

THE EVALUATION OF AN ACTIVE NETWORKING APPROACH FOR SUPPORTING THE QOS REQUIREMENTS OF DISTRIBUTED VIRTUAL ENVIRONMENTS

TATIANA BALIKHINA and ALI MAQOUSI *Faculty of Information Technology,
University of Petra, Amman, Jordan*

DAVID DUCE *School of Technology, Oxford Brookes University*

FRANK BALL *School of Design Engineering and Computing, Bournemouth University*

Abstract. This paper describes work that is part of a more general investigation into how Active Network ideas might benefit large scale Distributed-Virtual-Environments (DVEs). Active Network approaches have been shown to offer improved solutions to the Scalable Reliable Multicast problem, and this is in a sense the lowest level at which Active Networks might benefit DVEs in supporting the peer-to-peer architectures considered most promising for large scale DVEs. To go further than this, the key benefit of Active Networking is the ability to take away from the application the need to understand the network topology and delegate the execution of certain actions, for example intelligent message pruning, to the network itself. The need to exchange geometrical information results in a type of traffic that can place occasional, short-lived, but heavy loads on the network. However, the Level of Detail (LoD) concept provides the potential to reduce this loading in certain circumstances. This paper introduces the performance modelling approach being used to evaluate the effectiveness of active network approaches for supporting DVEs and presents an evaluation of messages filtering mechanisms, which are based on the (LoD) concept. It describes the simulation experiment used to carry out the evaluation, presents its results and discusses plans for future work.

Keywords: Distributed-Virtual-Environments; Active-Networking; Level-of-Detail; Packet-Filtering.

1. INTRODUCTION

In many respects large-scale Distributed-Virtual-Environments (DVEs) are similar to distributed multi-user games. They provide 3D graphical computer generated environments with stereo sound, supporting real-time collaboration between a potentially large numbers of users distributed around the world. They seek to provide a sense of presence and engagement akin to that experienced in real-world collaboration. DVEs were originally supported by local area networks; however, large scale DVEs will generally require the support of high performance wide area networks.

It is apparent that large scale DVEs offer a significant challenge to both end-system and network performance, and effective use of resources will need to be considered. Furthermore, the need to exchange geometrical information that provides the 3D graphical representation of objects in the virtual world in a timely, and in some cases real-time, manner, leads to Quality of Service (QoS) requirements that are additional to those in other types of multimedia application. An exchange of geometrical information usually requires a large quantity of data to be transferred reliably in a relatively short period of time. Generally, each transfer will then be followed by a relatively long period of silence resulting in a bursty type of traffic source. Smoothing out the bursts is possible, but is not an effective strategy as

this will reduce responsiveness and may interfere with the requirements of interactivity.

It has been proposed that large scale DVEs should be based on a Peer-to-Peer architecture since the client server model is less effective for large scale systems [Diehl, 2001]. However, making Peer-to-Peer communication scalable requires the use of multicast groups [Tanenbaum, 1996], which in turn need to be supported by a Scalable Reliable Multicast framework (SRM) [Floyd et al, 1997].

The Level of Detail (LoD) concept is based on the principle that the detail at which an object needs to be presented to the user depends on the distance between that user and the object in the virtual world, viz. the visibility relationship [Balikhina *et al*, 2002a]. By representing the object at different levels of detail in a way that can be built up additively, the visibility relationship can be exploited to reduce the resource requirements for communication, memory, and computation.

Within the network, LoD filtering can be employed in conjunction with SRM so that only the appropriate level of detail need be sent along each branch in the multicast tree. The filters are configured according to the individual requirements of the different destinations within the multicast tree, and this configuration may change dynamically in response to user movements within the virtual world. Both SRM and LoD filtering require processing resource within

the network additional to that needed for the end-to-end routing of packets. Therefore the concept of Active Networking in which computational resources are placed directly in the network to support end-user processing requirements would appear to be advantageous for supporting the communications requirements of DVEs. The aim of the work presented in this paper is to investigate the advantages that the active networking approach may offer for managing the demands of DVEs with particular focus on the benefits offered by LoD filtering.

We have investigated the performance benefits that could result from applying Level of Detail (LoD) filtering to the geometrical-information-traffic generated by a DVE, assuming that the DVE is supported by an active network infrastructure. This work was carried out using simulation models, which were developed using the OPNET simulation package. The investigation focused on the performance and resource utilization benefits that would be offered to both the hosts and the network. However, this paper is focused mainly on network traffic issues and presents the second in a series of simulation studies [Balikhina et al, 2002b] into the effectiveness of LoD filtering for improving the profile of the network traffic produced by a DVE.

The remainder of this paper is structured as follows: section 2 introduces DVEs and identifies the types of network traffic they are likely to produce; section 3 provides a brief description of active networks; section 4 outlines the LoD concept and briefly explains the process of LoD filtering; section 5 introduces the simulation model; section 6 describes the evaluation experiment that have been carried out using the simulation model and presents the results; and finally section 7 concludes and outlines future work.

2. DISTRIBUTED VIRTUAL ENVIRONMENTS.

DVEs provide 3D graphical computer generated environments with stereo sound, supporting real-time collaboration between potentially large numbers of users distributed around the world. A DVE system has four basic components [Singhal and Zyda, 1999]

- Graphics engines and displays
- Communication and control devices
- Processing systems
- A data network

It was at first only possible to use such environments where local area networks were available, though the push to higher performance public data networks

is now changing that picture. The prediction is that there will eventually be large-scale, persistent, distributed virtual environments, inhabited by millions of entities. This leads to the question of what architecture to adopt in order to support such environments in order that the same architecture will support growth from 10s to potentially millions of users.

The challenges raised by DVEs are threefold:

- How to provide network support for the mixed traffic types demanded by DVEs, typically voice, video, geometrical-information and data, with their very different characteristics with regard to network properties such as delay, jitter and loss rate?
- How to manage group participation in a DVE? Not all data are relevant to all participants in a DVE.
- How to manage persistent state within a DVE and manage state updates?

The work presented in this paper is primarily aimed at addressing the first of these three challenges. However, by consequence it also makes a contribution to the second.

DVEs are attracting increasing attention from the simulation community (for example battleground simulation) and it is possible that this could lead to the emergence of large-scale, persistent, distributed virtual environments, inhabited by millions of entities resulting in an even greater demand for network support, and presenting problems of scalability. Persistent DVEs are defined as DVEs that are “never ending” or “always on” [Abrams, 1999]. This raises the challenge of failure management and maintenance, and enhancement to components and services without the need to close down the environment. Together with a high degree of user interaction, these challenges places more demand on both processing power and communications bandwidth. In order to meet these challenges appropriate system and communications architectures are required to support DVEs.

2.3 System and Communications Architectures

A wide range of system architectures are possible, however, there are two extreme positions on a spectrum of possibilities [Diehl, 2001]:

- Client-server (with centralized state)
- Peer-to-peer (with decentralized state)

It is probably true to say that the client-server architecture is the most common nowadays.

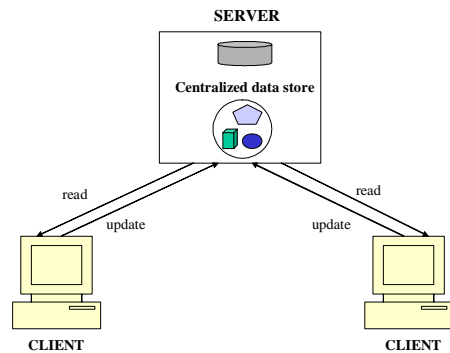


Figure 1. Client-server architecture

A single server manages the state of the virtual environment. A client sends an update to the server which is then propagated to other clients. All communication goes through the server, which becomes both a bottleneck and a single point of failure. As in Web applications, client caching strategies also cause difficulties and require sophisticated protocols to maintain consistency and freshness of state. However, the client/server approach is relatively simple to implement and provides a perfectly satisfactory solution for DVEs up to some size limit. This approach is also taken in Internet games servers such as Kali, TEN and Mplayer.

The peer-to-peer approach, on the other hand, distributes the state amongst all clients. Each client has a partial copy of the state. When a change occurs, a client has to send the changes to all other clients. If a client fails or leaves the environment, its data are lost, but other clients can continue without it (under some assumptions about how data are distributed). Peer-to-peer communication is a more promising approach to large scale DVEs than the client-server approach.

The key idea in making such architectures scalable is the use of multicast groups [Tanenbaum, 1996]. Multicast protocols enable a sender to send data to a set of hosts in a multicast group, avoiding unnecessary duplication of transmission by sending only one copy of the data along each arc in the routing tree between sender and receivers.

The DIVE (Distributed Interaction Virtual Environment) system, illustrated in simplified form in Figure 2, is built in this way [Frecon and Stenius, 1998] as is the MASSIVE system [Greenhalgh, 1998] and the earlier draft ANSI standard for protocols for distributed interactive simulation - DIS [ANSI, 1993].

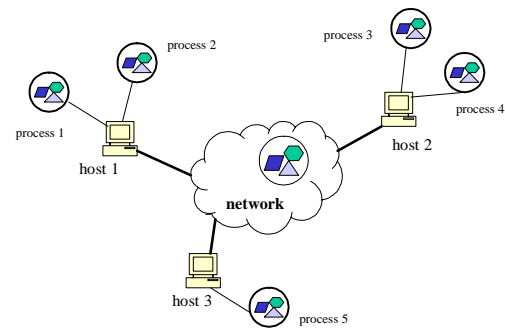


Figure 2. Peer-to-peer architecture

Multicast delivery systems were first used for the delivery of time critical data such as video and audio streams, where it is pointless to retransmit lost packets. The design of protocols to support scalable multicast for data-types, which demand reliable delivery is a very non-trivial problem.

The Scalable Reliable Multicast framework (SRM) [Floyd et al, 1997] provides basic functionality for scalable application level reliability in multicast environments. SRM is fully decentralized (no ring or central controller) and handles arbitrarily large groups of participants.

2.4 DVE Generated Network Traffic

DVEs need to exchange three types of information:

- 1) State-update or control information, which usually results in short messages that are sent either periodically, or in response to some event.
- 2) Continuous media, i.e. audio and video.
- 3) Geometrical information that describes the objects in the virtual environment. Generally this type of information is exchanged when a new user joins the virtual world or when objects mutate.

Each of these types of information results in a different type of network traffic.

The requirements of both control messages and continuous media traffic are already well supported by a number of network-level Quality of Service (QoS) support mechanisms. However, the requirements for exchanging geometrical information are quite different. An exchange of geometrical information in response to some change in the virtual environment (e.g. an object mutation or the arrival of a new player) usually requires a large quantity of data to be transferred in a relatively short period of time. Usually the transfer will then be followed by a relatively long period of silence. Therefore the type

of traffic generated by the transfer of geometrical information is potentially very bursty in nature. Although the bursts could be smoothed out, doing so will reduce the responsiveness of the DVE and may interfere with the requirements of interactivity. Generally, there will be a need to allocate relatively high bandwidth and support large queue sizes in order to accommodate this type of traffic. It is unlikely that these requirements could be accommodated easily by existing IP networks, and furthermore they would present a significant challenge to future multisevice networks. This provides a motivation to consider alternative approaches, and the concept of active networking would seem to be a potential candidate.

3. ACTIVE NETWORKS

In conventional networks, computational resources only exist outside the core of the network. The primary function of the network is end-to-end routing of user traffic. The Active Network approach challenges this assertion [Tennenhouse *et al*, 1997]. The key idea in Active Networks is that by placing computational resources directly within the network specifically to support end-user processing requirements, higher performance will be achievable than by conventional means, service deployment will be ameliorated and new classes of service will become possible. It has to be said that this is a contentious area of network research, not least for network operators who have to be convinced that the management, reliability and security of such entities are tractable problems. Nevertheless, there are research programmes in this area in both the USA and UK, and growing evidence that there are tangible benefits in the approach. We take a broader view of what constitutes an active network approach, and believe that the concept could also include use of an active network overlay that employs active elements which are placed strategically at certain points within a non-active network infrastructure.

Scalable Reliable Multicast is one class of protocol that has been shown to benefit from the presence of small amounts of session specific computation at key routing nodes, increasing throughput, decreasing latency and reducing congestion [Floyd *et al*, 1997] [Kasera *et al*, 2000]. Other suggestions include provision of transcoding services within a network, to address issues of bandwidth heterogeneity amongst destinations in a multicast group. The generalization here is that the application knows what transcoding processing to apply, the network knows where best to apply it.

The main focus of our investigation into active network support for large scale DVEs is to address the possibility of applying active filtering to the level of detail (LoD) concept [Balikhina *et al*, 2002a, 2002b, 2004]. To a certain extent we can see an analogy

between active video scaling for multicast, and interest filtering using transmission packets, which contain geometrical information for definite level of details. In the case of video, scaling the quality of received video will be defined by receiver's ability to obtain video stream (e.g. through congested line). The decision to forward, or not to forward, a packet is based on information the active router receives from the users regarding available bandwidth. However, for interest filtering, the distance between objects in the virtual world identifies the appropriate LoD needed for representation of the objects. Therefore in this case the forwarding decision is based on the spatial positioning of objects in within the virtual world.

4. THE LEVEL OF DETAIL CONCEPT AND LOD FILTERING IN AN ACTIVE NETWORK

Normally, a user agent has to load and render the whole scene. However, methods for spatial positioning divide a three-dimensional space into subspaces and compute a visibility relation or use a predefined one. By means of the visibility-relationship, both the bandwidth requirements and the computation time needed to render a scene can be reduced, since only visible subspaces need be transmitted and rendered. If the visibility of an object to a player depends on the viewer's distance from the object, then the browser can choose the appropriate scene graph for the object. The closer the viewer gets to the object, the more details are rendered. Multiple representations of one object with decreasing resolution (LoDs) are used to reduce the rendering cost for distant or otherwise less important objects. LoDs are additive, i.e., as the viewer moves closer to the object the next higher LoD can be added to the existing LoD.

As only one level of detail for a given object need be displayed at any one time, instead of transmitting the entire geometrical data relating to a scene, it is only necessary to transmit those packets that contain the appropriate LoD for each particular receiver. Therefore each packet that is transmitted will only contain information relating to one particular LoD, and each receiver will be sent a batch of packets for each LoD it requires. It should be noted that generally this type of application will generate two separate classes of network traffic: control information that is used to manage, and notify, changes in state; and geometrical information that describes the objects in the virtual environment. LoD filtering is only applied to geometrical information.

We assume that geometrical information is distributed across the hosts on which the DVE is implemented, and that the host upon which a user is resi-

dent holds information relating to the user's position in the virtual environment. The users of the DVE will belong to a multicast group within the physical network, and each user's host will be the source of a multicast tree to the hosts of other users. As users move within the virtual environment changes in their relative positions will determine the appropriate level of detail needed to view the objects within the virtual environment (to simplify discussion we shall assume that the only objects in the virtual environment are the representations of its users). If objects move closer to each other then higher LoDs will need to be transmitted and the new levels that are required may be different for each member of a multicast group. With LoD filtering each source of a multicast tree will transmit all packets that contain up to the highest level of detail needed within the multicast group served by the multicast tree. Active nodes along the multicast tree will filter out packets that contain a LoD higher than that required on their downstream branches. As users move, or new users join the virtual environment, the application will generate control messages to inform both the active nodes and the sources of the appropriate level of detail for the recipients. In addition to user movements, transmission of geometrical information can also result for object mutation, however, the transmission process take place in the same manner.

In contrast to the active approach with LOD filtering a non-active approach would generally require full geometrical information to be transferred when a new user joins (or an existing users rejoins) the virtual environment or when an object mutates. However, for the non-active case in general, transmission of geometrical information does not result from user movement within the virtual environment.

Generally, the amount of geometrical information required to represent an object in full detail can be quite large, and if updates within the DVE are to be sufficiently responsive then the transmission of geometrical information should be carried out as quickly as possible. Also given that the transfer of geometrical information will not generally be a continuous process this implies that the resulting traffic will be very bursty in nature and will place a high demand on the network for short intervals of time. Whilst such bursts could be smoothed out, this will result in longer transfer times, which in turn will reduce responsiveness. Furthermore, this type of traffic could be sensitive to packet loss, as the need for retransmission would also degrade responsiveness. Clearly, the LoD concept and LoD filtering have the potential to reduce burst sizes, however, the extent to which this can be effective will depend on the degree of user movement and the nature of user interaction within the virtual environment. The number of levels used can vary for different implementa-

tions, however, from the geometrical viewpoint there is little benefit to be gained by using more than 10 levels. In our work we consider 9 levels of detail (9LoD) and 3 levels of detail (3LoD) as these represent two different scales of the LoD concept whose performance can easily be compared using identical patterns of user movement.

In order to evaluate the effectiveness of LoD filtering we have built a simulation that can model a range of user interaction and the resulting transmission of geometrical information traffic.

5. THE SIMULATION MODEL

The simulation model comprises two basic elements, hosts and routers that can be used to study DVEs of various scale and configuration. However, the scale of these models is obviously constrained by the availability of resources such as memory and processing power in the simulation environment.

The Host elements produce control packets and geometrical-information packets, the generation of both being driven by the simulated interaction of the user within the virtual environment. User interaction is simulated through events and actions that are generated by a process that can model different categories of user movement. The router element models an active network node that also has multi-service capability. A simple Class Based Queuing (CBQ) system is provided at each output link. In this way DVE geometrical information traffic can be allocated its own class (Geo_Info Class), assigned its own portion of the link bandwidth and therefore will be isolated from other classes of traffic. For simplicity the router is internally non-blocking, so queues will only build up in the output stages. The router can distinguish between active and non-active packets, which allows DVE packets to be sent using either active or non-active mode. In this way it is possible to compare the performance advantages gained from the active approach, by running the same scenario in both active and non-active mode.

Currently, the only objects that exist in the virtual environment are the representations of the individual users. The spatial relationship between users within the virtual environment is modelled using a 9*9 grid that subdivides the virtual environment. Users can be placed either randomly or selectively, on squares within this grid and then at discrete intervals time they either remain in position or move to any adjacent square according to some rules and random factors, the parameters of which can be adjusted to simulate various categories of user interaction. The set of initial placements together with the set of movements that take place over a predefined interval of time will generate a particular scenario that can

be used to determine the events that drive the generation of both control packets and geographical-information packets. When modelling 9LoD these events occur as users change position within the 9*9 grid, whilst in the case of 3LoD events occur when users change position within a 3*3 grid that is overlaid on the 9*9 grid. The scenarios are recorded and therefore the same scenario can be used to drive the simulation in either active or non-active mode, and in the active case for both 3LoD and 9LoD. Hence consistence is provided when comparing the performance of the three different options.

So far we have identified four general types of distinctly different movement patterns, which we have categorised as:

- a) *Random*: in which users move randomly within the grid and can leave and rejoin the virtual environment at the edge of the grid.
- b) *Far*: in which the user moves around in regions of the grid that are distant from each other but is not permitted to leave the virtual environment
- c) *Centre*: in which user movement clusters around the centre of the grid.
- d) *Rejoin*: in which user movement is concentrated near to the outside of the grid with a moderately high probability that a user could depart and later rejoin the virtual environment.

To help categorise the movement scenarios produced by our user movement simulator we have built a visualisation tool that plots the movements over time. Each user is represented by a different coloured square and the brightness of the square relates to how recently the user was in that position. Figures 3 and 4 shows the movement patterns for the *Far* and *Random* categories respectively.




Figure 3 Movement pattern *Far*




Figure 4 Movement pattern *Random*

The plots serve to show the general pattern of user movements during a particular scenario, and by observation we can then decide to which category of movement pattern the scenario belongs.

6. THE EVALUATION EXPERIMENT

For this evaluation experiment we used the simulation elements described in the previous section to build a model of a small scale DVE that comprises four hosts and five routers, as shown in figure 5.




Figure 5 Network Configuration

The service rate allocated to the Geo_Info class between the hosts and the routers was 2Mbit/s. The link rates between all the routers are 100Mbit/s, and the service rate assigned to the Geo_Info class was varied between 2 and 4 Mbit/s in steps of 0.5 Mbit/s.

All packets that convey geometrical information, whether that information was progressive (i.e. using LoD) or otherwise, were a fixed size of 8kbits. To simulate the processing delays associated with accessing the geometrical information and preparing the packets for transmission, the individual packets were generated according to an exponential distribution.

Within the model the transmission of a block of packets was triggered by a certain event, generally the receipt of a control packet, with the size of the block being determined by the highest LoD required. The mean packet inter-generation time used to simulate information retrieval and packet processing was set at 0.01s, giving a mean traffic generation rate of 0.8 Mb/s for each source.

The objective of this experiment was to compare the performance of the non-active and active (3LoD and 9LoD) approaches with regard to the demand placed on the routers with respect to queuing requirements, and therefore we observed the maximum queue occupancy and utilisation of the queues associated with the Geo_Info Class. For this experiment we assumed sufficiently large queue sizes so that we can observe the behaviour of the traffic unaffected by packet losses.

The simulation was run for both the active and non-active cases, and the active case considered geometric representations that used both three levels of detail and nine levels of detail. Simulations for these three cases were carried out using scenarios for the *Random*, *Far*, *Centre*, and *Rejoin* categories of movement patterns using a simulation time of 1 hour. The choice of 1 hour for the simulation time was determined during preliminary experiments in which we found that extending the simulation time beyond 1 hour made no significant difference to the results [Balikhina, 2005]. The parameters of interest, i.e. the maximum queue occupancy and queue utilisation pattern, were observed and recorded for the Geo-Info class output queue on each of the five routers.

6.1 Results from the Simulation

Obviously the results differed from one queue to another, however, they all showed the same general trend regarding the comparison between the active and non-active approaches. Therefore we only present the results gathered from one particular Geo_Info queues, that of router 3, these results being typical and representative of the general behavior at all the queues. However, for those interested, the results from the other queues, together with results from other separate, but related, experiments, are presented in [Balikhina, 2005].

6.1.1 Maximum Queue Occupancy

The maximum queue occupancies v the assigned service rates are presented below in Figure 6.

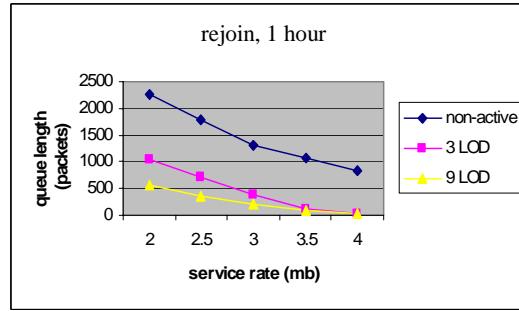


Fig. 6a

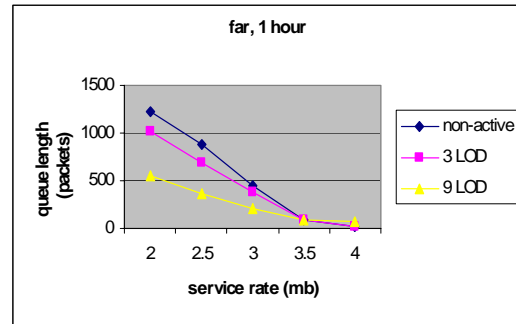


Fig. 6b

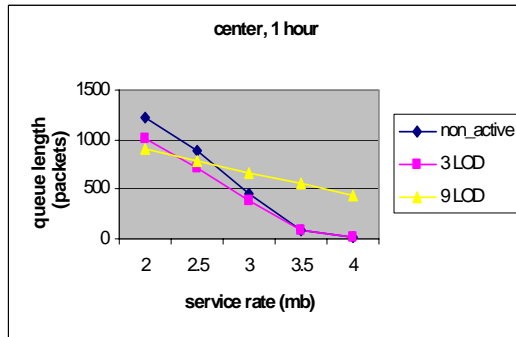


Fig. 6c




Fig. 6d

Figure 6 Maximum Queue Occupancy v Service Rate

In the case of the *Rejoin* scenario 9LoD offered a significant reduction in the maximum queue occupancy throughout the entire range of service rate allocation, and 3LoD performed better than the non-active case for service rate allocations of less than 3.5Mb/s (see figure 6a). In this case most of the user movement took place around the edges of the grid with users generally being remote from one another. Therefore with the active approach, many instances of movement, or rejoining, resulted in only a few levels of detail being sent. However, with the non-active approach full geometrical information needed to be sent for each instance of a rejoin, resulting in much higher values of maximum queue occupancy.

With the *Far* scenario the pattern of user movement was very similar to that of the *Rejoin* scenario, except users did not leave and rejoin the virtual environment. Therefore the performance of both 3LoD and 9LoD were similar to that in the *Rejoin* case (see figure 6b). As users are not leaving and rejoining the virtual world, with the non-active approach traffic is only generated at the initialisation stage, and this results in a significantly lower value of maximum queue occupancy than that for the *Rejoin* scenario. However, 9LoD performed significantly better, and 3LoD performed marginally better, than the non-active approach for service rate allocations less than 3.5Mb/s.

User movement in the *Centre* scenario is concentrated in at centre of the virtual world rather than at the edges, with the users generally moving in close proximity to each other. As with the *Far* scenario users do not leave and rejoin the virtual environment. Therefore the performance of the non-active approach is very similar to that for the *Far* scenario. However, because users are moving in close proximity to each other, for the active approach movement will generally require the transfer of a high number of levels of detail, therefore lowering performance (see figure 6c). In this case 9LoD only performs marginally better than the non-active approach for service rate allocation less than 2.5Mb/s, and performs much worse for allocations greater than 3 Mb/s. 3LoD offers little improvement on the non-active case but performs better than 9LoD for service rate allocation greater than 2.5 Mb/s. The reason why 3LoD performs better than 9LoD in this case can be explained by the fact that 3LoD uses a grid of lower resolution and therefore, for the same scenario, the events that trigger transfer of geometrical information will occur less frequently. However, in many cases both 3LoD and 9LoD will need to send geometrical information at the highest level of detail.

The *Random* scenario produces the combined effects of users leaving and rejoining the virtual world and

users moving in close proximity to each other. In many ways it is similar to a combination of the *Rejoin* and *Centre* scenarios except that movement is mainly concentrated near one edge of the grid, and also the probability of a user leaving and rejoining the virtual world is less than in the *Rejoin* case. The impact on the non-active approach of users rejoining is much less than in the *Rejoin* scenario therefore the performance of the non-active approach is very close to that observed in the *Far* and *Centre* cases. However, the combination of users rejoining and moving in close proximity to each other has a more significant effect on the active approach. In this scenario 9LoD gave worse performance than the non-active approach in all cases except for that of the lowest service rate allocation (see figure 6d). The performance of 3LoD was either worse or not significantly better than that of 9LoD throughout the range of service rate allocation. The better performance of the non-active approach in this scenario is due mainly to the following reasons:

- 1) Whilst the active approach needs to respond to both user movement and rejoins, the non-active approach only need respond to rejoins.
- 2) Users moving in close proximity to each other make it more likely that the active approach will need to send geometrical information at all levels of detail.
- 3) The amount of data required to represent a geometrical object at all levels of details is generally greater than that needed to represent the same object in the non-active approach

6.1.2 Queue Utilisation

As would be expected the traffic generated by the transfer of geometric information was seen to be very bursty, resulting in the need for large queue sizes, but giving a low overall utilisation of the queue. Obviously, queue utilisation decreases with increase in assigned service rate.

The results show that there was no significant difference in the overall utilisation of the queues between the active and non-active cases, except for the *Rejoin* scenario, in which the overall utilisation was slightly higher for the non-active case. However, even in this case, with an assigned service rate of 2Mb/s the queue was still seen to be empty for 89% of the time. There was, however, a significant difference in the pattern of utilisation over time between the active and non-active cases, and there was also a difference between the utilisation patterns of the 3LoD and 9LoD cases.

The different patterns of queue utilisation for an assigned service rate of 2Mb/s in the two contrasting

cases of *Far* and *Random* are shown below in figure 7.

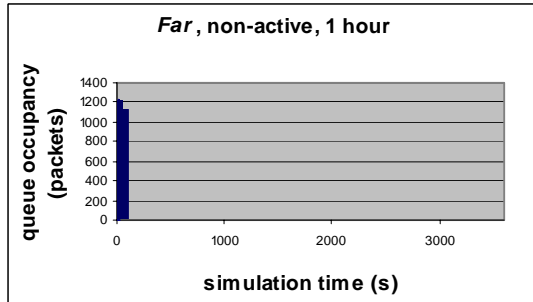


Fig. 7a

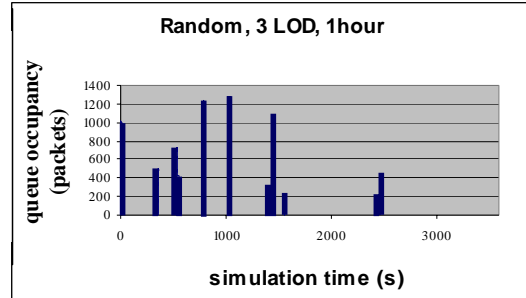


Fig. 7e




Fig. 7b

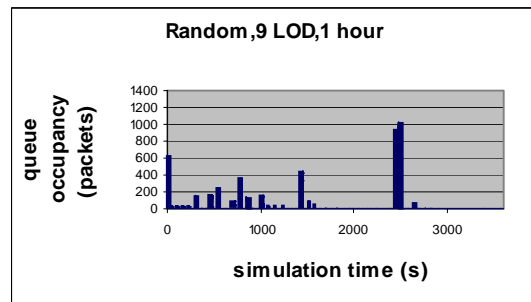


Fig. 7f

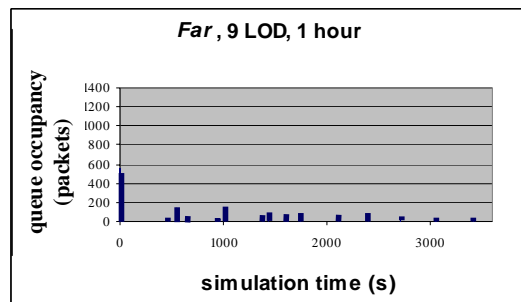


Fig. 7c

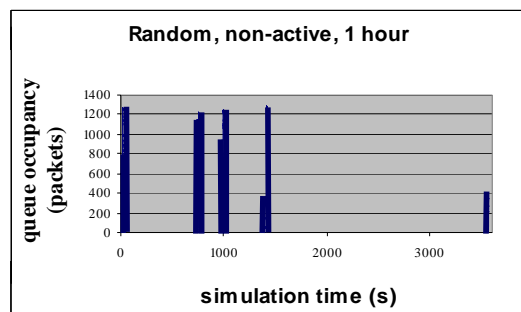


Fig. 7d

Figure 7 Patterns of Queue Utilisation

These figures show a very similar contrast to that observed between the *Centre* and *Rejoin* scenarios, with the pattern for *Centre* being very similar to that of *Far*, and likewise for *Rejoin* and *Random* except for the scale of queue occupancy.

In the *Far* scenario the non-active approach only produces traffic at the initialisation stage during which the queue occupancy is high. However, following the initialisation stage the queue remains empty for the remainder of the simulation time (see figure 7a). The utilisation pattern is similar for 3LoD but the initial queue occupancy is much lower (see figure 7b). For 9LoD there is a initial burst of queuing activity with a queue occupancy significantly lower than that of the 3LoD case, followed by bursts of queuing activity throughout the simulation period in which the queuing occupancy is significantly lower than that of the initial burst (see figure 7c).

With the *Random* scenario the non-active approach produces a similar initial burst of queuing activity to that of the *Far* case followed by bursts of queuing activity throughout the simulation period that are triggered by users rejoining the virtual world. In these subsequent bursts of queuing activity the level

of the queue occupancy is influenced by the time between rejoin events (see figure 7d). 3LoD also produced a initial burst of queuing activity similar to that in the *Far* scenario, however, in this case it was followed by bursts of queuing activity triggered by both user movement and users rejoining the virtual world. These burst occurred more frequently than for the non-active approach but in the majority of cases the level of queue occupancy was significantly lower than for the non-active approach. However, approximately 30% of the bursts resulted in a level of queue occupancy similar to that for the non-active approach (see figure 7e). The initial burst of queuing activity produced by 9LoD was once again similar to that of the *Far* scenario. This was also followed by bursts of queuing activity triggered by both user movement and users rejoining the virtual world. These burst occurred more frequently than for 3LoD, but in all but one case the level of queue occupancy was very much lower than for both 3LoD and the non-active approach (see figure 7f). In the one case where the level of queue occupancy was only slightly lower than the maximum values of both 3LoD and the non-active approach, an examination of the scenario script suggests that the high level of queue occupancy experienced during this burst was due to an adverse combination of movements and rejoins. The extent to which such adverse combinations could reduce the overall effectiveness of the Level-of-Detail approach is something that warrants further investigation.

The different patterns of queue utilisation observed in these results may affect situations where the bandwidth allocated to the Geo_Info class is shared on a lower priority basis by other classes of traffic. Generally, the non-active approach will result in fewer interruptions to the service of the lower priority classes than LoD filtering. However, the duration of each interruption will usually be longer. Also 9LoD will generally result in more frequent, but shorter interruptions, than 3LoD.

6.1.3 Summary of Results

The results show that for reducing the value of maximum queuing occupancy LoD filtering is not advantageous in all cases, and its performance relative to the non-active approach is sensitive to both movement scenario and service rate allocation. Also it can be seen that in a limited number of circumstance using 3LoD can give a better result than using 9LoD.

For scenarios where users are generally not moving in close proximity to each other, e.g. the far scenario, LoD filtering, particularly 9LoD, offers a significant reduction in the value of maximum queue occupancy, except for cases where a comparatively

high service rate is allocated. If such scenarios also include a significant amount of rejoining activity, i.e. the *Rejoin* scenario, then the results indicate that LoD filtering will remain beneficial for significantly greater levels of service rate allocation

Generally, LoD filtering performs less well than the non-active approach for scenarios where the users are moving in close proximity to each other, particular if this is combined with rejoin activity, e.g. the Random Scenario. However, even in this case, for comparatively low values of service rate allocation (2Mb/s in this experiment), 9LoD still offers some benefits.

It is also important to note that even in the cases where LoD filtering offers significant improvement, geometrical information traffic still requires relatively large queue sizes unless a comparatively high service rate is allocated. As the advantages offered by LoD decreases with increase in service rate then using a relatively high service rate allocation with the non-active approach could be seen as being equally effective to LoD filtering. Given the overall low utilisation of the queues the bandwidth allocated to the Geo_Info class could be made available to other classes of traffic on a lower priority basis at times when the Geo-Info queue is empty. In this way the bandwidth would not be wasted but absorbed by the lower priority classes. However, there will always be a limit to the amount of bandwidth that can be allocated in total to the Geo_Info class, and higher levels of allocation to individual multicast groups will mean that less multicast groups can be accommodated.

7. CONCLUSIONS AND FUTURE WORK

This paper has presented some performance modelling work that is related to a more general investigation into how Active Network ideas might benefit large scale Distributed-Virtual-Environments (DVEs). It has introduced a simulation-based approach that is being used to evaluate the effectiveness of messages filtering mechanisms, which are based on the Level of Detail (LoD) concept. The LoD concept has been briefly explained, and the simulation model that was developed to carry out this study has been described. Evaluation experiments using the simulation model have been outlined and the results have been presented

The results show that in certain cases there are significant benefits offered by the LoD filtering approach, whilst in others it performs less well than a non-active approach, particularly for higher levels of service rate allocation. Furthermore, even in cases where LoD filtering results in a significant reduction

in the maximum queue occupancy, the maximum queue lengths are still quite high. We believe that LoD filtering has a role to play in supporting the communication requirements of DVEs and that it may also offer benefit at the application level since it also has the potential to reduce memory requirements in the hosts. However, we also believe that on its own LoD filtering it is unlikely to be sufficient to fully meet the needs and QoS requirements of Geometrical information traffic and other mechanisms may also be needed to provide further support.

Recently the simulation model was modified to increase the number of hosts to seven, and then used to repeat the majority of the experiments described above. However, these results showed that apart from scale, the pattern of behaviour remained the same as those presented in this paper. With the modified simulation model it was not possible to repeat all the experiments for a simulation time of one hour due to resource limitations, and this suggests that we have reached the limit of scalability for the model using our current simulation environment.

In future work we will continue our study of the active approach by further addressing the issues raised in this paper and investigating the possibility of combining other mechanisms, including caching and predictive forwarding, with LoD filtering to further improve traffic profiles and offer better support to geometrical information traffic.

We will also investigate new techniques to allow us to extend the scale and complexity of the simulation models so we can continue to study the problem using a larger number of users and a larger scale network. Furthermore, we will seek to link the output from studies of user behaviour into our methods of generating scenarios of user interaction, further develop the visualisation tool, and possibly extend the range of categories for patterns of user movement.

References

Abrams H.H, "Extensible Interest Management for Scalable Persistent Distributed Virtual Environments", PhD Thesis, Naval Postgraduate School, Monterey CA USA, 1999.

ANSI "Standard for information technology, protocols for distributed interactive simulation", DIS-ANSI IEEE Standard 1278-1993, American National Standards Institute, 1993.

Balikhina T, Ball F and Duce D, "Distributed Virtual Environments – An Active Future? ", *Proceedings of the EGUK 2002 conference*, June 11-13 2002, Leicester, UK. M. Jones (ED), IEEE Computer Science Press.

Balikhina T, Ball F and Duce D, "Active Network Support for Distributed Virtual Environments", *Proceedings of 3rd Postgraduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting, PGNet 2002*, pp328-334 (June 2002) Liverpool, UK.ISBN 1 902550 086.

Balikhina T, "Active Network Support for Distributed Virtual Environments", PhD Thesis, Oxford Brookes University, August 2005.

Diehl S, *Distributed Virtual Worlds*, Springer, 2001.

Floyd S, Jacobson V, Lin C.G, McCanne S and Zhang L. "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing", *IEEE/ACM Transactions on Networking*, December 1997.

Frecon E and Stenius M, "DIVE: A Scalable Network Architecture for Distributed Virtual Environments", *Distributed Systems Engineering Journal*, Special Issue on Distributed Virtual Environments, 5(3), 1998.

Greenhalgh C, *Large Scale Collaborative Virtual Environments*, Springer, 1998.

Kasera S, Bhattacharyya S, Keaton M, Kiwiore D, Zabele S, Kurose J and Towsley D, "Scalable Fair Reliable Multicast Using Active Services", *IEEE Network Magazine*, 14(1), 2000.

Singhal S and Zyda M, *Networked Virtual Environments – Design and Implementation*, Addison-Wesley, New York, 1999.

Tanenbaum A.S, *Computer Networks*, Prentice-Hall, 1996.

Tennenhouse D, Smith J, Sincoskie W., Wetherall D and Minden G, "A Survey of Active Network Research", *IEEE Communications Magazine*, 35(1), pp. 80-86, 1997.

BIOGRAPHY



Tatiana Balikhina received an MSc. in computer engineering from Electrical Engineering Institute, St. Petersburg, Russia, 1989 for her research in image processing. Following graduation she was involved in developing software for business applications, and was a database administrator and main developer for finance department at University of Petra (UOP) from 1998 – 2002. In 2004 she joined the Faculty of Information Technology as a teaching assistant. She was awarded a PhD in computer science from Oxford Brookes University in 2005 for her work on network architectures for large scale distributed virtual environments. Since 2005 she has been an assistant professor in the Faculty of Information Technology at UOP and is currently involved in research relating to network architectures for large scale distributed virtual environments, network performance, computer and Internet security awareness.



Ali Maqousi received an MSc. in computer engineering from Electrical Engineering Institute, St. Petersburg, Russia, 1989. He was a network administrator and part-time teacher assistant at University of Petra (UOP) from 1993 – 1997 and full-time teacher assistant from 1999 -2003. He was awarded a PhD in computer science from Oxford Brookes University in 2003 for his work on providing Quality of Service (QoS) in packet switched networks. Since 2003 he has been as assistant professor in the Faculty of Information Technology at UOP) and is

currently involved in research relating to multiservice networking, network performance, computer and internet security awareness, and network architectures for large scale distributed virtual environment



David Duce is a Professor in Computer Science in the Department of Computing, School of Technology at Oxford Brookes University, where he leads the Multiservice Systems research group. He has contributed extensively to the development of standards for computer graphics including the ISO/IEC standards GKS, PHIGS and most recently PNG. His current research interests are in distributed cooperative visualization and web graphics, He is currently Vice Chairman of Eurographics.



Frank Ball is a Reader in Computing in the School of Design Engineering and Computing at Bournemouth University and a member of the Multimedia Communications Research Group. Since 1988 he has been involved in research relating to Quality of Service (QoS) in Multimedia Communication Systems, performance modelling of Multiservice Networks., and the development and evaluation of network level QoS support mechanisms. More recently his research interests have extended to include the development of multilevel approaches to performance modelling and evaluation.