Monitoring and Analyzing Performance of Networked Virtual Environments: the Case of EVE

C. Bouras[1, 2] E. Giannaka[1, 2]

[1]Computer Engineering and Informatics Department, University of Patras GR-26500 Rion, Patras, GREECE

[2]Research Academic Computer Technology Institute 61 Riga Feraiou STR, GR-26221, Patras, GREECE

E-mail: bouras@cti.gr, giannaka@cti.gr

Abstract

This paper focuses on the performance monitoring of a networked virtual environment, called EVE, in order to identify the elements that degrade its network performance. In particular, based both on the architectural and communication model that EVE adopts in order to cope with the dynamic behavior of the network, we performed a series of experiments, taking into account the number of the connected users, the network load as well as the connection type of the users.

1. Introduction

Networked virtual environments, which were brought in the market's foreground the last decade, presented a great interest on performance issues. The sizes of the virtual worlds, the graphics and the variety of the provided services for the achievement of a higher sense of realism resulted in demanding applications in the sense of recourses and bandwidth. Thus, the performance has been many times downgraded and consistency as well as reliability could not be retained by the application.

This paper presents an effort for the performance monitoring and analysis of a networked virtual environment called EVE [1], which stands for Educational Virtual Environment. EVE constitutes a multi-user 3D platform, implemented for collaborative, distant learning. In this paper we will try to define the boundaries in which the platform can operate efficiently and reliably, through a series of experiments conducted under the network simulator ns-2 [2].

The remainder of this paper is organized as follows. We will begin by describing EVE's architecture as well as the communication model that the platform adopts. Section III is engaged with the simulation model that was used for the experiments on EVE. These experiments are described in section IV. Finally we conclude with the results that are raised from the experiments conducted on EVE along with the planned future work.

2. EVE's architectural and communication model

EVE's architecture is based on a client-multi server platform model. The structure of the server side for EVE is directed to the decrement of the load for the processes that actualize and transmit the updates of the Virtual Worlds. Therefore, the server side architecture comprises of two basic components, one main server, the Message Server and the Application servers. In particular, the Message Server is responsible for the manipulation of the virtual worlds that constitute the training area of the platform and supports the illusion to the participants that they share a common space by updating the view of the world every time that a shared object is modified. Regarding the Application Servers, they are responsible for providing specific functionality to the participants of the virtual world. In the current form of EVE there are two application servers available, a chat server and an audio server.

As far as it concerns the communication model that EVE adopts, it categorizes the messages exchanged in three basic categories [4]: a) the messages related with the initialization of the virtual world and the initial connection of a client to a server as well as the messages exchanged between the servers of the platform, b) the position messages that are related with the avatars' position in the virtual environment and c) the important messages, which correspond to messages that are vital for the consistency of the networked virtual environment. For simplicity reasons, we consider as important messages all messages except from the position messages.

The platform adopts three types of communication protocols, the TCP, the UDP and the H.323 protocol. In particular, for the following messages: a) the server-to-server communication, b) the initial connection of a client to the message server, where the current status of the world is transmitted, c) the messages that are vital for the consistency of the networked virtual environment, the TCP protocol has been adopted. As far as it concerns the position messages, which carry information about the avatars' position in the virtual world, the UDP communication was selected, as a possible failure in their delivery does not create important scene inconsistencies to the participants. Finally, the audio server uses H.323 as its main protocol. The audio server, which is provided by EVE, is in fact an H.323 MCU, which supports audio conferencing among the platform users.

3. Simulation Model of EVE

For the realization of the experiments and for approaching real situations the topology selected is depicted in Figure 1.

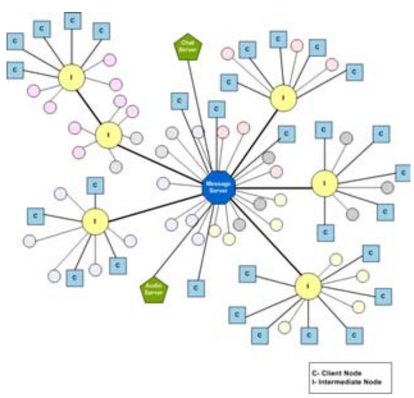

Figure 1. Topology of the experiments

In particular, the topology encounters thirty concurrent users, which are distributed in two concurrent virtual worlds. These users are scattered and could either be directly connected to the Message Server, or be situated one or two hops away from it.

Regarding the characteristics of the network, where the experiments will be conducted, we assume that the Message Server and the Application Servers are connected through a 2Mb line. For the nodes that intermediate between the client nodes and the Message Server we demarcate a line of 10 MB. To these lines we applied background traffic, which is characterized by an average throughput of 7000 Kbps for a line of 10 MB.

As mentioned, the network is a dynamic medium, which is characterized by an extremely unpredictable behavior [3]. For the conduction of the experiments we took into account the following factors that can affect the network performance: a) the number of the connected users that corresponds to the number of the simultaneous connections that the platform can support efficiently, b) the network load that refers to the amount of traffic that passes through the network and mainly affects the bandwidth, since its increment leads to decrement of the available bandwidth, c) the connection type that is related to the available bandwidth that characterizes the user's connection to the network, when s/he accesses the platform.

4. Experiments

In the framework of the paper presented, we conducted a series of experiments in order to investigate the way that the alteration variables, mentioned above, affect the platform's performance. For each of the alteration variables we measured the drop rate for the communication directed from the Message Server to the clients.

It should be mentioned that in the majority of the experiments conducted the basic topology used and usually selected as comparison metric, is the topology of the thirty simultaneous users, connected through ISDN connections. In addition the client monitored is usually situated one hop away from the Message Server and suffers from background traffic, which fills a line of 10 MB at an average of 70%.

4.1. Number of Connected Users

For this series of experiments a topology of 10 connected users was designed, which is presented in Figure 2. In this topology, we assume that all users are connected to the same virtual world and consequently the messages sent are forwarded to all of the connected clients.

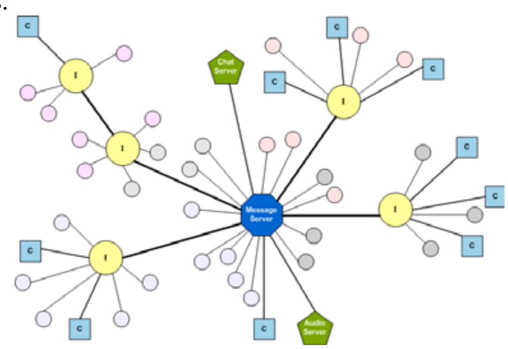

Figure 2. Topology for 10 connected clients

In Figure 3 the graphical representation of the drop rate of each topology for each of the packet types is depicted. For the drop rates of the simulated topologies we observe that for a certain period of time, at the beginning of the simulation time, the drop rate is almost equal to zero and it gradually increases until it stabilizes, at almost the middle of the simulation time.

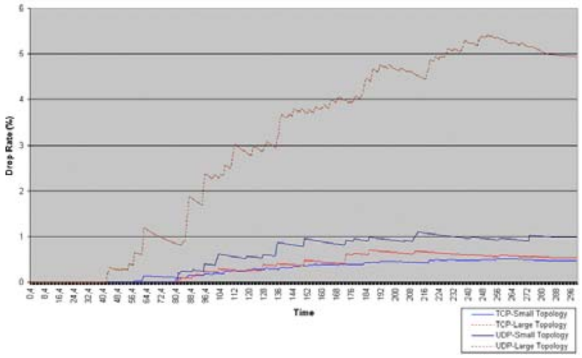

Figure 3. Comparison of the drop rate for TCP and UDP packets

However, if we observe, in Figure 4, the line that connects the Message Server to the Intermediate node of the client, we notice that in the first 40 seconds the drop rate seems higher, which in the case of the smaller topology reaches the 35% while in the case of the larger topology reaches 53%. This percentage is very high for both topologies and justifies the zero drop rate of Figure 3, since a large amount of packets was dropped in the intermediate line and the packets that managed to pass through this line could easily be supported by the available bandwidth. In addition, if we exclude the first 8 seconds in Figure 3, where the drops for both TCP and UDP packets are extremely high, for the rest of the simulation time the losses for the TCP packets are at about 2%. From the same representation we observe that the greater amount of packets lost, due to the effect of the background traffic, are UDP. Even though this percentage at the beginning reaches an amount of 30%, however as time passes this amount is settled and reaches an average of 12%.

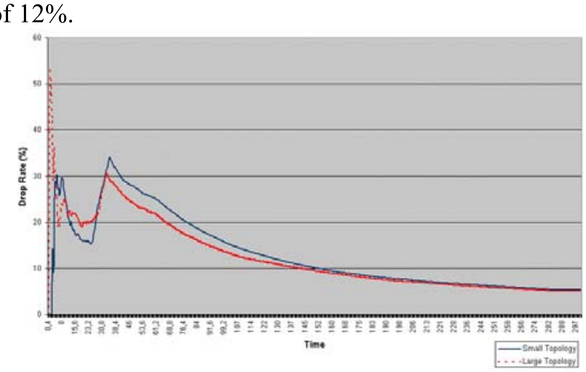

Figure 4: Drop rate of the intermediate lines due to background traffic

As mentioned previously, the UDP packets, represent the users' movement within the virtual world and the packets form a flow from the intermediate positions that the user's avatar follows. Consequently, the loss of this kind of information concerns mainly some of the intermediate positions of the users, which are re-allocated with the transmission of the new position message.

4.2. Network Load

In this series of experiments we took into account three network states: a) a network without background traffic that is considered ideal, b) a background traffic A that in a line of 10Mb achieves an average throughput of 7000 Kbps, c) a background traffic B that in a line of 10Mb, achieves an average throughput of 8000 Kbps. The traffic applied to these lines concerns mainly FTP, Telnet, Web and SMTP traffic.

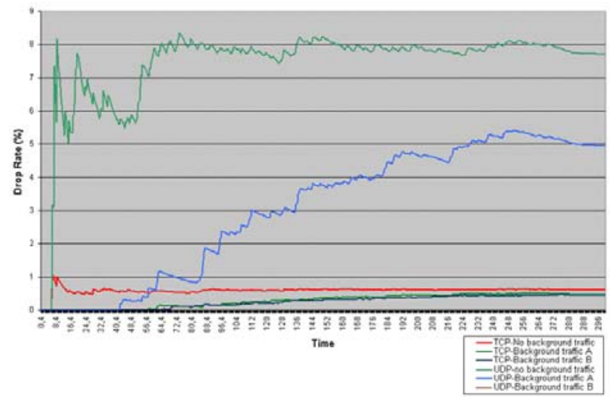

Figure 5. Drop rate of TCP and UDP packets for each case of network load

In Figure 5 we observe that the drop rate for the cases that the network is empty is greater that the one noticed when background traffic is applied. This fact can be explained as follows: in the case of an empty network all the messages designated to the client pass through the 10 MB line successfully. However, when they reach the 128 KB line congestion is created which results to losses of information. These losses mainly concern UDP packets, as displayed. Regarding the drop rate of the TCP packets, the rate is very small and minor for the maintenance of the consistency.

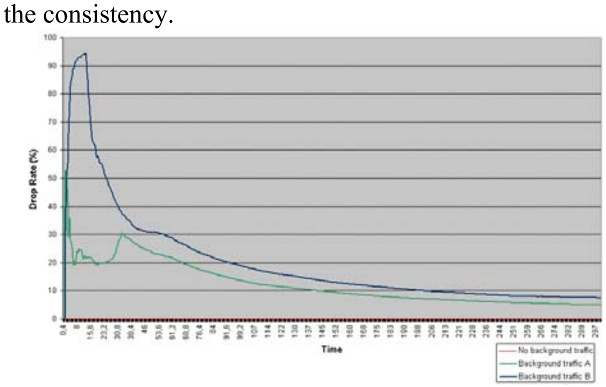

Figure 6. Drop rate of the intermediate line for each network load

As far as it concerns the losses in the case of the background traffic A and B, we notice that the drop rate is notably smaller than the rate of the empty network. This can be justified by Figure 6, where we observe that the losses that the background traffic causes reaches, in the case of background traffic B a rate of 95%. This increased rate results to the decrement of the messages that manage to pass the terminal line of the client.

4.3. Connection Type

For the realization of this series of experiments we took into account the following connection types: a) users use a modem in order to access the platform, b) users are connected through ISDN, having available a bandwidth on order of 128 KB, c) users are connected through ADSL, with a maximum bandwidth of 512 KB and d) users are provided with a line of 2 MB, which is considered optimum for the conduction of the experiments.

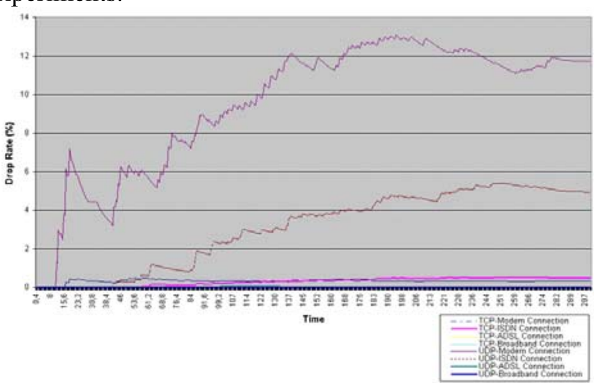

Figure 7. Drop rate for TCP and UDP packets in regard to the connection type

As depicted in Figure 7, for the modem connection, the drops can reach the rate of 13%. Regarding the ISDN connection the drop rate is notably reduced to a rate of 6%. In addition, we observe that the major part of the dropped information concerns UDP packets, which are related to the avatars movement. As far as it concerns the ADSL as well as the Broadband connection, they do not present losses in the terminal line.

5. Conclusions and Future Work

This paper was mainly targeted at the performance monitoring of a multi-user networked virtual environment, called EVE. Based on both the architectural and communication model that the platform adopts in order to handle the dynamic nature of the network we conducted of a series of experiments on ns-2.

From the experiments conducted, the following results were extracted about the networking performance of EVE: a) the current version of EVE can support efficiently up to two concurrent virtual worlds, that is a total of 34 connected users, b) the platform operates efficiently even under a serious amount of background traffic, which can reach a percent of 70%, regarding the completeness of the intermediate line, c) the performance, in regard to the connection type that the users adopt in order to access the platform is very efficient for ISDN, ADSL and Broadband connections, d) the losses noticed in the terminal line that connects the Intermediate nodes to the clients are very small and cannot cause series inconsistencies to the view of the virtual environment, e) the greatest percentage of lost packets, in all experiments conducted, corresponds to UDP packets, which are used for the position messages of the avatars. Thus, even in the case of a high loss of UDP packets, the most severe impact that the end user will meet is a rough transition of the avatar from one point to another. The above results prove that EVE can support efficiently the provided services, ensuring reliability and safety in the transmission of information.

During the primer design of the platform, we took into account the possible need for extension due to increased demand and participation. The new architectural model is still based on two types of servers, the message server and the application servers. The main concept of this architectural model is the fact that each message server is a back up server for a number of the rest message servers. Thus, if a failure happens, the clients supported and hosted by the “damaged” message server can be distributed to the other servers of the system. Thus, one of the basic next steps is to implement the extended model of EVE and perform additional simulations so as to evaluate the extended architectural schema and extract conclusions on how the optimization of its performance could be achieved.

6. References

[1] EVE (Educational Virtual Environment) prototype, http://ouranos.ceid.upatras.gr/vr/

[2] NS by example, http://nile.wpi.edu/NS/

[3] C. Bouras, E. Giannaka, “Performance Monitoring on Networked Virtual Environments”, 5th International Conference on Internet Computing (IC 04) Las Vegas, Nevada, USA, June 21 - 24 2004, pp. 302 - 308

[4] C. Bouras, E. Giannaka, T. Tsiatsos, “Issues for the Performance Monitoring on Networked Virtual Environments”, 7th ConTEL – International Conference on Telecommunications, Zagreb, Croatia, June 11 – 13 2003, pp. 725 – 728