Visibility-based Interest Management in Collaborative Virtual Environments

Mojtaba Hosseini
DISCOVER Lab, University of Ottawa
School of Information Technology and Engineering, Ottawa, Canada
1(613) 5625800-6248
mojtaba@discover.uottawa.ca

Steve Pettifer
Advanced Interface Group
University of Manchester
Manchester, UK
44(161) 2756259
srp@cs.man.ac.uk

Nicolas D. Georganas
DISCOVER Lab, University of Ottawa
School of Information Technology and Engineering, Ottawa, Canada
1(613) 5625800-6248
georganas@discover.uottawa.ca

ABSTRACT
One of the challenges facing developers of large-scale and content-rich Collaborative Virtual Environments is the lack of bandwidth to support the exchange of information between participants. It is usually not feasible for all participants to receive all the data produced by all other participants. It is therefore necessary to filter out some of such data for every user according to their interest. This paper presents an interest management mechanism based on each user's visibility of others. Its main features include occlusion awareness and visibility-based filtering, no need for a server and renderer-based visibility calculation.

Categories and Subject Descriptors
H.5.1 [**Multimedia Information Systems**]: Artificial, augmented, virtual realities.

General Terms
Management, Measurement, Performance, Design.

Keywords
CVE, AoIM

1. INTRODUCTION
Whether it is the number of users in Collaborative Virtual Environments (CVE) or the amount of data each user generates in a CVE, increase in either or both can quickly limit the bandwidth available for the application. It is usually not feasible for all participants to receive all the data produced by all other participants. It is towards the aim of filtering data outside the interest of a given user, that many Area of Interest Management (AoIM) techniques have been proposed and implemented.

This paper presents a novel approach to interest management based on the results of visibility culling calculations performed by each user's rendering mechanism such that each user will primarily receive updates regarding only users that are visible.

ACM CVE '02, September 30-October 2, 2002, Bonn, Germany.

2. VISIBILITY AS A MEASURE OF INTEREST
In order to filter messages in a CVE based on visibility of users, we distinguish between the different categories of messages generated by each user as: Position data (those that contain the present position, not including orientation, of the user in 3D space), Visual-type data (those that contain information that a user would be interested in knowing and receiving only if its source is visible, such as the video of a textured avatar) and Non-visual type data (those that contain information that a user would be interested in knowing and receiving even if the source is not visible, such as a user's voice).

Provided that the position data and visual-type data are transmitted over separate channels, we propose to use visibility as a measure of interest and filter visual-type data based on the received position data. That is to say, if the location of participants is available, then each participant can subscribe to the visual-type data of only those users that are visible. Using visibility as a measure of awareness has already been exploited in other systems [2][3][4]. The novelty in our approach is how, or more precisely where, visibility calculations take place.

Unlike the DVE systems cited, we propose that each participant should calculate their own visibility of other participants (instead of a server or other central interest manager). An interesting outcome of this proposal is that there is actually no need to perform such calculations at all. This is because visibility culling of objects in the scene is already done in real-time by the renderer of each participant. If the results of visibility culling performed by the renderer are available to the application (irrespective of what algorithm is used by the renderer to determine visibility), the need to duplicate them for interest management vanishes. Therefore, not only visibility calculations are done in a completely distributed manner (with no need for servers), but also there is no added burden of performing such calculations at all.

3. VISIBILITY-BASED AOIM
This section distinguishes our work from other DVE interest management schemes as well as highlights the advantages and features of our proposal.

1. *Locale/region level interest management:* In contrast to spatial partitioning techniques (NPSNET [7], DIVE [6], SPLINE [1]), our approach is not concerned with filtering user position data at the world level but strives to filter out non-position

data at the locale/region/cell level. In fact, the proposed approach augments the spatial partitioning techniques. While the former focus on limiting the awareness of each user to those users that are spatially close, the latter strives to further reduce the data exchange between spatially close users. This is appropriate for applications where each user generates enough traffic that receiving it from every user in even one spatial unit (Locale/region) requires bandwidth above what is available. CVEs employing video textured avatars and/or articulated avatars controlled by various body tracking input devices are examples of such applications. The main difference between our approach and other Locale-level interest management schemes such as VELVET [8] and MASSIVE [5] are points 2 and 4 below.

2. *Visibility based interest management:* Unlike MASSIVE, our approach can provide flexible terrain-aware area of interest calculation since visibility determines interest. It is flexible because the 'area of interest' of each user is calculated dynamically in real time by the renderer as opposed to the more static definition of aura/nimbus/focus. Our approach is terrain-aware since visibility of users implicitly takes into account the environment or terrain and thus reacts appropriately to the occlusion of users by the environment/terrain or other users.
3. *Server-less interest management:* With regards to other networked virtual environments using visibility as a measure of interest such as RING [4], DEVA [9] and others [2][3] as well as other Locale-level awareness management schemes such as MASSIVE, the work presented here differs in that a server is not needed to calculate interest collision, thus eliminating a need for a potential bottleneck and a single point of failure. Instead each client performs their own interest management calculations at the Locale/region level and subscribes to the relevant sources of data.
4. *Renderer based visibility calculation:* since visibility based interest management is performed on each participant's machine (as opposed to a server), and since every participant's renderer performs visibility calculations in real-time for rendering purposes, the results of these calculations can be used to determine the corresponding user's interest. In other words, not only are the interest calculations completely distributed, but they are not necessary at all.

It should be mentioned that the implementation of our proposal requires the rendering mechanism used for the CVE application to make the results of its visibility calculations available to that application. Figure 1 illustrates a simplified design for the implementation of the proposed interest management scheme.

4. CONCLUSION AND FUTURE WORK

After a suitable interest management scheme has limited the communication of CVE participants to users in a region or Locale, the further filtering of data other than position may be necessary if every participant generates significant amount of data. This paper presented a novel client-side visibility-based interest management scheme at the Locale level that does not require a server or any awareness collision calculations and is also terrain and occlusion aware. Future work includes investigating the possible advantages of our proposal (terrain and occlusion awareness) as well as addressing potential drawbacks through its implementation and testing.

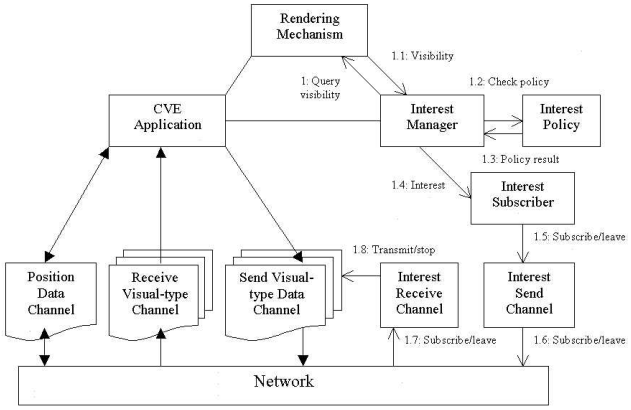

Figure 1 Systems design for visibility-based AoIM

5. REFERENCES

[1] Barrus, J.W., Waters, R.C., Anderson, D.B. Locales and beacons: efficient and precise support for large multi-user virtual environments. Proceedings of 1996 IEEE Virtual Reality Annual International Symposium (VRAIS'96), pp. 204-213.

[2] Beeharee, A., Pettifer, S., West, A. A framework for optimising network usage for plausible distributed virtual environments. Proceedings of New Trends in Cognitive Science, 2001.

[3] Capin, T.K., Pandzic, I.S., Magnenat-Thalmann N., Thalmann, D. Networking Data for Virtual Humans. Avatars in Networked Virtual Environments, Wiley, 1999.

[4] Funkhouser, T.A. RING: A client-server system for multi-user virtual environments. Proceedings of Computer Graphics (1995 SIGGRAPH Symposium on Interactive 3D Graphics), pp. 85-92.

[5] Greenhalgh, C., Benford S. MASSIVE: a Distributed Virtual Reality System Incorporating Spatial Trading. Proceedings of 1995 International Conference on Distributed Computing Systems (DCS'95), pp. 27-34.

[6] Hagsand, O. DIVE: Platform for multi-user virtual environments. IEEE Multimedia 3 (1) (1996) 30-39.

[7] Macedonia, M., Pratt, D., Zyda, M. NPSNET: Network software architecture for large scale virtual environments. Presence 3 (4) (1994) 265-287.

[8] Oliveira, J.C., Georganas, N.D. VELVET: An adaptive hybrid architecture for very large virtual environments. Proceedings of 2002 IEEE International Communication Conference (ICC'02).

[9] Pettifer, S., Cook, J., Marsh, J., West, A. Deva3: Architecture for a large scale virtual reality system. In Proceedings of ACM Symposium in Virtual Reality Software and Technology 2000, pages 33-39.