. Since architectures based on networked servers seem to be not scalable enough to support massively multi-player applications, peer-to-peer (P2P) architectures have been proposed as an efficient and truly scalable solution for this kind of systems. However, in order to design efficient DVEs based on peer-to-peer architectures these systems must be characterized, measuring the impact of different client behaviors on system performance.

This paper presents the experimental characterization of peer-to-peer distributed virtual environments in regard to well-known performance metrics in distributed systems. Characterization results show that system saturation is inherently avoided due to the peer-to-peer scheme, as it could be expected. Also, these results show that the saturation of a given client exclusively has an effect on the surrounding clients in the virtual world, having no noticeable effect at all on the rest of avatars. Finally, the characterization results show that the response time offered to client computers greatly depends on the number of new connections that these clients have to make when new neighbors appear in the virtual world. These

results can be used as the basis for an efficient design of peer-to-peer DVE systems.

1. Introduction

In recent years, large scale distributed virtual environments (DVEs) have become a major trend in distributed applications, mainly due to the enormous popularity of multi-player online games in the entertainment industry. These highly interactive systems simulate a 3-D virtual world where multiple users share the same scenario. Each user is represented in the shared virtual environment by an entity called *avatar*, whose state is controlled by the user through a client computer. The system renders the images of the virtual world that each user would see if he was located at that point in the virtual environment. Hundreds and even thousands of client computers can be simultaneously connected to the DVE through different networks, and even through Internet. DVE systems are currently used in many different applications [13], such as civil and military distributed training, collaborative design and e-learning. Nevertheless, the most extended example of DVE systems are commercial, massively multi-player online game (MMOG) environments [4, 2, 15].

Architectures based on networked servers have been the major standard for DVE systems during the last years [13, 3]. In these architectures, the control of the simulation relies on several interconnected servers. Client computers are assigned to one of the servers in the system. When a client computer mo-

*This work has been jointly supported by the Spanish MEC and European Commission FEDER funds under grants Consolider Ingenio-2010 CSD2006-00046 and TIN2006-15516-C04-04.

defines the state (usually the position, but it can also modify the appearance or other stateful information) of an avatar, it also sends an updating message to its server, which in turn must propagate this message to other servers and clients. In this sense, concepts like areas of influence (AOI) [13] have been proposed for limiting the number of neighboring avatars that a given avatar must communicate with. All these concepts define a neighborhood area for avatars, in such a way that a given client computer controlling a given avatar i must notify all the movements of i (by sending an updating message) only to the client computers that control the avatars located in the neighborhood of avatar i . The avatars in the AOI of avatar i are denoted as *neighbor avatars* of avatar i . However, in spite of these techniques networked-server architectures do not properly scale with the number of existing users, particularly for the case of MMOGs (where up to hundreds of thousands of clients can be simultaneously connected to the system) [1].

Peer-to-peer architectures (P2P) seem to be the most adequate scheme in order to provide good scalability for large scale DVE systems, and several online games based on P2P architectures have been designed [10]. However, P2P architectures must face the awareness problem, consisting of ensuring that each avatar (for the sake of shortness, in the rest of the paper we will use the term avatar to denote the client computer controlling that avatar) is aware of all the avatars in its neighborhood [14]. Solving the awareness problem is a necessary condition to provide each participant with a consistent view of the environment. In this sense, different strategies for providing awareness in DVE systems based on P2P architectures have been proposed [7, 11]. All these strategies impose both additional communications between different avatars and additional computations (performed by the client computers controlling these avatars). In order to design truly efficient P2P DVE systems, the impact that these (and other) computations and communications have on the real system performance must be measured.

Based on that motivation, in this paper we propose the characterization of peer-to-peer DVE systems. Since the behavior of an avatar can be determined by both the number of neighbor avatars in its surroundings and its movement rate [12], we have measured the effects that such parameters have on two well-known performance metrics for distributed systems: latency and throughput. However, in order to avoid clock skewings, we have measured the response time (defined as the round-trip delay for the messages sent by each client) instead of latency. Although appearance and other kind of informations can also be exchanged between avatars in a DVE, the effects that adding these exchanges have on the system can also be achieved by increasing the movement rate and/or the number of neighbors, since at last they result either in more messages exchanged among more destinations or in more messages exchanged during the same period.

The results show that, unlike in networked-server architectures, the system throughput increases with the number of client computers connected to the system. That is, the peer-to-peer scheme is fully scalable, as it could be expected from a fully distributed scheme. Also, the results show that the saturation of a given client exclusively has an effect on the surrounding clients in the virtual world, having no effects at all on the rest of avatars. Finally, characterization results show that the response time offered to client computers greatly depends on the number of new connections that these clients have to establish when new neighbors appears in their surroundings. These results can be used as the basis for an efficient design of peer-to-peer DVE systems.

The rest of the paper is organized as follows: Section 2 details the proposed characterization setup that allows to experimentally study the behavior of peer-to-peer DVE systems. Next, Section 3 presents the evaluation results. Finally, Section 4 presents some concluding remarks and future work to be done.

2. Characterization Setup

We propose the evaluation of P2P DVE systems by simulation, as it was done for DVE systems based on networked-server architectures [12]. The evaluation methodology used is based on the main standards for modeling collaborative virtual environments, such as FI-PA [5], DIS [8] and HLA [9]. Since DVE systems are inherently based on networks, the metrics used for evaluating the performance of these systems include the two main metrics used for evaluating network performance. These metrics are latency and throughput. Nevertheless, in order to avoid clock skewing we have selected throughput and response time (defined as the round-trip delay for the messages sent by each client) as the performance metrics to be characterized. Concretely, we have developed a simulator modeling a DVE system based on a peer-to-peer architecture. The simulator is written in C++ and it is composed of two applications, one modeling the clients and the other one modeling the central loader. All clients must initially join the system through the central loader. Both applications use different threads for managing the different connections they must establish. Such connections are performed by means of sockets.

Each client has a main thread for managing the actions required by the user and different threads for communicating with its neighbor clients. For each neighbor, two threads are executed, one for listening and one for sending messages. Similarly, the central loader has two threads for communicating with each client in the system and also a main thread. It must be noticed that once a client has joined the system, that client does not need to communicate with the central loader any more. Since the goal of this characterization is to study how the system evolves as clients interact with the environment, rather than analyzing how new clients join the system, in our simulator each client is initially provided with the IP addresses of its initial neighbors.

A simulation consists of each avatar performing 100 movements. An iteration of the who-

le system consists of all avatars making a movement. Each avatar notifies its neighbors as well as the central loader when it reaches the 101th iteration, and then it leaves the system. We have chosen the number of 100 iterations (movements) for a simulation because it is the number of movements that the most distant avatar needs to reach the center of the square virtual world. The virtual world was a 2D square whose sides were 100 meters long. Each time an avatar moves, it sends a message to all its neighbor avatars (the client computer controlling that avatar sends a message to the client computers controlling the neighbor avatars). These destination avatars then send back an acknowledgment to the sending avatar, in such a way that the sending avatar can compute the round-trip delay for each message sent. We have denoted the average round-trip delay for all the messages sent by an avatar as the Average System Response (ASR) for that avatar (for that client computer).

Two different characterization setups have been made, depending on the metrics to be studied. In order to study the system throughput (the maximum number of clients the system can support while maintaining full awareness and reasonable latency levels), we have used a cluster of 14 nodes. One of these PCs hosted the central loader, and the rest of the 13 PCs hosted the clients in the system in a uniformly distributed way. Each node was a dual Opteron processor running Linux (SuSE 10.1 distribution). Although a real system does not require communication between clients and the central loader, we have implemented a monitoring algorithm to check that the system supports a full awareness rate. This algorithm consists of each client dividing its cycle time in two phases. In the first phase, clients move following a given movement pattern (described in Section 3) and they communicate each other their new location in the virtual space by exchanging messages. In the second phase, each client sends the central loader a message containing information about its new location and also about which other clients it considers as its neighbors. In this way, the central loader can compute the awareness

rate in real time. We have used a movement rate of one client movement each 3.6 seconds. From this period, 2.1 seconds are dedicated to the first phase and 1.5 seconds are dedicated to the second phase. Finally, we used an AOI size of 10 meters.

In order to study the system latency, we have used a different characterization setup. In this case, we have used personal computers interconnected by a fast Ethernet network, hosting the avatars in a uniformly distributed way. The purpose of such setup is to establish an upper number for the average value of the ASR when client computers are slightly loaded. Finally, in order to study the effects that the saturation of a single client computer has on the rest of clients, we have used the same setup except that one of the PCs hosts a single client, and the rest of the PCs host the rest of the clients in a uniformly distributed way. In these two setups we have not monitored the awareness rate, and the movement cycle time has been reduced accordingly.

3. Evaluation Results

We have simulated the behavior of a set of independent avatars in a generic DVE system based on a P2P architecture. These avatars are located within a seamless 3D virtual world following three different and well-known initial distributions: uniform, skewed and clustered [12]. Starting from these initial locations, in each simulation avatars can move into the scene following one of three different movement patterns: Changing Circular Pattern (CCP), HP-All (HPA) and HP-Near (HPN) [12]. CCP considers that all avatars in the virtual world move randomly around the virtual scene following circular trajectories. HPA considers that there exists certain “hot points” where all avatars approach sooner or later. This movement pattern is typical of multiuser games, where users must get resources (as weapons, energy, vehicles, bonus points, etc.) that are located at certain locations in the virtual world. Finally, HPN also considers these hot-points, but only avatars located within a given radius of the hot-points approach these locations. In




Figure 1: Final distribution of avatars for a) CCP, b) HPN, and c) HPA movement patterns applied to an initial uniform distribution of avatars.

order to illustrate these movement patterns, Figure 1 shows the final distribution of avatars that a 2-D virtual world (represented as a square map) would show if these movement patterns were applied to a uniform initial distribution of avatars. For evaluation purposes, we have considered the nine possible combinations of the three initial distributions of avatars in the virtual world and the three movement patterns.

3.1. Scalability

The first feature to be proved in a peer-to-peer scheme is its scalability. With this purpose, we have performed more than 4000 experiments to study the behavior of DVE systems based on a P2P architecture. In all of them, the system has provided a full awareness rate while saturation has not been reached by any CPU. We have made tests with different world sizes S (the number of connected clients), although for the sake of shortness we will show a single case with 520 avatars.

Concretely, Figure 2 shows a representative example of the experiments performed in order to study the system throughput. On the X-axis this Figure shows the iteration number of the simulation performed, and on the Y-axis it shows the average value in seconds for the ASR of all the avatars. Each point in the plots represents the average ASR of ten different executions of the same simulation. In all of them, each cluster node hosted 40 clients.

Figure 2 shows that the ASR remains always below the threshold of 250 ms. required to provide QoS to users [6] and also with a full awareness rate. Since in these experiments each




Figure 2: Average ASR values for different configurations of a DVE with 520 avatars

cluster node hosted 40 avatars, we can state that there would not be any problem (the CPU utilization would be far from 100 %) when executing one avatar in one PC.

Additionally, we have analyzed the system behavior in a transverse way, that is, we have studied the average ASR values for different population sizes. Although we have performed this analysis for all the combinations of initial distributions and movement patterns, for the sake of shortness we show here a single case. All the cases showed similar results. Concretely, Figure 3 shows the average ASR values provided to avatars when following the Uniform-CCP scheme.




Figure 3: Average ASR values for different populations of avatars under a Uniform-CCP combination

Figure 3 shows three flat lines at a different levels, showing that while the population in-

creases in an order of magnitude (from 100 to 1000 avatars) the average ASR increases by a factor of two, and it is still far from the QoS threshold of 250 ms.. Therefore, these experiments prove that the peer-to-peer scheme is fully scalable, regardless of the number of client computers connected to the system. Although these results could be expected due to the inherently distributed nature of the peer-to-peer scheme, it is worth mention the remarkable differences with the results provided by the networked-server scheme [12].

3.2. Levels of Interaction

Another aspect of P2P DVEs to be characterized is the behavior of the system under different levels of interactions, that is, under different communication rates and different numbers of neighbors to communicate with. Concretely, we have measured the average value of the ASR provided to all avatars for different movement patterns and initial distributions of avatar in the virtual worlds. Since the workload that a given avatar adds to the system depends on both the movement rate of the avatar and also on the number of neighbor avatars in the virtual world [12], we have measured the ASR provided by the system for different values of these two parameters. In these tests we used a wide range of populations sizes. However, for the sake of shortness we will show in this Section the results for a system with 101 avatars. Concretely, Figure 4 shows the results for all of the distributions with a cycle period T of 0.6 seconds (all avatars make a movement every 0.1 seconds) and an AOI of 10 meters, while Figure 4 shows the results for the same AOI but now with a movement rate of 0.1 seconds. Then, Figure 5 show the results for an AOI of 20 meters and a movement rate of 0.1 seconds. Both Figures show on the X-axis the iteration number (a simulation is composed of 100 iterations), and it shows on the Y-axis the average value of the ASR provided to all the avatars.

In Figure 4 it can be seen that the three plots corresponding to the HPA movement pattern and the plot for the HPN pattern with a skewed initial distribution of avatars show a

linear increasing with the number of iterations. That is, for these patterns the client computers are not capable of sending, returning and processing the number of messages generated by avatars when they move in a period cycle. This is due to the fact that in the HPN and particularly in the HPA pattern, avatars tend to crowd the hot points, increasing the number of neighbors in the AOI in the subsequent iterations. Since they must send more messages in the same cycle period, the client computers become saturated and the average ASR increases.




Figure 4: Average ASR values for $T=0.1$ s. & $AOI=10$ m.

Figure 5 shows the results for the same population of avatars and the same patterns when the AOI size is doubled. The results shown in this Figure are similar to those shown in Figure 4.

Figure 5 shows that for five combinations of movement patterns and initial distributions of avatars the average ASR values linearly increase with the iteration number. This behavior indicates that the client computer cannot process all the messages sent and received at that rate. These results show that increasing the AOI size has less significant effects on system latency than the movement rate of avatars. When comparing this Figure with Figure 4 the only significant difference is that in the former one there are five plots with a non-zero slope, while in the latter one there are four plots.

Additionally, Figure 5 shows an unexpected




Figure 5: Average ASR values for $T=0.1$ s. & $AOI=20$ m.

behavior. Effectively, it can be seen that the plot with the highest ASR values is the CCP movement pattern with a skewed initial distribution of avatars in the virtual world (SKW CCP). This is a strange behavior, because the CCP pattern describes a circular pattern and therefore the number of avatars within the AOIs of avatars does not increase as simulation proceeds.

In order to find the reasons for this strange behavior, we have studied the average percentage of messages sent by any avatar in each iteration, in respect with the total number of avatars in the system S . Although the results are not shown here due to space limitations, the CCP pattern provided one of the lowest percentage of messages, therefore discarding the number of messages as the cause of the strange behavior. Also, we studied the number of neighbors, but the percentages of neighbors in each iteration for the SKW CCP combination neither were the higher ones.

Finally, we studied the number of new connections started by avatars in each iteration. Figure 6 shows the average percentage of connections started by avatars (in an accumulative measurement), with respect to the total number of avatars in the system S .

Figure 6 shows that the SKW CCP plot is one of the combinations with the highest number of new connections per iteration, and therefore these results can explain the behavior shown in Figure 5. Effectively, the CCP

is a circular movement pattern. When this pattern is applied on a skewed distribution of avatars, in each iteration there are a lot of avatars that enter and leave the AOI of each avatar. This event occurs for a lot of avatars, forcing the corresponding client computers to establish new connections and also to finish some of the current ones. The overhead required to perform these actions makes the client computers to enter saturation, the connections cannot be established immediately and therefore the average ASR increases.




Figure 6: Average fraction of new connections for AOI=20m and T=0.1

As a result, we can state that not only the number of current neighbors, but also the number of new connections made along time by avatars affects the ASR provided to other avatars, particularly if the client computers are close to saturation.

3.3. Client Saturation

Additionally, we have studied the effects that a slow client computer or network link(s) can have on the rest of client computers. For this study, we have tested the 9 different combinations of initial distributions of avatars and movement patterns on a real DVE system based on P2P with 101 PCs. In these experiments, we have forced a given avatar (controlled by a single client computer) to move faster and faster, until the number of movements per second is so high that the client computer cannot send all the messages to the corresponding neighbors within the cycle period. Under these

conditions, we have measured the latency (the ASR) provided to different avatars. We have used the consistency algorithm used in [11]. This algorithm distinguishes among different kinds of neighbors of a given avatar i . First, the first level neighbors (L1) are those avatars located inside the AOI of avatar i . Next, the second level neighbors (L2) are those avatars that are L1 neighbors of the L1 neighbors of i and are not inside AOI of avatar i . Therefore, we have computed the average ASR value provided to i , the average ASR value provided to its L1 neighbors, the average ASR value provided to its L2 neighbors and the average ASR value provided to rest of the avatars in the system.

Due to space limitations, we will show the results only for a single combination of initial distributions and movement patterns of avatars. However, all the combinations produced very similar results. Figure 7 shows the results for the combination of a uniform distribution of avatars and an HPAll movement pattern. While the plot for AV0 has a parabolic shape, the shape of the plot for L1 neighbors is linear with a low slope. The plot for the L2 neighbors is similar to the plot for L1 neighbors, except that from the 70th iteration up it becomes negligible. The plot for the rest of avatars is a flat line with a negligible value. These plots indicate that the effects of the saturation of AV0 (the client computer controlling AV0) are one order of magnitude lower in its L1 and L2 neighbors, while they are negligible in the rest of avatars. Therefore, with all these results we can state that the effects of the saturation of a given avatar are limited to L1 and L2 neighbors. The rest of avatars are not affected.

4. Conclusions

In this paper, we have proposed the experimental characterization of peer-to-peer distributed virtual environments, in order to set up the basis for an efficient design of this kind of DVE systems. The characterization results show that system saturation is inherently avoided due to the peer-to-peer scheme, as it could be expected. Also, characterization re-




Figura 7: Effects of a client saturation on different kinds of avatars under a Uniform-HPALL combination

sults show that the saturation of a given client exclusively has an effect on the surrounding clients in the virtual world, having no effects at all on the rest of avatars. This result can be crucial for designing an awareness technique or to achieve an acceptable time-space consistency.

Finally, characterization results show that the response time offered to client computers greatly depends on the number of new connections that these clients have to make when new neighbors appear in their areas of interest. That is, the saturation of an avatar (of the client computer controlling that avatar) can be produced by the sudden appearance (or disappearance) of several neighbors in its AOI. These results can be used as the basis for an efficient design of peer-to-peer DVE systems.

Referencias

- [1] T. Alexander. *Massively Multiplayer Game Development II*. Charles River Media, 2005.
- [2] Anarchy Online: : <http://www.anarchy-online.com>.
- [3] N. Beatrice, S. Antonio, L. Rynson, and L. Frederick. A multiserver architecture for distributed virtual walkthrough. In *Proceedings of ACM VRST'02*, pages 163–170, 2002.
- [4] Everquest: <http://everquest.station.sony.com/>.
- [5] FIPA. Fipa agent management specification, 2000. Available at <http://www.fipa.org/specs/fipa00023/>.
- [6] T. Henderson and S. Bhatti. Networked games: a qos-sensitive application for qos-insensitive users? In *Proceedings of the ACM SIGCOMM 2003*, pages 141–147. ACM Press / ACM SIGCOMM, 2003.
- [7] S. Y. Hu and G. M. Liao. Scalable peer-to-peer networked virtual environment. In *Proceedings ACM SIGCOMM 2004 workshops on NetGames '04*, pages 129–133, 2004.
- [8] IEEE. *1278.1 IEEE Standard for Distributed Interactive Simulation-Application Protocols (ANSI)*, 1997.
- [9] F. Kuhl, R. Weatherly, and J. Dahmann. *Creating Computer Simulation Systems: An Introduction to the High Level Architecture*. Prentice-Hall PTR, 1999.
- [10] S. Mooney and B. Games. *Battlezone: Official Strategy Guide*. BradyGame Publisher, 1998.
- [11] P. Morillo, W. Moncho, J. M. Orduña, and J. Duato. Providing full awareness to dves based on peer-to-peer architectures. *Lecture Notes on Computer Science*, 4035:336–347, 2006.
- [12] P. Morillo, J. M. Orduña, M. Fernández, and J. Duato. Improving the performance of distributed virtual environment systems. *IEEE Transactions on Parallel and Distributed Systems*, 16(7):637–649, 2005.
- [13] S. Singhal and M. Zyda. *Networked Virtual Environments*. ACM Press, 1999.
- [14] R. B. Smith, R. Hixon, and B. Horan. *Collaborative Virtual Environments*, chapter Supporting Flexible Roles in a Shared Space. Springer-Verlag, 2001.
- [15] Startcraft: <http://www.blizzard.com/starcraft>.