EGVE Symposium (2008)
B. Mohler and R. van Liere (Editors)

Spotlight Interest Management for Distributed Virtual Environments

I. Dunwell & J.C. Whelan

Department of Computer Science, University of Hull, UK

Abstract

This paper presents a novel refinement to visual attention-based interest management in distributed virtual environments (VEs). It is suggested that in the context of a desktop VE where only limited immersion occurs, using proximity in virtual space as a primary measure of relevance may be less effective than considering the characteristics of visual interaction with the two-dimensional display. The method seeks to utilise a spotlight model of human attention in place of a proximity measure, capable of giving extremely distant clients near the centre of the display priority. In order to evaluate the technique, a series of user experiments are described which seek to study the participant's ability to detect change between techniques in a proprietary collaborative virtual environment. Two groups of users are shown to exhibit a blind preference for the spotlight method, and failed to detect a significant change when available bandwidth was reduced using this approach. The technique may be integrated alongside existing saliency-based interest management paradigms as an alternative to the distance-based factor.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Virtual Reality I.3.7 [Computer Graphics]: Three Dimensional Graphics and Realism

1. Introduction

Modern distributed virtual environment (DVE) technology can be found in many applications, ranging from entertainment and online gaming through to military simulations. The need for efficient distribution and delivery of network data in DVEs has long been recognised as essential for providing a high quality of simulation, and is a key factor in providing immersive large-scale environments [MZP*94]. The process of providing such distribution in an effective fashion is commonly termed Interest Management (IM).

A trend common to both IM and high-performance rendering research is the consideration of theories of human attention as a means to perform optimisation. Existing IM approaches, described in Section 2, commonly weight proximity to the user as the predominant factor in establishing relevance. The approach described in this paper considers the impact of the user interface common to desktop virtual environments on this assumption, in particular the ability of the user to reorient the viewpoint as an overt response to stimuli and the implications of interaction with a two-dimensional display. It seeks to exploit the simplicity of a spotlight model of human attention in order to refine the proximity-based approximation by considering also the likelihood of a focal point around the centre of the display, and subsequent deployment of covert attention. The two-step implementation seeks to initially estimate the most likely target of a users covert attention, and subsequently applies a spotlight-based measure to determine relative saliencies for objects within the virtual environment. These saliences can in turn be transcribed to network resource allocations, offering a practical means to apply the method to existing environments.

The results of a series of experiments, described in Section 4, compare weightings based on proximity against those based on the spotlight approach. A method for evaluating attention-based interest management is developed in which test subjects indirectly indicate their preference by attempting to estimate the latency of a given configuration. Subsequent comparison of the proposed method to a proximity-based approach suggests improved perceived quality in situations where a large number of users coexist in a small virtual space.

2. Background

A common approach for interest management on a large-scale involves splitting the environment into a number of discrete locales, whose form has been defined via a number of approaches ranging from the static hexagonal [MT95] grids of NPSNET, to dynamic quadtree [WZ98] or user-defined regions as in the SPLINE system [BWA96].

However, even in the presence of locale-based interest management, problems continue to emerge when a large number of users enter a single locale. Such grouping behaviour is often essential to the applications of virtual worlds, be it for purposes of military engagement, conferencing, or socialisation. In this case users are presented with a view of often many hundred other users, and a grid based approach is too coarse to perform filtering without causing visual artefacts (such as clients "popping-up" as they enter the same cell). Greenhalgh [GB95] sought to resolve this by moving the interest management process within a single locale towards a user-centric "aura" based approach, wherein each client specifies it's own focus. Extensions on this technique have included the use of dead-reckoning to predict aura intersection and thus allocate resources in advance [MLS05], and the use of multiple levels of detail [PKK00]. The VELVET system has also considered independent manipulation of focus and nimbus to provide both scalability and heterogeneity [OG03]. However, such scalability induces a degree-of-blindness problem where visibility is not mutual. To overcome this it is necessary to update all users at least on a coarse level on the locations of all other users; hence the aura-based approach may be seen as allocating levels of detail rather than proving an absolute definition of interest.

It is in this respect that the relationship between interest management and human attention begins to become apparent. Extensions of the aura-based approach have successfully considered occlusion as a filtering mechanism [HPG02], yet a wealth of research into human visual attention demonstrates that presence in the visual field alone does not represent salience, as evidenced by phenomena such as change blindness [MR76]. More recent research into visual attention has sought to view the process via feature-integration theory as described by Treisman [TG80]; which suggests humans identify objects as a collection of stimuli. In the visual case Treisman suggests the visual system builds separate feature maps when viewing a scene, theorising the brain holds these maps internally and combines them into a saliency map representing relevance. Both top-down and bottom-up characteristics of objects define their salience within this map. Beeharee [BWH03] seeks to apply this approach to the interest management process, finding key bottom-up characteristics in objects to be colour and motion, and utilising them as a means to measure relevance. The approach is evaluated by a small but non-trivial sample of users who perceived no difference between the proposed and existing techniques despite a clear reduction in bandwidth using the proposed approach.

Such success, coupled with the demonstrated effectiveness of considering perceptual techniques as a means for improving graphics performance by adjusing polygon level of detail [BCP03], [PN04] illustrates the potential for further consideration of visual attention within the interest management process.

The next section presents an approach that may be used either alongside or in lieu of a feature-integration based approach. It considers the use of a composable and computationally efficient model for approximating attention, based primarily on the spotlight model of attention as described by Eriksen and Hoffman [EH73]. It primarily questions the notion of proximity to the user in virtual space as a principal factor in salience computation, an assumption common to existing methods which seek to saliency map virtual scenes [LDC06]. Given the fact that within a desktop environment, user focus is upon a perceived three-dimensional image on a two-dimensional display, rather than a true three-dimensional scene, subsequent implications on immersion and hence perception are well-documented [Zel92]. By considering the relationship between a user's covert spotlight of attention and the display, and their overt interaction, the proposed method seeks to predict the focus of attention by predominantly considering visual interaction with the display, rather than the proximity of virtual avatars. As such, it may be interpreted as redefining the depth contribution to a saliency map, whilst remaining compatible with the consideration of other factors such as motion and contrast.

3. Spotlight Interest Management

Distributed virtual environments commonly contain both immutable content and mutable content. The primary advantage of keeping a proportion of content immutable is the ability to download such content prior to run-time, reducing run-time network overheads. Consequently immutable content often includes terrain and other complex geometry, whilst mutable content includes avatars and interactive objects. In order to maintain a consistent distributed simulation, only the mutable content need be updated via the network. For clarity, this paper refers to all mutable objects (including avatars) as *entities*.

By considering only entities, the approach minimises unnecessary processing. It should be noted, however, that as environments increase in sophistication, the demand to support a greater volume of mutable content emerges. To faciliate the development and analysis of the technique with a large number of entities, a virtual environment is created which represents users as coloured boxes moving over a region of undulating terrain. By simulating additional users it becomes possible to populate the environment with several hundred entities simultaneously. The underlying architecture is a peer-to-peer multicast based approach, although the saliency measure itself can be transcribed to any architecture. A typical desktop-based interaction model using the mouse and keyboard is implemented, as described by Hand [HAN97] as common to many commercial environments. The system operates by interpreting reorientation of the viewpoint as "overt" shifts in attention (analogous to moving the head or eyes), whilst seeking to anticipate the deployment of "covert" attention - commonly descibed at the mental image of the scene - using a spotlight metaphor.

The approach is a two-step process; firstly the orientation of the viewpoint and relative positions of avatars are examined in order to estimate the most likely position of the covert focus, and then a saliency field is created which allows each networked entity to be assigned a relevance value. These relevance values are subsequently translated to network resource allocation - in the case of the peer-to-peer virtual environment created for testing, these resource allocations take the form of multicast group assignments.

3.1. Determining Focus

To estimate covert focus, the assumption is made that the user is predominantly fixated around the centre of the screen. Considering position alone, the relative probability value of an entity at p_0 being the focus may be approximated as a summation of two terms, one representing the distance between user and object (the traditional distance-saliency factor), and another representing the offset of the entity from the centre of the display:

$$p_f = a|(p_1 - p_0)| + b\left(\frac{|(p_2 - p_1) \times (p_1 - p_0)|}{|(p_2 - p_1)|}\right) \quad (1)$$

Where the point p_1 is the position of the user, and p_2 the point at which a line traced from the centre of their field of view into the screen intersects the clip plane. Equation 1 introduces coeffcients for both the distance and offset weightings, defined as a and b respectively. These provide weighting between the two factors influencing likelihood of focus; in the case $a >> b$ the predominant factor is proximity to the user. However in the case $b >> a$, more distant entities generate a higher relevance provided they have a close proximity to the centre of the visual display. Figure 1 illustrates the distribution of probabilities for the $b >> a$ case. Observe that from a top-down view the probability distribution can be seen as a cone pointing away from the user, tending towards a cylinder as a tends to zero.

Figure 2 demonstrates the effect of adjusting these variables in practice. By altering the a/b ratio it becomes possible to control whether focus is predominantly assigned to entities which are nearby, or close to the centre of the screen. It may be implied that, since many existing methods consider proximity to the user as a primary measure of relevance, the $a > b$ case represents a classic aura-based approach such as that of Greenhalgh [GB95], whilst $b > a$ emphasises the spotlight as a measure of attention. Hence rather than conduct experiments to evaluate the effectiveness of various a/b ratios, a more direct evaluation of the technique as a whole may be conducted by using extreme values which seek to place extremely high emphasis on the centre of the screen or nearby entities. This is discussed further in Section 6.

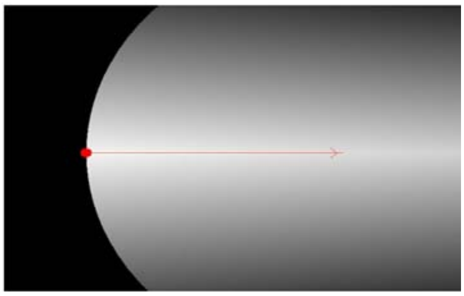

Figure 1: *Top-down two-dimensional slice of the resulting field of focal selection probabilties. The users position is illustrated as a dot with view direction along the arrow for the case b>a. Lighter regions show higher probabilities.*

To refine the approach further, and utilise the spotlight model of attention more rigorously, the next section goes on to describe the second stage of the process, which seeks to generate saliency values for other entities based on their proximity to the estimated target of attention on the display, and relative distances.

3.2. Applying a Spotlight

With a most likely object of focus established, it becomes possible to apply a "zoom-lens" type spotlight as described by Eriksen and Hoffman [EH73]. This is done by means of a radial-field type approach with components around both the target of focus, and the user. This exploits the three-dimensional nature of the scene - if the target entity is nearby, then the radial field implies a broad spotlight, conversely if the target entity is extremely distant, the field narrows when transcribed to the two-dimensional display. The component around the user is a practical consideration, applied to allow for entities extremely close yet outside the field of view to be considered higher relevance. This eliminates the chance of them undergoing excessive saliency (and thus resource allocation) shifts during rapid view reorientations.

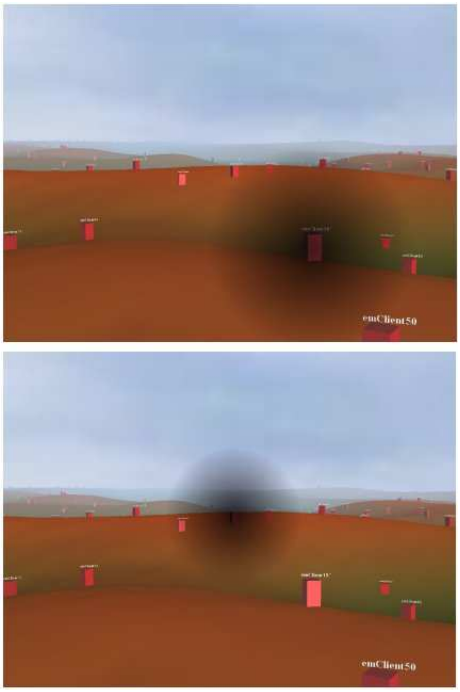

Figure 2: *Effects of a and b coefficients on focal selection. Top: $a = 5b$ Bottom: $a = b/5$*

Hence, we can describe the saliency s_p of an object at point p, relative to an observer at point p_{obs} as:

$$s_p = \frac{R_{obj}}{|(p_{obs} - p)|} + \frac{R_{obj}R_{focus}}{|(p_{focus} - p)|} \tag{2}$$

It also becomes possible to incorporate other measures of potential relevance such as colour and motion of other entities, by assuming they are pre-calculated into relevance values R_{obj} and measured against the relevance of the object selected as the focus, R_{focus}. However, for the purposes of directly evaluating the spotlight in a composable form (i.e. without requiring detailed entity characteristics to be defined), we consider an environment where position is the only variable, and hence $R_{obj} = R_{focus} = 1$. Figure 3 illustrates the resulting field of saliency values.

The effect of this approach is illustrated in Figure 4. The inset region showing the centre-right field of view shows how distant (yet central) entity A is afforded high saliency, whilst distant and offset entity B is granted lower relevance.

The emergent saliency values may be translated to both hardware and software resource allocation. A deliberate attempt is made to permit a clear separation between the saliency values, and the allocation of network resources. By providing such separation, the technique may be more easily analysed with respect to its approximation of visual attention, and also more easily compared to existing systems.

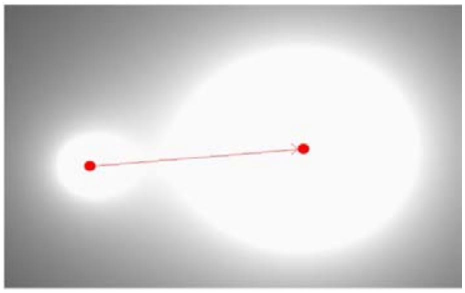

Figure 3: *Top-down illustration of the spotlight-based saliency field around client $left$ and target entity right. Note the proximity (aura around client) and spotlight-based (aura around target) components.*

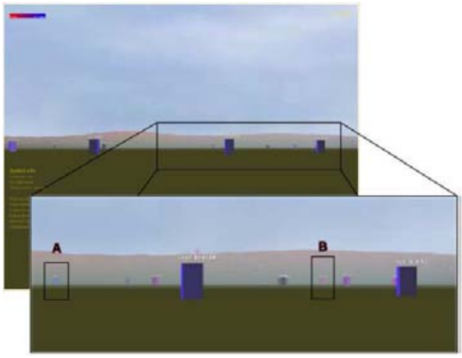

Figure 4: *Relevance weighting for distant entities, illustrated by colouration. High relevance entities are shaded blue through to red as saliency decreases (scale shown at top left)*

In the specific implementation used, a multicast architecture allows for saliency values to be transcribed to multicast assignments to groups at various levels of resolution. Each client subscribes firstly to an extremely low resolution global group (which would represent a single locale in a larger environment). This eliminates the degree-of-blindness problem mentioned in Section 2. The saliency measure is then used to invoke group subscriptions at higher resolution for the most salient clients. Figure 5 shows a simple example of this assignment for a line of clients; the highest resolution is afforded to those clients central to the field of view, whilst offset clients subscribe to increasingly lower-resolution groups.

The following and final sections of this paper discuss the method used to evaluate the technique, and subsequent results.

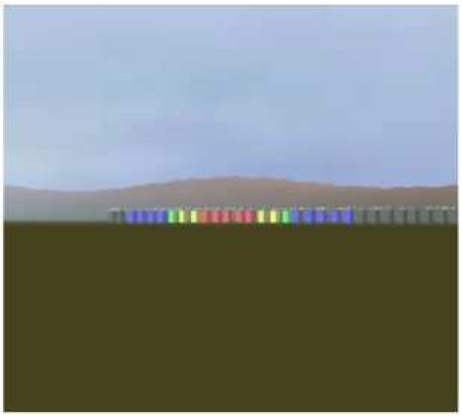

Figure 5: *Region of interest and multicast assigment. Highest resolution group is shaded red, with increasingly lower groups shaded orange, yellow, green, blue and black.*

4. Experiment Design

The purpose of the experiments conducted was principally to explore the effectiveness of a spotlight-based approach in comparison to a proximity-based measure of relevance. An open region of terrain was used for the experiments, which represents a "worst case" scenario for interest management, in which little visual occlusion occurs due to terrain and hence users are typically viewing a large number of other clients simultaneously. Hence in this situation interest management is most critical in providing continued quality of simulation.

The assumption common to existing methods that the proposed approach disputes is that proximity to the user in virtual space is a principal measure of relevance. Hence the experiments were designed in order to examine this factor alone. In order to investigate this dispute, a proximity-based technique is evaluated alongside the spotlight-based approach. In an attempt to provide quantifiable results, a series of experiments are conducted with groups of simultaneous users (shown in Figure 6), and 200 simulated clients exhibiting random motion.

Existing rendering or IM approaches aiming to examine perceptually-based refinements employ a wide range of techniques for their evaluation, as a result of the innate complexities of the visual attention process. Beharee [BWH03] uses a detection of change approach, where users participate with simulations using existing and proposed approaches. However, in this case, the responses of users will likely be influenced by many factors, such as their location and actions in the virtual world, and may not directly relate to the interest management performance. An alternative technique such as comparing predicted results with eye-tracking data (such as that of Parkhurst and Niebur [PN04]) can offer good data regarding the effectiveness of perceptually-optimised rendering, which could potentially be applied to IM. The key difference in the IM case, however, is that the cost of reassigning priorities to entities (e.g. multicast subscriptions), coupled with the high frequency of visual saccades, means simply detecting a single object of focus is inadequate; rather, the effectiveness of the method at predicting multiple targets must be considered.

The approach used to evaluate the technique reflected on these limitations. Rather than ask users to detect change on a yes/no basis, an experiment was devised wherein users were cued to quantify their preference for a technique numerically, by indicating what they felt the latency (interval between position updates) for other avatars a given simulation was. Subsequent analysis of variance (ANOVA) of the results can then offer evidence regarding whether change was detected.

It was necessary to first provide subjects with an illustration of the visible effects of increased latency in order to clarify their task; this was achieved using a large-screen display showing a demonstration of a large group of avatars moving with steadily increasing latencies. The subjects were then asked to perform latency estimations whilst performing two sets of tasks - firstly a grouping activity, in which they sought to seek out the other real-world users and form a small group within the environment, and secondly a tag game, in which a randomly chosen client had to chase the other real-world clients through the crowd. These activities were chosen to promote both rapid grouping and dispersal of clients, in order to stress the interest management process as much as possible. To test the capability of users to perform such evaluation, two controlled extremes were used: firstly a simulation limited only by local-area network performance, and secondly a simulation of an Internet-based scenario with limited bandwidth. In both cases no interest management was applied. Additional experiments were introduced that retained the Internet-based simulation whilst using spotlight and proximity-based interest management, and also a series which reduced the available bandwidth whilst using a spotlight-based approach.

To reduce potential bias arising from the sequence of the experiments, the order in which interest management approaches were applied was randomised between the grouping and tag activities. The experiment was also repeated in a different sequence with a second group of users.

5. Results and Discussion

Figure 7 shows the results obtained for the sequence of experiments described in Section 4. Subjects demonstrated a clear ability to distinguish accurately between the two extreme cases reinforcing the validity of the experiment. With interest management applied, the downstream packet rate (and hence bandwith) was restricted to a total of 80 packets/second across all multicast groups, effectively only al-

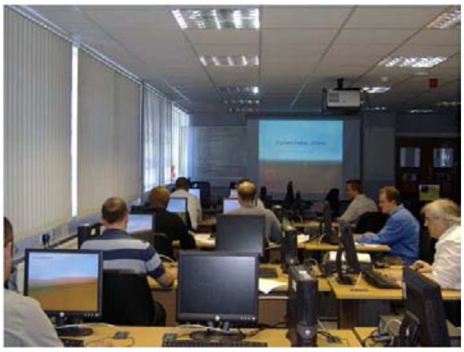

Figure 6: *Users participating in an experiment*

lowing a single entity an optimum update rate of 40 packets/second to match the framerate. Entities measured at lower saliency are given steadily decreasing update rates; in the case of the user viewing 20 clients, this means the ten entities deemed least salient are updated with the lowest resolution of 0.5 packets/second.

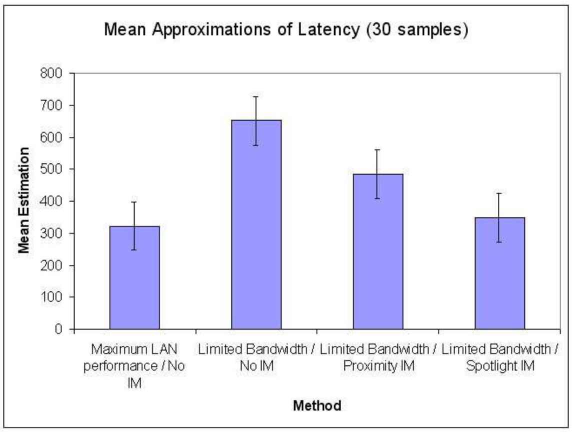

Figure 7: *Graph showing the mean estimations of subjects for a series of experiments. Error bars show standard error.*

A data set was created by combining the results for grouping and tag experiments, and merging the two iterations of the experiment with different users. Statistical analysis of this data set offers some indication of whether change was detected, or the difference in the means is by chance alone. The variance in expressed values was substantial (note the standard error illustrated in Figure 7); though it is important to reiterate that the experiment purely used latency estimation as a means for users to express their preference rather than as an absolute measure. A one-tailed t-test (alpha = 0.05) of the total observations by subjects between spotlight and proximity-based techniques suggests a significant difference was detected in favour of the spotlight approach ($t = 2.18$, $t_{crit} = 1.67$). Coupled with the evidence that all subjects correctly identified the best and worst-case scenarios (LAN performance and limited bandwidth without IM), these results offer some support to the hypothesis that the spotlight approach is a better measure of relevance than proximity.

An additional ANOVA considering task, method and observer as factors suggested highly significant effects of observer of the result (i.e. individual results were relative rather than absolute estimations, as would be expected), whilst task had much less impact. Correcting the ANOVA to take the individual differences into account (dividing each observation by the mean of the observer's estimates) still produces evidence that a change is detected in favour of the spotlight technique at 0.05 alpha.

An additional series of experiments considered the effects of reduction of bandwidth under the proposed technique; an ANOVA of the three experiments (at 100%, 80% (64 packets/sec) and 40% bandwidth (32 packets/sec)) suggests a significant difference was *not* detected at 0.05 alpha ($F = 0.66$ $F_{crit} = 3.10$) despite the significant reduction in network usage. Though the limited sample size and limited activities (grouping and tag) prevents any absolute conclusions from this data, this suggests the technique was substantially effective at masking bandwidth reduction from the subjects.

A final consideration, shown in Figure 8, is the effect of task on effectiveness. In this case, a significant difference was noticed only for the case of no IM. This reinforces the assumption that effective IM is crucial for providing a more consistent experience in a distributed VE, though further analysis would be required to reach conclusions regarding the impact of task on the different techniques.

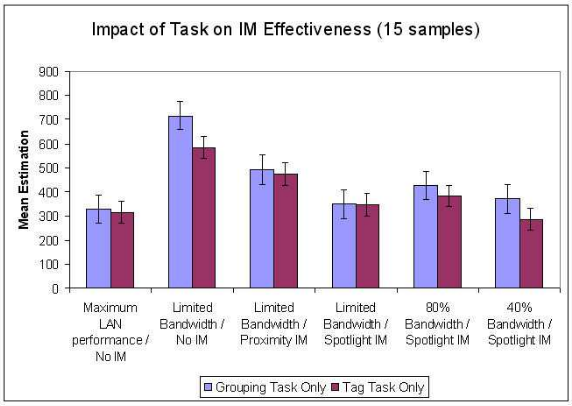

Figure 8: *Graph showing the difference in mean estimation between tasks*

.

It is important to note the limitations in the experiment when attempting to reach conclusions. The foremost limitation is the sample size; although the t-test suggests that significant difference was observed. A subsequent question is whether bias arose from the sequence of the experiments and values being expressed relative to the previous experiment - randomisation of the methods both between activities and the two user groups aimed to reduce any such influence.

Though the processing overheads for the spotlight approach are minimal, a drawback of the technique is that saliency shifts can happen frequently in crowded areas when the viewpoint is quickly reoriented. Figure 8 shows a comparison in the consequences for multicast subscriptions for a region with 200 simulated clients and a user rotating about the spot. Although during the experiments conducted this resulted in no apparent performance degradation, it should be considered as an additional overhead. Considering a means to reduce this side-effect, such as hysteresis, offers scope for future work.

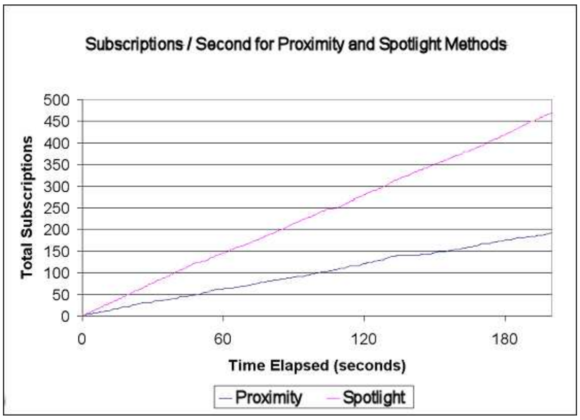

Figure 9: *Multicast subscription rate comparison*

6. Conclusions

The approach described in Section 4 illustrates how a spotlight-based filter may be applied with a minimum of iteration and scene analysis, and thus infers little impact on overall system performance. Many factors in existing interest management approaches, such as the consideration of bottom-up characteristics described by Beeharee [BWH03] may be integrated alongside the spotlight-based technique by replacing only the distance-saliency factor, and hence the approach may be viewed as complimentary to many existing techniques, save those which seek to define proximity as an absolute measure of relevance.

A predominant requirement of modern virtual environment technology is composability, as defined by Singhal and Zyda [SZ99]. Detailed analysis on a per-object level goes against the concept of composability; if, in order to perform property-based filtering we need detailed information on object properties stored in advance, we restrict the virtual environment to content which is defined in such a detailed format, or perform computationally-expensive feature extraction analysis at run-time. Whilst Beeharee represents avatars as coloured blocks, establishing colour or motion-based saliency for a more realistic, animated avatar rapidly becomes a complex problem. The spotlight approach offers some further advantages in this context, since it operates on a per-entity rather than per-pixel basis and only requires information on the fundamental attributes of position and orientation.

The system places constraints on user interface devices, since a relationship exists between the ability of the interest management to detect covert attention shifts based upon overt shifts carried out by the user. A spotlight model would be expected to perform less successfully in situations where view shifts are impossible (e.g. observers watching a demonstration), since in the absence of the ability to deploy overt attention through interacting with the environment, observers would be expected to shift their focus around the display far more frequently. It is also worth noting that in the case of more immersive display and interaction technology, proximity would be expected to increase in significance. Considering the method using such technology offers a potential for further studies.

This paper has presented a means for applying a spotlight approximation of human attention to the interest management process. The suitability of the model to a computationally-efficient and composable implementation leads to a model which may be easily integrated alongside existing approaches. Experimental results provide some evidence of the effectiveness of the technique, and, alongside the success of other attention-based approaches, affirm the considerable potential for future work. Such work, aimed at converging theories regarding visual attention with those of interest management and rendering, promises to offer richer, more compelling DVE experiences.

7. Acknowledgements

The authors would like to thank Dr. Craig Gaskell and Derek Wills for their input throughout the research, two anonymous reviewers and Dr. Philip Wheeler for their helpful suggestions regarding the analysis of results, and finally the staff and students of the University of Hull for giving up their time to participate in the experiments.

References

[BCP03] BROWN R., COOPER L., PHAM B.: Visual attention-based polygon level of detail management. In *GRAPHITE '03: Proceedings of the 1st international conference on Computer graphics and interactive techniques* *in Australasia and South East Asia* (New York, NY, USA, 2003), ACM, pp. 55–ff.

[BWA96] BARRUS J. W., WATERS R. C., ANDERSON D. B.: Locales: Supporting large multiuser virtual environments. *IEEE Computer Graphics and Applications 16*, 6 (1996), 50–57.

[BWH03] BEEHAREE A. K., WEST A. J., HUBBOLD R.: Visual attention based information culling for distributed virtual environments. In *VRST '03: Proceedings of the ACM symposium on Virtual reality software and technology* (New York, NY, USA, 2003), ACM, pp. 213–222.

[EH73] ERIKSEN C., HOFFMAN J.: The extent of processing of noise elements during selective encoding from visual displays. *Perception & Psychophysics 14*, 1 (1973), 155–160.

[GB95] GREENHALGH C., BENFORD S.: MASSIVE: A distributed virtual reality system incorporating spatial trading. In *International Conference on Distributed Computing Systems* (1995), pp. 27–34.

[HAN97] HAND C.: A survey of 3d interaction techniques. *Computer Graphics Forum 16*, 5 (1997), 269–281.

[HPG02] HOSSEINI M., PETTIFER S., GEORGANAS N. D.: Visibility-based interest management in collaborative virtual environments. In *CVE '02: Proceedings of the 4th international conference on Collaborative virtual environments* (New York, NY, USA, 2002), ACM, pp. 143–144.

[LDC06] LONGHURST P., DEBATTISTA K., CHALMERS A.: A gpu based saliency map for high-fidelity selective rendering. In *Afrigaph '06: Proceedings of the 4th international conference on Computer graphics, virtual reality, visualisation and interaction in Africa* (New York, NY, USA, 2006), ACM, pp. 21–29.

[MLS05] MORGAN G., LU F., STOREY K.: Interest management middleware for networked games. In *I3D '05: Proceedings of the 2005 symposium on Interactive 3D graphics and games* (New York, NY, USA, 2005), ACM, pp. 57–64.

[MR76] MCCONKIE G., RAYNER K.: *Identifying the span of the effective stimulus in reading: Literature review and theories of reading*, 2 ed. Newark, DE: International Reading Association, 1976.

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

[MZP*94] MACEDONIA M. R., ZYDA M. J., PRATT D. R., BARHAM P. T., ZESWITZ S.: Npsnet: A network software architecture for large-scale virtual environment. *Presence 3*, 4 (1994), 265–287.

[OG03] OLIVEIRA J. C. D., GEORGANAS N. D.: Velvet: An adaptive hybrid architecture for very large virtual environments. *Presence: Teleoperators & Virtual Environments 12*, 6 (2003), 555–580.

[PKK00] PARK C., KO H., KIM T.: Multi-resolution spatial model for large-scale virtual environment. In *VRST '00: Proceedings of the ACM symposium on Virtual reality software and technology* (New York, NY, USA, 2000), ACM, pp. 92–96.

[PN04] PARKHURST D., NIEBUR E.: A feasibility test for perceptually adaptive level of detail rendering on desktop systems. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization* (New York, NY, USA, 2004), ACM, pp. 49–56.

[SZ99] SINGHAL S., ZYDA M.: *Networked virtual environments: design and implementation*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1999.

[TG80] TREISMAN A., GELADE G.: A feature-integration theory of attention. *Cognitive Psychology 12*, 1 (1980), 97–136.

[WZ98] WATSEN K., ZYDA M.: Bamboo - a portable system for dynamically extensible, real-time, networked, virtual environments. In *VRAIS '98: Proceedings of the Virtual Reality Annual International Symposium* (Washington, DC, USA, 1998), IEEE Computer Society, p. 252.

[Zel92] ZELTZER D.: Autonomy, interaction, and presence. *Presence: Teleoper. Virtual Environ. 1*, 1 (1992), 127–132.