Perceptually-Oriented Interest Management In Large-Scale Networked Virtual Environments

I. Dunwell & Dr. J. C. Whelan

Center for Internet Computing, The University of Hull, UK

Abstract

Amongst the most significant challenges in developing large-scale multi-user virtual environments is the efficient filtering of data to each user - a process commonly described as "interest management". This work-in-progress paper presents a broad summary of existing approaches, placing an emphasis upon the relationship between interest management and human perception. Subsequently, an introduction to the challenges in evaluating the success of interest management, given such a relationship to perception, is presented. The initial development of a test environment aimed at overcoming some of the challenges in providing a platform for such evaluation is then described, together with discussion of a perceptually-oriented approach to interest management which relies on the description of perception as a dynamic field, formed by analysis of the user's focus.

1. Introduction

As virtual environments become more sophisticated in terms of both content and interactivity, the potential for a single, extensible virtual environment (VE) capable of supporting massive numbers of users has frequently captured both the attention of researchers and the imagination of the general public alike. The implementation and advancement of such massively-multi-user networked virtual environments (MMVEs) presents many challenges in a wide range of research areas.

Amongst these challenges, one of the most significant and complex is the filtering of network traffic to ensure an optimal use of available bandwidth. This filtering, or "interest management" (IM), involves ensuring only content relevant to each user is transmitted via the network. Current approaches aiming to provide IM are wide-ranging, and typically integrated into existing simulations. Such integration often comes at the cost of a certain degree of re-usability and makes evaluation of the generic potential of an approach difficult. Following a limited summary of existing approaches and their history, the challenges in such evaluation are discussed in Section 3. This discussion is followed by an outline of a test system intended to provide the grounds for such an evaluation, and, given such facilities, the concept of providing a perceptually driven approach centered around a user's focus is presented.

2. Background

Some of the earliest taxonomies of IM approaches [Mor96] identified the concept of intrinsic filtering - filtering based upon that attributes of entities within an environment. Further classification of these techniques leads to a subsequent distinction emerging between grid and aura-based approaches. Grid techniques, pioneered largely by SIMNET [MT95], revolve around partitioning virtual environments into regions of interest (commonly referred to as locales [BWA96]), with clients being considered relevant to all other clients within the same locale. By comparison, aura-based approaches such as those used in early systems such as MASSIVE [Gre98], define a region centered upon each client. If the interest regions of two clients intersect then they are considered to be mutually visible.

These two techniques have been refined and combined over recent years, in response to the rapidly emerging need for efficient IM within large-scale virtual environments. Commercial endeavors (primarily within the gaming industry) have brought large-scale multi-user virtual environments to the general public, creating online economies with substantial turnovers, recently described as being in excess of the economies of a significant number of real-world countries [Cas01]. As hardware continues to advance, significant potential for further applications remains. Hence the commercial need for efficient IM is becoming increasingly apparent.

One of the most obvious, and successful, IM techniques is the combination of multiple layers of interest management within a single system. Zyda et al. [AWZ98] demonstrate a three-tiered approach within the Bamboo toolkit for distributed environments, illustrating the effectiveness of combined IM paradigms. The combination of grid and aura-based approaches is commonly used - at their most fundamental level, grid and aura-based techniques both provide a simplistic and efficient means for filtering data within a large-scale VE. More recent refinements have sought to add additional aspects to the filtering process, such as considering visibility more closely [HPG02]. Whereas more basic aura and grid based systems filter spatially, irrespective of surrounding geometry, further culling based on visibility presents a clear example of the potential for refinements based on a better understanding of a users perceptual needs. The system is successful, noting a significant performance gain with no perceived loss in quality. Hence the potential for such visibility-oriented interest management in a large-scale virtual environment centered around a replicated-database model is clear. However, the potential for perceptual consideration is broader than visibility-based culling would suggest, since an interest manager for a truly extensible environment must take into account other aspects, such as extensibility, sound, and interactivity.

Thus whilst visibility alone is unquestionably a valuable criterion by which to perform filtering, the notion of perceptual relevance offers far more significant scope for refinement. Much as Reddy [Red01] applies perceptual techniques to rendering, with a significant performance gain, so may perception be considered in IM. Beharee et al. successfully demonstrate the potential for the use of perceptually oriented IM in a variety of simulations, ultimately concluding with a simulation of an 800-client virtual city [BWH]. The perceptual culling is based around the notion of change blindness; in effect aiming to filter out items of data in which the user is unlikely to identify a change. Change blindness itself is a significant field of research, and a large volume of work offers a variety of further possibilities for enhancement [Tse04].

Beharee et al. focus heavily on obtaining results with no noticeable decrease in quality of simulation. In a massive-scale virtual environment, situations in which available bandwidth is insufficient to offer such a result can prove commonplace, and hence analysis of which perceptual qualities offer the greatest scope for 'graceful degradation' is also of considerable importance. With regards to extensibility, as environments become more sophisticated - particularly in terms of object interactivity - the replicated data model common to the majority of existing systems (whereby simulation content is installed or downloaded to each client prior to run-time), prevents expansive and highly interactive environments. Subsequently an increasing level of importance can be placed upon creating IM capable of filtering textures, geometry and other world content to each user. Furthermore, the nature of input devices to virtual environments allows for a unique opportunity to analyse the behaviour of each client, and tailor the filtering process in response. Such behavioural analysis can range from following user behaviour and mouse input to more sophisticated data capture techniques such as eye-tracking [Jac95]. Since this effectively permits a users behaviour to be analysed at run-time, it becomes possible to consider filtering approaches which are capable of tailoring themselves more closely to behavioural input.

Whilst each of these aspects presents individual potential for refinement of IM, a recurring theme when examining existing approaches is a significant difficulty in comparison and evaluation, due to both the qualitative nature of perceptually oriented results and the application dependence of many IM approaches. The subsequent section aims to define this challenge in more detail, and goes on to describe a test system developed to address it.

3. Evaluating Interest Management

3.1. Generic IM

Singhal (a pioneer of the aforementioned Bamboo toolkit) stated confidently as early as 1998 that "Eventually there will be a persistent VE shared simultaneously by billions of participants" [ZS00]. This vision remains a long-term goal of MMVE developers, due in large part to the commercial potential for having such an immense audience co-existing within a single environment. Providing such a system implies two main factors need to be addressed; firstly the technical barriers in place must be overcome (such as bandwidth and rendering limitations), and secondly, on a more conceptual level, the environment must be sufficiently ubiquitous and extensible to fulfill the needs of all users. Given such needs may range from simple social interaction, through simulation and collaborative working, such fulfillment is no simple task.

Clearly, the role of interest management will be pivotal in providing such an environment; without the means to efficiently and rapidly control communication between clients such an application would prove impossible. Primarily, extensibility in such an environment implies a degree of dynamism within an interest management routine, due to the differing needs of users and expanding environment. However, the wealth of existing environments with broad ranging applications and near-identical interest management approaches strongly suggests that despite a need to interest management to suit a specific environment, generic principles do exist. The approach to interest management within the e-Agora system [Mr02] shows one such approach to generic interest management, in this case attempting to utilise 'general variables' to store different types of data.

Whilst there are many different aspects that may be considered in establishing a generic 'theme' of virtual environment, some of the most fundamental include:

1) The concept of each client as a virtual embodiment of a human. This brings with it a variety of notions, such as the centering of the region of interest around the client, and a relationship between the scene rendered to the client and the information required.

2) The finite amount of bandwidth available, implying a need for graceful degradation and, more fundamentally, scalability, in areas of high traffic (such as a crowded virtual room).

3) The implicit relationship between the software and hardware underlying the simulation and the simulation itself. Multiple levels of interest, for example, offer diminishing returns with increasing numbers of levels due to the fact they require formation of additional multicast groups [PB95].

4) A basis for the simulation in the real world. Whilst virtual environments exist which do not aim to provide such simulation, for the purposes of this context we choose to focus on those aiming to provide a 'realistic' simulation on a human avatar within a virtual world. This has broad implications regarding the relationship between the positions and properties of objects relative to the client and their perceptual relevance.

We may build upon these concepts to create interest management paradigms that, whilst not truly application-independent, lend themselves to a significantly broad range of applications. In particular we choose to focus on environments where the client is considered to be a virtual human, and thus their relationship with the VE considered parallel to the relationship of a human within the real world. As an important side-note, whilst the employment of any interest management technique has obvious overheads in terms of processing, as hardware advances and the depth of interactivity within virtual environments increases, bandwidth can be seen to be fast emerging as the true bottleneck for such environments. Hence despite a clear need for processor efficiency in any interest management approach, its application may not necessarily imply any tangible negative effects upon the simulation. It may, however, prove beneficial to accept some loss of simulation quality if a corresponding gain in bandwidth reduction is possible. Balancing - or providing some degree of scalability towards the two distinct criteria of bandwidth reduction and impact upon quality of simulation is hence important to any IM approach. Further defining these criteria is the focus of the next section.

3.2. Bandwidth vs. Quality of Simulation

Of crucial importance to the development of any interest management approach for such a general application is the provision of means to test and evaluate differing approaches against one another. Due to the breadth of existing large-scale VE applications, and the scope for future development, evaluating the success of differing approaches in a generic fashion proves a complex task.

One clear criteria for evaluation is the decrease in bandwidth consumption achieved. However, this must be weighed against the impact upon the quality of simulation. This is a broad-ranging concept, which must take into account a large number of qualitative variables in order to attempt to establish how a simulation 'feels'. Obviously, some of the common effects of interest management, such as entities 'popping up' near the user have a strong negative impact on this quality. Similarly, the effects on frame rate of a highly CPU-intensive interest manager may be factored into this broad notion. Thus we can loosely define the success of any interest management approach in a generic context to be the trade off between the reduction in bandwidth consumption compared to the perceived decrease in quality of simulation.

In an attempt to focus on such interest management issues generic to large scale VEs, and to better measure this balance between bandwidth and quality of simulation, a system is currently under development intended to simulate limited aspects of the data model of a peer-to-peer, multicast-based, massive-scale virtual environment. This simulation also includes behavioral simulations of large numbers of clients. This system is intended to allow for the evaluation of these two main concepts of quality of simulation and bandwidth consumption via the following approaches:

1) Qualitative evaluation of the perceptual effects and hence quality of simulation by means of a rendering engine. Users can be placed within the virtual environment and navigate through it freely, observing the surrounding scenery and clients. It is hoped by analyzing the response of a suitably large sample of users, some conclusions can be reached regarding the success of differing techniques in this respect. Of particular emphasis is the success of techniques as bandwidth constraints force quality of simulation to deteriorate.

2) Quantitative prediction of the bandwidth consumption for a variety of client behavioral models. This output may be compared alongside the qualities of simulation in order to better appreciate the effects of differing interest management approaches in as quantitative a method as possible.

4. Developing a Test Simulation

4.1. Overview

The test system developed (via a combination of Java and Java3D) bases itself around the notion of a massive scale, peer-to-peer, virtual environment, such as that of DIVE [FS98]. The system is highly modular, with the emphasis being placed on allowing components to be interchanged and modified as easily as possible.

Each client behaves as an autonomous entity running in its own thread, and, as a consequence of the modular nature, may have its motion governed either by user input or an AI routine. Current AI routines are simplistic, ranging from purely random motion to a 'point of interest' approach whereby points may be inserted into the virtual environment which attract clients, hence allowing for a more realistic simulation of client distribution within a virtual space (since clients would seldom distribute as evenly as a random motion algorithm would suggest).

Output on total traffic between clients is streamed to files, in order to assess client load distributions and bandwidth consumption. The environment is also rendered in real-time 3D (as shown in Figure 1) for purposes of evaluating quality of simulation. The rendering engine is currently primitive; capable only of showing both overviews of the system and a view from the perspective of a single client. Future work aims to extend this into providing a user-controlled client within a more realistic environment to better assess the quality of the simulation for a given interest management approach.

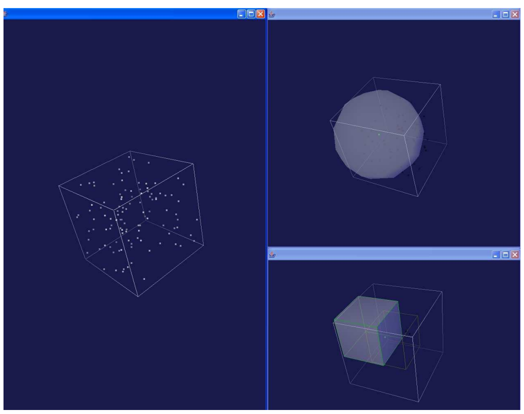

Figure 1: *Universe overview (left) alongside visualizations of both aura (top right) and grid based (bottom right) interest management implementations, focused on a single client*

Quantitative evaluation of bandwidth consumption is performed by keeping track of the number of total update packets sent through the simulation as it runs, which would in a real-world application correspond directly to packets sent via the network (each containing an update on nearby client attributes). By directly keeping track of the total number of packets sent over time by means of a data log, the effects of events within the simulation can clearly be observed.

4.2. Comparison of Existing Approaches

Figure 2 shows output for two basic grid and aura-based interest management implementations, alongside output with no interest management. In all cases clients were initialized at the center of the virtual universe and allowed to randomly distribute outwards, hence providing an illustration of the effects of a shift from an area of high client density to a lower one.

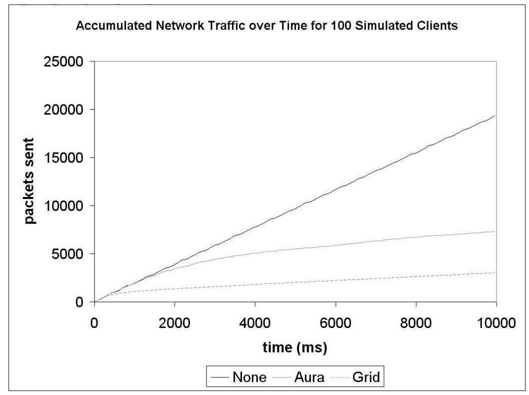

Figure 2: *Packets sent / time for 3 interest management simulations*

As can be observed, initially with all clients at the center of the universe and hence within each other's respective areas of interest, bandwidth use increases linearly with time in all 3 simulations. As clients distribute, the simulations running interest management deviate from this linear relationship, using less bandwidth as distant clients cease requesting updates. Whilst the illustration shown is for the case of 100 clients, identical behavior can be observed in simulations containing both 10 and upwards of 1000 clients. Whilst in this case the results may be similarly proven mathematically, the system is ultimately intended to allow for the interaction of human users, where such proof would be impossible.

Using such output, it becomes possible to compare bandwidth usage characteristics of differing interest management techniques. However, there is currently an extremely limited facility for evaluating the quality of simulation, and hence it proves difficult to draw any conclusions from bandwidth analysis alone. Establishing such facilities is a primary development goal, and hence as mentioned in the previous section bandwidth will ultimately be examined directly alongside quality of simulation. Future work intends to allow for contrasting techniques to be evaluated by studying the rendered output of the system and noting the change in quality of simulation whilst total bandwidth use remains constant. This differs from many approaches, which seek to provide evidence for bandwidth reduction while maintaining a constant quality of simulation. In practical applications the available bandwidth is often the constant factor, and thus placing emphasis on providing scalability such that quality of simulation may be optimised for any given bandwidth should prove advantageous.

In order to further develop techniques a consideration needs to be made of the nature of quality of simulation (as discussed in section 3), and ways in which its attributes may be exploited to provide such improvement. A fundamental aspect of quality of simulation of particular interest is its close relationship to the perception of the user. If a user simply 'feels' a simulation is more accurate, then an improvement can be assumed to have been achieved. Existing interest management techniques often provide some appreciation of perception as a simple spatial attribute - if a object is within a certain distance of the user, then it is perceptually important. More refined approaches, such as that of Beharee et al., take fuller advantage of the wealth of research regarding perception, and consider its implication to interest management. The next section introduces an alternative approach to the implementation of such perceptual considerations.

5. A Perceptually-Oriented Interest Management Approach

Currently, a method is being developed which bases itself on the aura concept of interest management. However, the shape and description of this aura is modified to accommodate a series of perceptual concepts:

1) Spatial proximity alone is largely irrelevant as a measure of relevance in the situation of a client observing a distant object. Many IM approaches offer poor performance with regards to distant objects, as they focus primarily on the space around the user.

2) It is possible to provide a good approximation of a users interests by studying their interaction within the environment - specifically, how they orient their viewpoint using mouse or head-mounted input devices.

3) Such interests may shift rapidly, particularly in the case of military simulations or games.

Existing perceptually-oriented IM approaches base themselves largely around the model of attention as a multi-modal process [PJR01]. In this scenario, objects are filtered based on intrinsic low-level properties such as their colour, or higher-level simulation dependant attributes. However, a wealth of alternative models for attention exist which offer scope for continued research. In particular, the spotlight model of attention [Tre86] may be seen to offer the potential to provide additional filtering based on user viewpoint orientation. It is this model which forms the basis of the technique currently being developed to work alongside existing methods and offer further refinement.

In order to implement such a model, a system is developed wherein each client is assigned a dynamic aura, a concept not new in itself (Velvet [OG02], for example, implements such a system successfully). However, the shape of the aura is derived by attempting to determine the visual focus of the client, based upon their orientation and the surrounding geometry. This method provides a unique emphasis on the focus of the user, rather than their entire field of vision or immediate aura, and is of particular relevance to large-scale military simulations wherein the target of the user is frequently of more importance than other nearby entities.

A basic illustration of this concept is shown in Figure 3. Both the region around the user (1) and their focus (2) are considered of interest and thus assigned high level of detail (LOD) priority, although significantly greater emphasis is placed around the focus. Thus a relatively distant client (3) receives a high priority due to their position near the focus. The focus itself is determined by use of a conical pick segment, drawn from the center of the user's field of view, which ultimately intersects with either geometry of another client. In the case of a geometry intersection, the local area is analysed for nearby clients and focus shifted to the nearest. It is hoped that users will naturally learn to realign the center of the screen according to their interest (using mouse-driven input), hence making this measure of focus valid. It also integrates well with military simulations and first-person style games, which tend to place the user's weapon crosshair central on the screen.

Subsequent to the establishment of focus, a field of values representing perceptual relevance (termed P-values) is established, based upon field equations arising from the concepts discussed above and generally obeying computationally-simple $1/r^2$ behaviour. The system thus can be described as taking advantage of 'inattentional blindness' [MR98], as objects in the middle-ground generate lower P-values to those near the user or point of focus. These values may be subsequently quantised to provide support for multiple levels of detail (as illustrated in Figure 3), and thus translated into multicast group subscriptions. The generic nature of the P-values allows a separation between providing a measure of perceptual relevance and performing the information culling central to interest management, which is useful for adding additional weighting or filtering prior to data transmission. Current work is considering more closely other perceptually important aspects of behaviour, such as sudden movement, in an attempt to predict focus shifts prior to user interaction - in effect building upon existing methods which seek to predict interactions based purely on object motion [ML03]. Further research intends to consider allowing event-triggered focus shifts (such as a sound near the user) to combat 'degree of blindness' effects (scenarios where clients are not mutually visible) as noted by the Velvet system.

A further advantage of the P-value field description is an ability to downsize in areas of high network traffic not only by a reduction in the volume of each aura, but also by directly changing the thresholds for the translation between values and their corresponding multicast groups. This permits emphasis to shift from displaying large numbers of clients at low resolution or smaller numbers with increased quality without necessitating any recalculations of aura structure or definition. The approach is generally de-

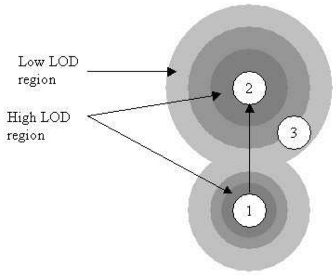

Figure 3: *Perceptual Relevance Model*

signed for scenarios where clients are viewing or interacting with large numbers of other clients in open environments, providing filtering for large scale scenarios such as virtual cities or combat simulations. Thus such techniques which offer a low processing overhead alongside scalability, and remain capable of increasing CPU demand to filter more accurately (by increasing the rate at which the aura is recalculated), are advantageous.

Preliminary implementations within the test system described in Section 4 indicate a sufficiently low CPU load to be viable; however, to accurately evaluate the success of the method a measure of quality of simulation is required (as noted previously). This is the focus of future work.

6. Summary

This paper has presented a brief review of the state-of-the-art regarding interest management in MMVEs, alongside providing a discussion of the potential for further analysis of perceptual relevance, and the shortcomings of existing perceptual techniques in a generic context. It has also provided some discussion into the criteria that need to be established for the comparative evaluation of interest management techniques, and outlined the development of a test system intended to provide facilities for such evaluation of performance is a wide range of scenarios. Finally, an approach has been described which seeks to provide both flexibility and scalability alongside a perceptually-oriented IM paradigm.

It is hoped such refinements to interest management will allow not only for more efficient use of bandwidth alongside improved quality of simulation, but will also allow for far greater extensibility, as content may be streamed at run time rather than downloaded in advance. Such environments would allow for far greater creativity and interactivity than existing environments, and subsequently offer far more potential.

References

[AWZ98] ABRAMS H., WATSEN K., ZYDA M.: Three-tiered interest management for large-scale virtual environments. In *Proceedings of the ACM symposium on Virtual reality software and technology* (1998), pp. 125 – 129.

[BWA96] BARRUS J. W., WATERS R. C., ANDERSON D. B.: Locales: supporting large multi-user virtual environments. *IEEE Computer Graphics and Applications 16*, 6 (Nov. 1996).

[BWH03] BEEHAREE A. K., WEST A. J., HUBBOLD R.: Visual attention based information culling for distributed virtual environments. In *Proceedings of the ACM symposium on Virtual reality software and technology* (2003), pp. 213 – 222.

[Cas01] CASTRANOVA E.: Virtual worlds: A first-hand account of market and society on the cyberian frontier. In *CESifo Working Paper Series* (Dec. 2001), no. 618.

[FS98] FRECON E., STENIUS M.: DIVE: a scalable network architecture for networked virtual environments. *Distributed Systems Engineering 5* (Sept. 1998), 91–100.

[Gre98] GREENHALGH C.: Awareness management in the MASSIVE system. *Distributed Systems Engineering 5*, 3 (Sept. 1998), 129–137.

[HPG02] HOSSEINI M., PETTIFER S., GEORGANAS N.: Visibility-based interest management in collaborative virtual environments. In *Proceedings of the 4th international conference on Collaborative virtual environments* (2002), pp. 143–144.

[Jac95] JACOB R. J. K.: Eye tracking in advanced interface design, virtual environments and advanced interface design. *Oxford University Press, Inc., New York* (1995).

[ML03] MORGAN G., LU F.: Predictive interest management: An approach to managing message dissemination for distributed virtual environments. In *Proceedings of the First International Workshop on Interactive Rich Media Content Production: Architectures, Technologies, Applications, Tools* (2003).

[Mor96] MORSE K. L.: *Interest Management in Large-Scale Distributed Simulations*. Tech. rep., University of California, Irvine, 1996.

[MR98] MACK A., ROCK I.: *Inattentional Blindness*. MIT Press, 1998.

[Mr02] MÁSA M., ŐÁRA J.: Generalized interest management in virtual environments. In *Proceedings of the 4th International Conference on Collaborative Virtual Environments* (2002), pp. 149–150.

[MT95] MILLER D. C., THORPE J. A.: SIMNET: The advent of simulator networking. In *Proceedings of the IEEE* (1995), pp. 83(8):1114–1123.

[OG02] OLIVEIRA J., GEORGANAS N.: Velvet: An adaptive hybrid architecture for very large virtual environments. *IEEE International Conference of Communications 1*, 2 (2002), 491–495.

[PB95] PRATT D. R., BRUTZMAN D. P.: Exploiting reality with multicast groups: a network architecture for large-scale virtual environments. In *VRAIS '95: Proceedings of the Virtual Reality Annual International Symposium (VRAIS'95)* (1995), IEEE Computer Society, pp. 38–45.

[PJR01] PASHLER H., JOHNSTONE J. C., RUTHRUFF E.: Attention and performance. *Annual Review of Psychology 52* (2001), 629–51.

[Red01] REDDY M.: Perceptually optimized 3d graphics. *IEEE Computer Graphics and Applications 21*, 5 (2001), 68–75.

[Tre86] TREISMAN A.: Features and objects in visual processing. *Scientific American 254* (1986), 114–124.

[Tse04] TSE P. U.: Mapping visual attention with change blindness: New directions for a new method. *Cognitive Science 28*, 1 (2004), 241–258.

[ZS00] ZYDA M., SINGHAL S.: *Networked Virtual Environments: Design and Implementation.* ACM Press, 2000.